A NEW L-CURVE FOR ILL-POSED PROBLEMS

LOTHAR REICHEL * AND HASSANE SADOK †

Dedicated to Claude Brezinski.

Abstract. The truncated singular value decomposition is a popular method for the solution of linear ill-posed problems. The method requires the choice of a truncation index, which affects the quality of the computed approximate solution. This paper proposes that an L-curve, which is determined by how well the given data (right-hand side) can be approximated by a linear combination of the first (few) left singular vectors (or functions), be used as an aid for determining the truncation index.

Key words. Ill-posed problem, truncated singular value decomposition, truncation criterion, L-curve.

1. Introduction. Let $A : \mathcal{X} \to \mathcal{Y}$ be a compact linear operator between the Hilbert spaces $\mathcal{X}$ and $\mathcal{Y}$. Let $\langle \cdot, \cdot \rangle$ denote the inner products in $\mathcal{X}$ and $\mathcal{Y}$, and let $\| \cdot \|$ denote the associated norms, i.e., $\|f\| = \langle f, f \rangle^{1/2}$ for $f \in \mathcal{X}$ and $f \in \mathcal{Y}$. We are concerned with the computation of the minimal-norm solution $\hat{x} \in \mathcal{X}$ of equations of the form

$$Ax = \hat{y} \tag{1.1}$$

by the singular value decomposition. The right-hand side $\hat{y} \in \mathcal{Y}$ is assumed to live in the range of A, denoted by $\mathcal{R}(A)$. Then the equation (1.1) is well defined and has a unique solution of minimal norm. We will assume that $\dim(\mathcal{R}(A)) = \infty$.

Introduce the singular triplets $\{\sigma_j, u_j, v_j\}_{j=1}^{\infty}$ consisting of the positive singular values σ_j of A and the associated left and right singular functions $u_j \in \mathcal{Y}$ and $v_j \in \mathcal{X}$, respectively. The singular values are enumerated according to $\sigma_1 \geq \sigma_2 \geq \sigma_3 \geq \ldots > 0$. The singular values and functions satisfy

$$\langle u_j, u_\ell \rangle = \langle v_j, v_\ell \rangle = \begin{cases} 1, & j = \ell, \\ 0, & j \neq \ell, \end{cases}$$

and

$$\begin{aligned} Av_j &= \sigma_j u_j, \qquad A^* u_j = \sigma_j v_j, \qquad j = 1, 2, 3, \ldots \ , \\ Ax &= \sum_{j=1}^{\infty} \sigma_j \langle x, v_j \rangle u_j, \qquad \forall x \in \mathcal{X}, \\ A^* y &= \sum_{j=1}^{\infty} \sigma_j \langle y, u_j \rangle v_j, \qquad \forall y \in \mathcal{Y}, \end{aligned}$$

where $A^* : \mathcal{Y} \to \mathcal{X}$ denotes the adjoint of A. It follows from the compactness of A that the singular values cluster at zero.

*Department of Mathematical Sciences, Kent State University, Kent, OH 44242, USA. E-mail: reichel@math.kent.edu. Research supported in part by an OBR Research Challenge Grant.

†Laboratoire de Mathématiques Pures et Appliquées, Université du Littoral, Centre Universtaire de la Mi-Voix, Bâtiment H. Poincarré, 50 Rue F. Buisson, BP 699, 62228 Calais Cedex, France. E-mail: Hassane.Sadok@lmpa.univ-littoral.fr.

The minimal-norm solution of (1.1) can be expressed as

$$\hat{x} = \sum_{j=1}^{\infty} \frac{\langle \hat{y}, u_j \rangle}{\sigma_j} v_j;$$

in particular, $\hat{x} \in \mathcal{X}$ implies that the Picard condition

$$\sum_{j=1}^{\infty} \frac{|\langle \hat{y}, u_j \rangle|^2}{\sigma_j^2} < \infty \tag{1.2}$$

is satisfied. Due to the clustering of the σ_j at zero, the Fourier coefficients

$$\hat{c}_j = \langle \hat{y}, u_j \rangle, \qquad j = 1, 2, 3, \dots , \tag{1.3}$$

of the right-hand side $\hat{y}$ have to converge to zero sufficiently rapidly in order for the inequality (1.2) to hold.

The right-hand side of ill-posed problems that arise in applications in science and engineering often is obtained through measurement and typically is contaminated by measurement error. Let y^δ denote the available approximation of the error-free right-hand side $\hat{y}$. We would like to determine an approximation of the solution $\hat{x}$ of (1.1) by determining an approximate solution of

$$Ax = y^\delta. \tag{1.4}$$

The error $y^\delta - \hat{y}$ in y^δ will sometimes be referred to as "noise." Due to this error, y^δ might not be in $\mathcal{R}(A)$; in fact, y^δ is not even guaranteed to live in $\mathcal{Y}$.

The Truncated Singular Value Decomposition (TSVD) determines an approximation x_k^δ of $\hat{x}$ of the form

$$x_k^\delta = \sum_{j=1}^{k} \frac{\langle y^\delta, u_j \rangle}{\sigma_j} v_j. \tag{1.5}$$

The difference $x_k^\delta - \hat{x}$ depends on the truncation index k; generally $\|x_k^\delta - \hat{x}\|$ decreases as k increases and is fairly small, but increases rapidly with k when k is large. Let $\hat{k} \geq 1$ denote the smallest index, such that

$$\|x_{\hat{k}}^\delta - \hat{x}\| = \min_{k \geq 1} \|x_k^\delta - \hat{x}\|. \tag{1.6}$$

When a reliable estimate of the norm of the error

$$\delta = \|y^\delta - y\| \tag{1.7}$$

is known, a suitable value of the truncation index k often can be determined by the *discrepancy principle*, which selects the smallest index $k = k_\delta$, such that $\|Ax_{k_\delta}^\delta - y^\delta\| \leq \tau\delta$. Here $\tau > 1$ is a constant independent of δ. Note that $\|Ax_k^\delta - y^\delta\|$ decreases monotonically as k increases, and therefore k_δ increases monotonically as δ decreases to zero. It can be shown that $x_{k^\delta}^\delta \to \hat{x}$ as $\delta \searrow 0$; see Engl [8, Section 4.3] and Hansen [11, Section 7.2] for discussions on properties and applications of the discrepancy principle. Here we only remark for future reference that the index k_δ does not explicitly depend on $\|x_k^\delta\|$.

In many applications an estimate (1.7) of the norm of the error in y^δ is not available. Then the discrepancy principle cannot be applied to determine a suitable value of the truncation index k. The present paper describes a new L-curve, which can be used to choose a suitable value of the truncation index in this situation. This L-curve is determined by the Fourier coefficients

$$c_j^\delta = \langle y^\delta, u_j \rangle, \qquad j = 1, 2, 3, \dots \ , \tag{1.8}$$

of the contaminated right-hand side. The L-curve of the present paper seeks to assess the error in the right-hand side y^δ by measuring how well y^δ can be approximated by the first k left singular functions $u_1, u_2, \dots, u_k$. Since this L-curve depends on the norm of the residual error

$$\|y^\delta - \sum_{j=1}^{k} c_j^\delta u_j\| = \min_{\alpha_j} \|y^\delta - \sum_{j=1}^{k} \alpha_j u_j\|,$$

we refer to it as the *residual L-curve*. This L-curve differs from the L-curves considered by Lawson and Hanson [14], Hansen [9], and Hansen and O'Leary [13], which also take the size of the approximate solutions x_k^δ into account; see Hansen [11] for insightful discussions on the latter kind of of L-curves.

This paper is organized as follows. Section 2 describes the residual L-curve and discusses an application to the approximation of a smooth function, which has been contaminated by noise. The L-curve analyzed and advocated by Hansen [11] and Hansen and O'Leary [13] is also described. Section 3 discusses the solution of ill-posed problems. Computed examples that illustrate the application of the residual L-curve to the determination of a suitable truncation index k for the TSVD-solution (1.5) are presented, and comparisons with the L-curve described in [11, 13] are provided. We use methods developed by Castellanos et al. [7] and Hansen et al. [12] for determining the "vertex" of a discrete L-curve. Section 4 contains concluding remarks.

2. The residual L-curve. Define the residual errors

$$r_k^\delta = y^\delta - \sum_{j=1}^{k} c_j^\delta u_j, \qquad k = 1, 2, 3, \dots \ ,$$

where the Fourier coefficients c_j^δ are defined by (1.8). By the orthonormality of the u_j, we have

$$\|r_k^\delta\|^2 = \|y^\delta\|^2 - \sum_{j=1}^{k} |c_j^\delta|^2. \tag{2.1}$$

Thus, $\|r_k^\delta\| \ll \|r_{k-1}^\delta\|$ implies that the Fourier coefficient c_k^δ is of large magnitude, i.e., that y^δ is rich in the component u_k.

The residual L-curve is the graph obtained by connecting consecutive points in the sequence

$$p_k = (k, \ln \|r_k^\delta\|), \qquad k = 1, 2, 3, \dots \ , \tag{2.2}$$

by straight lines. This curve often looks roughly like the letter "L", because typically the Fourier coefficients c_j^δ decrease rapidly to zero in magnitude for small values of j,

but decrease slowly (if at all) to zero in magnitude for larger values of j. The reason for this behavior of the c_j^δ is that the Fourier coefficients $\hat{c}_j$, given by (1.3), of the error-free right-hand side, typically converge to zero quite rapidly with increasing index j. This is a consequence of the Picard condition (1.2). Therefore, while the "signal" $\hat{y}$ dominates the noise $y^\delta - \hat{y}$ in y^δ, we can expect rapid decrease in magnitude of the Fourier coefficients c_j^δ of the contaminated right-hand side y^δ as well. We remark that similarly to the discrepancy principle, the residual L-curve does not explicitly depend on how quickly the singular values decrease with increasing index.

The error $y^\delta - \hat{y}$ typically is oscillatory with many sign changes. For many ill-posed problems, the number of sign changes of the left singular functions u_j increases only slowly with j; in fact, it is quite common that u_j only has $j-1$ sign changes. Therefore, the left singular functions u_j with a small to moderately sized index j are unable to approximate the error $y^\delta - \hat{y}$ well. This results in a very slow decrease of the norm of the residual error (2.1) when the signal ceases to dominate the residual error. The rapid decrease of $\|r_k^\delta\|$ for small values of k (when $\hat{y}$ dominates r_k^δ) and the subsequent slower decrease of $\|r_k^\delta\|$ (when r_k^δ is dominated by the error $y^\delta - \hat{y}$) gives rise to an L-shape of the graph (2.2) with a more or less pronounced "vertex." This vertex helps us separate the signal from the noise in y^δ. Let k_L denote the index of the point at the vertex. Then, ideally, r_k^δ is dominated by noise for $k > k_L$, but not for $k \leq k_L$. This suggests that we should use $x_{k_L}^\delta$ as an approximation of $\hat{x}$.

We note that Bakushinskii [1] shows that the design of a convergent regularization method requires knowledge of the norm of the error δ in y^δ; see also [8, Theorem 3.3]. The residual L-curve is an aid for determining δ.

Example 2.1. We apply the residual L-curve to the approximation of an analytic function which has been contaminated by noise. This example illustrates the ability of the residual L-curve to separate signal from noise. Thus, consider the approximation of the function $y(t) = 1/(2-t)$ on the unit circle in the complex plane by a polynomial. Introduce the inner product of functions g and h that are continuous on the unit circle,

$$\langle g, h\rangle = \frac{1}{2\pi}\int_0^{2\pi} g(e^{i\theta})\overline{h(e^{i\theta})}d\theta, \qquad i = \sqrt{-1}, \tag{2.3}$$

where the bar denotes complex conjugation. The powers of t are orthonormal with respect to this inner product. Since

$$\langle y, t^j\rangle = 2^{-j-1}, \qquad j = 0, 1, 2, \dots\ ,$$

the best polynomial approximant of degree k of $y(t)$ with respect to the norm $\|g\| = \langle g, g\rangle^{1/2}$ is given by

$$q_k(t) = \sum_{j=0}^{k} 2^{-j-1}t^j.$$

Assume that only a contaminated discrete approximation $y^\delta = [y_j^\delta]_{j=1}^n \in \mathbb{C}^n$ of $y(t)$ is available. Specifically, let $\hat{y} = [\hat{y}_j]_{j=1}^n \in \mathbb{C}^n$ be a tabulation of $y(t)$ at equidistant nodes, i.e., $\hat{y}_j = y(e^{2\pi i(j-1)/n})$, and let y^δ be obtained by adding a "noise" vector $w \in \mathbb{R}^n$ with normally distributed components with zero mean to $\hat{y}$. We would like to determine an accurate polynomial approximant of $y(t)$ from the available data y^δ.

The powers t^j are orthonormal also with respect to the discrete inner product

$$\langle g, h\rangle_n = \frac{1}{n}\sum_{j=0}^{n-1} g(e^{2\pi ij/n})\overline{h(e^{2\pi ij/n})}, \tag{2.4}$$

and we determine the polynomial approximants

$$q_k^\delta(t) = \sum_{j=0}^{k} \langle y^\delta, t^j \rangle_n t^j, \qquad k = 0, 1, 2, \ldots \ .$$

We use $n = 32$ in the computations for this example.

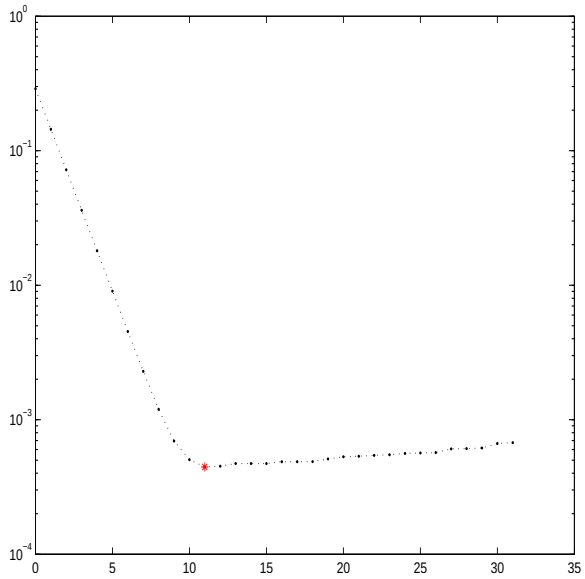

FIG. 2.1. *Example 2.1: Discrete error norm* $\|q_k^\delta(t) - y(t)\|_{320}$ *as a function of* k. *The polynomial* $q_{11}^\delta(t)$ *is seen to furnish the best approximation of* $y(t)$. *The error for this polynomial is marked by* $*$ *(in red).*

Figure 2.1 displays the discrete error norm

$$\|q_k^\delta(t) - y^\delta(t)\|_m = \left(\frac{1}{m} \sum_{j=0}^{m-1} |q_k^\delta(e^{2\pi ij/m}) - y^\delta(e^{2\pi ij/m})|^2 \right)^{1/2}, \qquad k = 0, 1, 2, \ldots \ ,$$

for $m = 10n$. The errors can be seen to decrease as k increases up to $k = 11$, but increase slowly with k for $k \geq 11$. Thus, the polynomial $q_{11}^\delta(t)$ furnishes the most accurate approximation of $y(t)$ with $\|q_k^\delta(t) - y^\delta(t)\|_{320} = 4.5 \cdot 10^{-4}$. We remark that the random vector w is of size $\|w\|_{32} = 6.8 \cdot 10^{-4}$, where $\|w\|_{32} = \langle w, w \rangle_{32}^{1/2}$. It is constructed as described in the beginning of Section 3 to yield a relative error of about $1 \cdot 10^{-3}$; cf. (3.2).

Define the residual vectors $r_k^\delta \in \mathbb{C}^n$ with components $y_j^\delta - q_k^\delta(e^{2\pi i(j-1)/n})$, $1 \leq j \leq n$, for $k = 0, 1, 2, \ldots$. Figure 2.2 displays the residual L-curve determined by the points $(k, \ln \|r_k^\delta\|_{32})$, $k = 0, 1, 2, \ldots$. The residual L-curve is seen to be L-shaped. We apply the triangle method by Castellanos et al. [7] to determine its "vertex." This method yields the vertex $(11, \ln \|r_{11}^\delta\|_{32})$, which is marked by $*$ in Figure 2.2; see the end of this section for further discussion on methods for determining the vertex.

Thus, the residual L-curve combined with the triangle method suggests that the polynomial $q_{11}^\delta(t)$ be used as an approximation of $y(t)$. Recall that Figure 2.1 shows

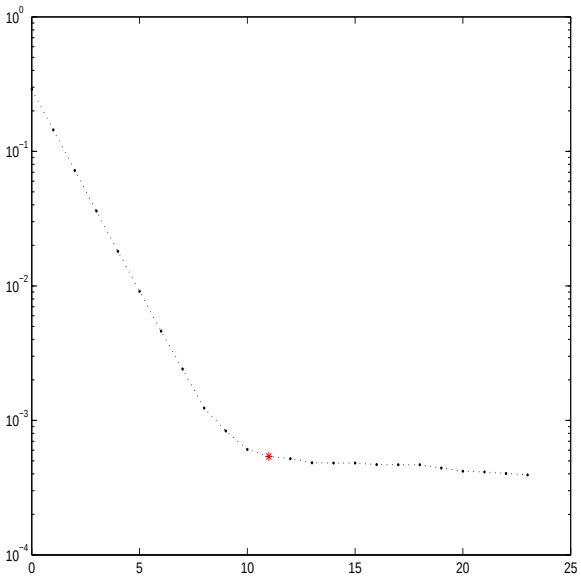

FIG. 2.2. *Example 2.1: Residual L-curve with the vertex determined by the triangle method marked by $*$ (in red).*

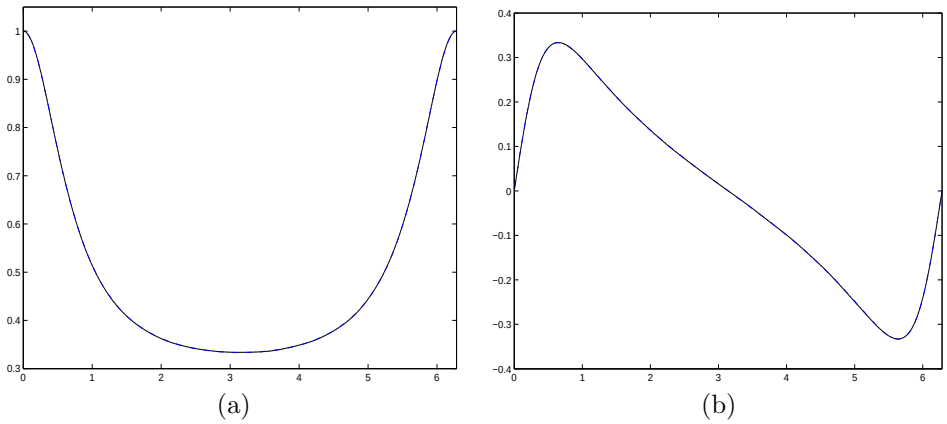

FIG. 2.3. *Example 2.1: Real and imaginary parts of the function $y(e^{i\theta})$ and polynomial $q_{11}(e^{i\theta})$ for $0 \leq \theta \leq 2\pi$: (a) real part of $q_{11}(e^{i\theta})$ (black continuous curve) and of $y(e^{i\theta})$ (blue dash-dotted curve), (b) imaginary part of $q_{11}(e^{i\theta})$ (black continuous curve) and of $y(e^{i\theta})$. The curves in each graph are too close to distinguish.*

$q_{11}^{\delta}(t)$ to be the best polynomial approximant of $y(t)$. Figure 2.3 displays the real and imaginary parts of $y(t)$ and $q_{11}^{\delta}(t)$.

This example illustrates the possibility of using the residual L-curve for determining polynomial approximants in the presence of errors in the data (function values). In Section 3 we describe applications of the residual L-curve to the solution of ill-posed problems with contaminated data. □

Hansen and O'Leary [13] consider the L-curve obtained by connecting consecutive

points in the sequence

$$p_k = (\ln \|x_k^\delta\|, \ln \|r_k^\delta\|), \qquad k = 1, 2, 3, \dots \,, \tag{2.5}$$

by straight lines; see also Hansen [9, 11] and Lawson and Hansen [14]. We refer to this graph as the *standard L-curve*.

We remark that the ordering of the coordinates for the point p_k is arbitrary. The ordering (2.5) is used in [14], as well as in [5, 6] in the context of Tikhonov regularization. The reverse ordering can be found in [7, 9, 11, 12, 13]. We use the ordering (2.5) in the present paper because this simplifies comparison with the residual L-curve determined by the points (2.2). The ordering (2.2) of the coordinates for the latter L-cure is natural, because we consider $\log \|r_k^\delta\|$ a function of k.

Differently from the discrepancy principle and the residual L-curve, the standard L-curve (2.5) explicitly takes the growth of the norm of the computed approximate solutions

$$\|x_k^\delta\| = \left(\sum_{j=1}^{k} \frac{\langle y^\delta, u_j \rangle^2}{\sigma_j^2} \right)^{1/2},$$

and thereby also the decrease of the singular values σ_k to zero, as k increases into account. The standard L-curve has been applied successfully to the solution of many linear discrete ill-posed problems. However, it does not always perform well, e.g., when the norm of the error in y^δ is small; see, e.g., Hansen [11, p. 190] for discussions on shortcomings and references.

While it is fairly straightforward to determine the rough location of the "vertex" in Example 2.1 by visual inspection, the determination of the location of the vertex of the standard and residual L-curves can be a nontrivial task to carry out for a computer program. A reason for this is that the vertex may not be very pronounced and this makes it difficult to define suitable computational criteria for identifying the vertex. Moreover, points p_k for the standard L-curve may be spaced irregularly; the distance between pairs of adjacent points (2.5) may be "tiny" for some values of k and this can make it difficult to define the vertex in a meaningful way. The latter difficulty does not arise with points (2.2) of the residual L-curve since the spacing of the horizontal coordinate is prescribed.

Castellanos et al. [7] recently proposed a scheme referred to as the "triangle method" for locating the vertex of the standard L-curve, and Hansen et al. [12] subsequently developed a faster method, referred to as the "pruning method," for the same purpose. We will apply the triangle and pruning methods to both the standard and residual L-curves in the computed examples of Section 3 and compare the vertices determined by these methods. Note that the coordinates of the data for the M-files by Castellanos et al. [7] and Hansen et al. [12], that implement the triangle and pruning methods, respectively, are switched compared with the notation of the present paper.

The residual L-curve may also be considered within a statistical framework. Let x^δ denote the least-squares solution of minimal norm of (1.4) and let L be a linear functional defined on $\mathcal{X}$. By the Gauss-Markov theorem, the best linear unbiased estimator of $L\hat{x}$ is given given by Lx^δ; see, e.g., Björck [2, Section 1.1.2]. However, the variance of x^δ is large when A is ill-conditioned. The residual L-curve seeks to determine a biased estimator of smaller variance than x^δ.

L-curve	vertex algor.	vertex index	error
residual	triangle	7	$1.1 \cdot 10^{-1}$
residual	pruning	7	$1.1 \cdot 10^{-1}$
standard	triangle	8	$1.5 \cdot 10^{-1}$
standard	pruning	11	$2.8 \cdot 10^{-1}$

TABLE 3.1
Example 3.1: $\eta = 2$, $\hat{k} = 7$.

3. Computed examples. The operators A of this section are $n \times n$ matrices, and the right-hand sides $\hat{y}$ and y^δ are n-vectors. The "noisy" right-hand sides are determined by generating a vector w with normally distributed entries with mean zero and variance one, and computing

$$y^\delta = \hat{y} + w\|\hat{y}\| \cdot 10^{-\eta}/\sqrt{n} \tag{3.1}$$

for some $\eta \geq 0$. In this section $\|\cdot\|$ denotes the Euclidean vector norm. The Central Limit Theorem yields that the relative error in y^δ satisfies

$$\frac{\|y^\delta - \hat{y}\|}{\|\hat{y}\|} \approx 1 \cdot 10^{-\eta}; \tag{3.2}$$

however, we will not use this fact in the numerical experiments. The purpose of the experiments is to demonstrate the application of the residual L-curve to selecting a suitable approximation x_k^δ of $\hat{x}$ in situations when an estimate of $\|y^\delta - \hat{y}\|$ is not available.

We compute approximate solutions x_k^δ, $k = 1, 2, 3, \ldots$, using the `tsvd` function from the MATLAB program package Regularization Tools by Hansen [10]. This package also contains MATLAB programs for discretizing several integral equations of the first kind. We use these programs for determining matrices A and right-hand sides $\hat{y}$ for several examples. The computed examples show the quality of the approximate solutions $x_{k_L}^\delta$ determined by locating the vertex of the residual and standard L-curves using the triangle and pruning methods. We report both the indices k_L and errors $\|x_{k_L}^\delta - \hat{x}\|$ obtained in this manner.

Similarly as for the infinite-dimensional situation considered in Section 2, the norm of the residual error $\|r_k^\delta\|$ typically decreases rapidly with k for small values of k, and slowly for larger values of k. However, due to the finite dimensionality of the problems considered in this section, $\|r_k^\delta\|$ may decrease rapidly again with k when k is close to n; in exact arithmetic $r_n^\delta = 0$. The latter rapid decrease can make it difficult for the triangle and pruning methods to determine an appropriate vertex index. We therefore in all experiments apply these methods to the data $\{p_k\}_{k=1}^{n/2}$ with p_k defined by (2.2) or (2.5). All computations were carried out using MATLAB with unit roundoff $\epsilon \approx 2 \cdot 10^{-16}$.

Example 3.1. Consider the solution of the Fredholm integral equation of the first kind

$$\int_{-6}^{6} \kappa(s,t)x(t)dt = y(s), \qquad -6 \leq s \leq 6, \tag{3.3}$$

discussed by Phillips [16]. Its solution, kernel, and right-hand side are given by

$$x(t) = \begin{cases} 1 + \cos(\frac{\pi}{3}t), & \text{if } |t| < 3, \\ 0, & \text{otherwise,} \end{cases}$$

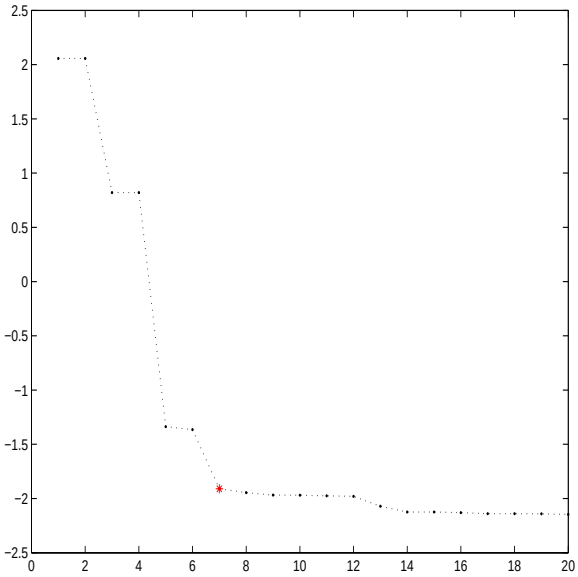

FIG. 3.1. *Example 3.1: Residual L-curve for $\eta = 2$ with the vertex determined by the triangle and pruning methods marked by $*$ (in red). The graph displays the 10-logarithm of $\|r_k^\delta\|$ as a function of k.*

L-curve	vertex algor.	vertex index	error
residual	triangle	7	$7.3 \cdot 10^{-2}$
residual	pruning	7	$7.3 \cdot 10^{-2}$
standard	triangle	19	$2.5 \cdot 10^{-1}$
standard	pruning	19	$2.5 \cdot 10^{-1}$

TABLE 3.2
Example 3.1: $\eta = 3$, $\hat{k} = 11$, $\|x_{11}^\delta - \hat{x}\| = 3.7 \cdot 10^{-2}$.

$$\kappa(s,t) = x(s-t),$$
$$y(s) = (6 - |s|)(1 + \frac{1}{2}\cos(\frac{\pi}{3}s)) + \frac{9}{2\pi}\sin(\frac{\pi}{3}|s|).$$

We use the code phillips from [10] to discretize (3.3) by a Galerkin method with orthonormal box functions as test and trial functions to obtain the symmetric indefinite matrix $A \in \mathbb{R}^{40\times 40}$ and a scaled approximate solution $\hat{x} \in \mathbb{R}^{40}$. We determine the error-free right-hand side from $\hat{y} = A\hat{x}$, and the contaminated right-hand sides y^δ are computed according to (3.1) for $\eta \in \{2, 3, 4\}$.

Tables 3.1-3.3 summarize the results obtained. The triangle method applied to the residual L-curve determines approximate solutions x_k^δ, whose error[1] is less than a factor two larger than the error in the best approximate solution $x_{\hat{k}}^\delta$. In our experiments the triangle method applied to the residual L-curve often determined approximate solutions x_k^δ with index k close to $\hat{k}$. For many problems, but not all, the index k determined is smaller than $\hat{k}$. Thus, the triangle method applied to the residual L-curve tends to deliver an over-regularized approximate solution. This is also true for the pruning method applied to the residual L-curve. In our experience, the triangle

[1] We refer to $\|x_k^\delta - \hat{x}\|$ as the error.

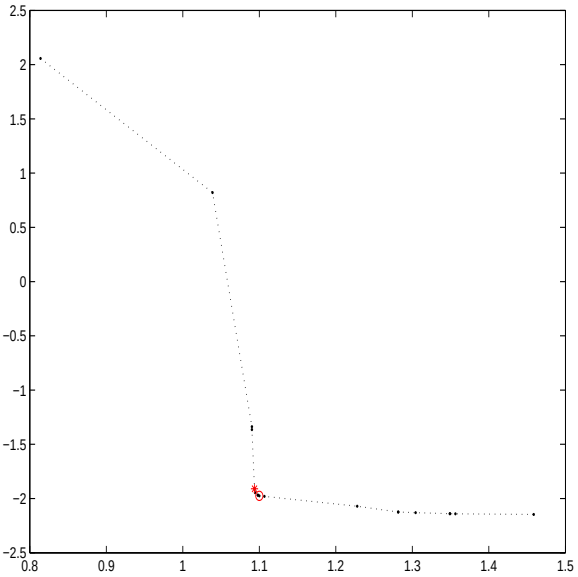

FIG. 3.2. *Example 3.1: Standard L-curve for $\eta = 2$ with the vertex determined by the triangle method marked by $*$ (in red) and the vertex identified by the pruning method marked by o (in red). The graph shows the* 10*-logarithm of $\|r_k^\delta\|$ as a function of the* 10*-logarithm of $\|x_k^\delta\|$.*

L-curve	vertex algor.	vertex index	error
residual	triangle	14	$2.3 \cdot 10^{-2}$
residual	pruning	5	$3.2 \cdot 10^{-1}$
standard	triangle	20	$3.1 \cdot 10^{-2}$
standard	pruning	19	$2.5 \cdot 10^{-2}$

TABLE 3.3
Example 3.1: $\eta = 4$, $\hat{k} = 12$, $\|x_{12}^\delta - \hat{x}\| = 1.2 \cdot 10^{-2}$.

and pruning methods applied to the standard L-curve often determine approximate solutions x_k^δ with index $k > \hat{k}$; hence, the approximate solutions x_k^δ obtained in this manner tend to be under-regularized.

Figure 3.1 shows the residual L-curve for $\eta = 2$ with the vertex determined by the triangle and pruning methods marked by $*$. Figure 3.2 is analogous for the standard L-curve. The approximate solutions selected by these methods, as well as $\hat{x}$, are displayed in Figure 3.3. The magnitude of the exact Fourier coefficients $\hat{c}_k$ and of the contaminated Fourier coefficients c_k^δ for $\eta = 2$ and $\eta = 4$ in (3.1) are displayed in Figure 3.4 for $1 \leq k \leq 40$. In particular, the figure shows that $\hat{c}_{2k} = 0$ for all $k \geq 1$.

The ability of the vertex of the residual L-curve to determine an index k close to the optimal index $\hat{k}$, defined by (1.6), depends on whether the magnitude of the contaminated Fourier coefficients c_{2k-1}^δ decreases rapidly, like the magnitude of the exact Fourier coefficients $\hat{c}_{2k-1}$ as k increases and $k \leq (\hat{k}+1)/2$, and ceases to decrease for $k > (\hat{k}+1)/2$. Figure 3.4 shows the coefficients c_{2k-1}^δ essentially behave in this manner. □

We remark that the location of the vertex may vary with the noise-vector w. The examples presented in this paper display typical performance of the residual L-curve.

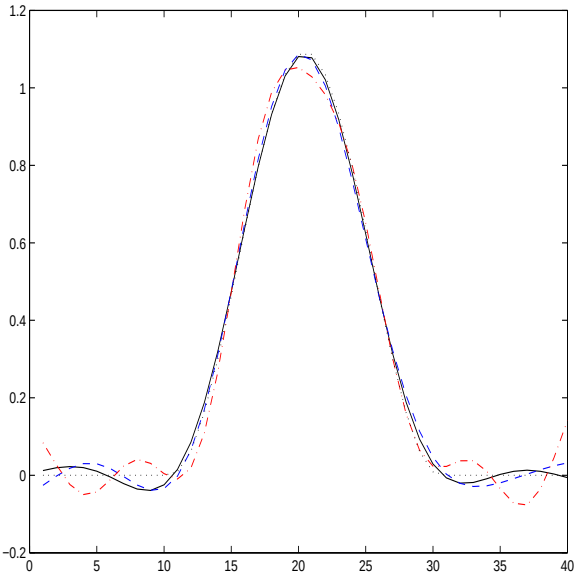

FIG. 3.3. *Example 3.1: Computed approximate solutions for* $\eta = 2$*:* x_7^δ *(black continuous graph),* x_8^δ *(blue dashed graph), and* x_{11}^δ *(red dash-dotted graph). The error-free solution* $\hat{x}$ *is displayed by a black dotted graph. The horizontal axis gives the entry number of the vectors.*

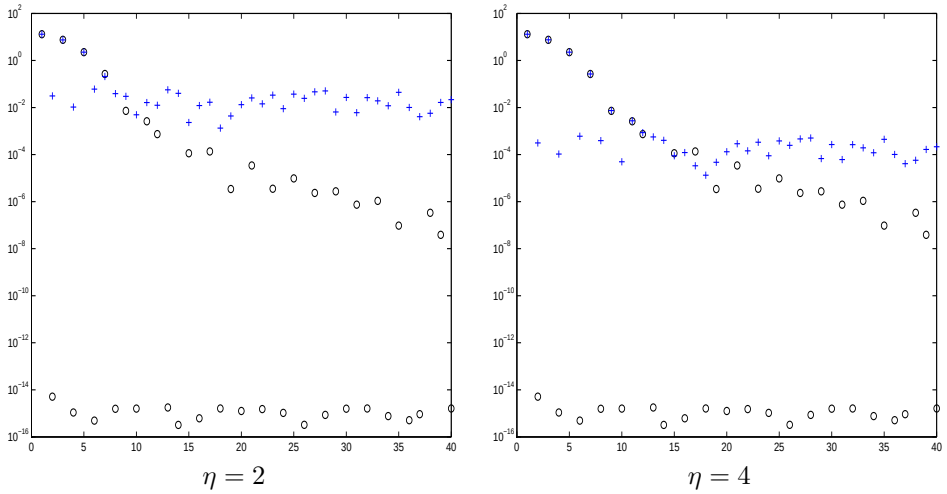

FIG. 3.4. *Example 3.1: Magnitude of noise-free Fourier coefficients* $\hat{c}_j$ *(1.3) marked by* (o) *in black and of the noise-contaminated Fourier coefficients* c_j^δ *(1.8) marked by* $(+)$ *in blue for* $\eta = 2$ *and* $\eta = 4$*.*

Example 3.2. We consider the integral equation

$$\int_{-\pi/2}^{\pi/2} \kappa(s,t)x(t)dt = y(s), \qquad -\frac{\pi}{2} \leq s \leq \frac{\pi}{2}, \tag{3.4}$$

L-curve	vertex algor.	vertex index	error
residual	triangle	4	1.1
residual	pruning	4	1.1
standard	triangle	7	2.0
standard	pruning	7	2.0

TABLE 3.4
Example 3.2: $\eta = 2$, $\hat{k} = 5$, $\|x_5^\delta - \hat{x}\| = 1.1$.

L-curve	vertex algor.	vertex index	error
residual	triangle	7	$3.6 \cdot 10^{-1}$
residual	pruning	7	$3.6 \cdot 10^{-1}$
standard	triangle	8	$6.9 \cdot 10^{-1}$
standard	pruning	8	$6.9 \cdot 10^{-1}$

TABLE 3.5
Example 3.2: $\eta = 3$, $\hat{k} = 7$.

where

$$\kappa(s,t) = (\cos(s) + \cos(t))^2 \left(\frac{\sin(\xi)}{\xi} \right)^2, \qquad \xi = \pi(\sin(s) + \sin(t)),$$

and the right-hand side $y(s)$ is chosen so that the solution $x(t)$ is a sum of two Gaussian functions. This integral equation is discussed by Shaw [17]. We use the code shaw from [10] to discretize (3.4) by a quadrature rule with 40 nodes. This yields the matrix $A \in \mathbb{R}^{40\times 40}$ and solution $\hat{x} \in \mathbb{R}^{40}$. The error-free right-hand side is determined by $\hat{y} = A\hat{x}$ and the contaminated right-hand sides by (3.1) for $\eta \in \{2, 3, 4\}$. The computed results are reported in Tables 3.4-3.6.

Figure 3.5 is analogous to Figure 3.4 and shows the magnitude of the exact and contaminated Fourier coefficients $\hat{c}_k$ and c_k^δ, respectively, for right-hand side vectors y^δ generated by (3.1) with $\eta = 2$ and $\eta = 4$.

Similarly as in Example 3.1, the triangle and pruning methods applied to the residual L-curve give smaller vertex indices than when these methods are applied to the standard L-curve. For $\eta \in \{2, 3, 4\}$, the approximate solutions determined using the residual L-curve are better approximations of $\hat{x}$ than the approximate solutions determined using the standard L-curve.

We remark that the residual and standard L-curves also can be applied in conjunction with other approaches to selecting a suitable approximate solution. For instance, we may use the approximate solution selected by the triangle or pruning methods as starting point for the selection criterion described by Morigi et al. [15], who advocate to choose an approximate solution x_k^δ with $\|x_{k+1}^\delta - x_k^\delta\|$ small. Let k_L be a vertex index determined by the triangle or pruning method. The function

$$f(k) = \|x_{k+1}^\delta - x_k^\delta\|, \qquad 1 \le k < n, \tag{3.5}$$

typically has many local maxima and minima. Let k' be a local minimum closest to k_L. In the present example, for $\eta \in \{2, 3, 4\}$, this selection criterion, with the vertex indices k_L obtained from the standard L-curve as input, determines the vertex indices selected by the triangle and pruning methods applied to the residual L-curve. For example, when $\eta = 3$, the triangle and pruning methods applied to the standard

L-curve	vertex algor.	vertex index	error
residual	triangle	7	$3.0 \cdot 10^{-1}$
residual	pruning	7	$3.0 \cdot 10^{-1}$
standard	triangle	9	$3.0 \cdot 10^{-1}$
standard	pruning	9	$3.0 \cdot 10^{-1}$

TABLE 3.6
Example 3.2: $\eta = 4$, $\hat{k} = 7$.

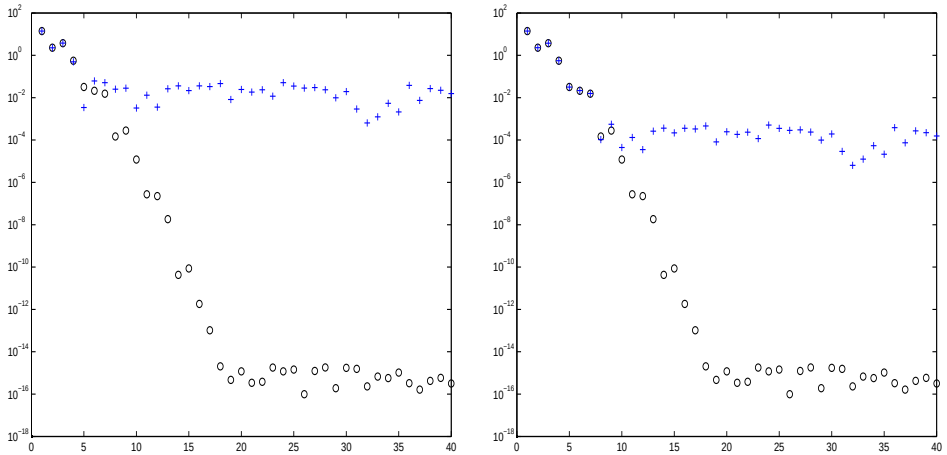

FIG. 3.5. *Example 3.2: Magnitude of noise-free Fourier coefficients* $\hat{c}_j$ *(1.3) marked by* (o) *in black and of the noise-contaminated Fourier coefficients* c_j^δ *(1.8) marked by* $(+)$ *in blue for* $\eta = 2$ *and* $\eta = 4$.

L-curve determine the vertex index 8; see Table 3.5. The closest local minimum of the function (3.5) is at $k = 7$. □

Example 3.3. The inverse Laplace transform

$$\int_0^\infty \exp(-st)x(t)dt = \frac{2}{2s+1}, \qquad s \geq 0, \tag{3.6}$$

has the solution $x(t) = \exp(-t/2)$. We discretize (3.6) using the function ilaplace in [10] to obtain the matrix $A \in \mathbb{R}^{20\times 20}$ and vectors $\hat{x}, \hat{y}, y^\delta \in \mathbb{R}^{20}$ similarly as in Example 3.1.

Tables 3.7-3.9 show the performance of the triangle and pruning methods applied to the residual and standard L-curves. As in the previous examples, the residual L-curve yields vertices with a smaller index. The approximate solutions x_k^δ associated with vertices of the residual L-curve are more accurate approximations of $\hat{x}$ than the approximate solutions associated with vertices of the standard L-curve.

Figure 3.6 shows the residual L-curve for $\eta = 4$ with the vertex determined by the triangle and pruning methods marked by $*$. Figure 3.7 is analogous for the standard L-curve. The approximate solutions selected by these methods are displayed by Figure 3.8. Finally, Figure 3.9 is analogous to Figure 3.5 and shows the magnitude of the exact and contaminated Fourier coefficients $\hat{c}_k$ and c_k^δ, respectively, for $\eta = 2$ and $\eta = 4$ in (3.1). □

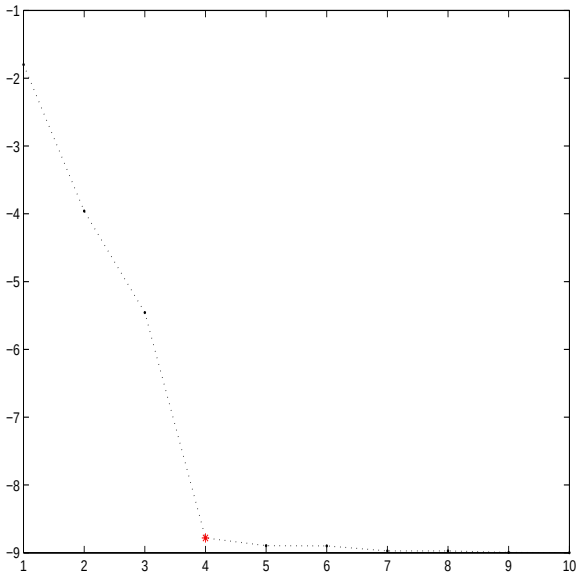

FIG. 3.6. *Example 3.3: Residual L-curve for $\eta = 4$ with the vertex determined by the triangle and pruning methods marked by $*$ (in red).*

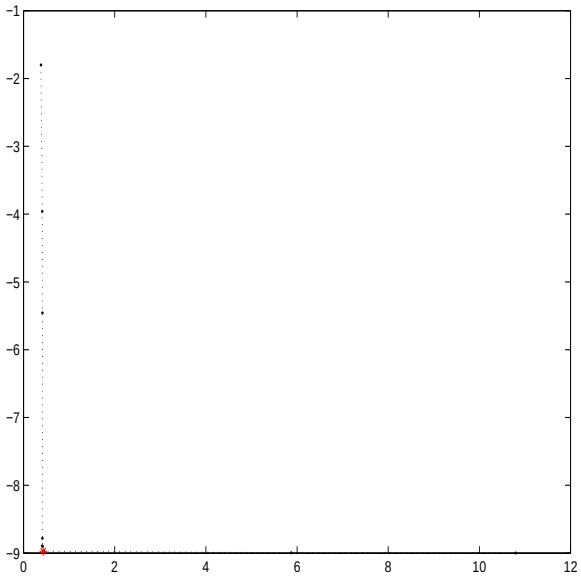

FIG. 3.7. *Example 3.3: Standard L-curve for $\eta = 4$ with the vertex determined by the triangle and pruning methods marked by $*$ (in red).*

4. Conclusion and extensions. This paper introduces a new aid, the residual L-curve, for determining a suitable truncation index for the TSVD method. The residual L-curve is designed to determine the amount of noise in the data (the available right-hand side). The computed examples illustrate that the vertex of the residual L-curve may determine better approximations of the error-free solution $\hat{x}$ than the

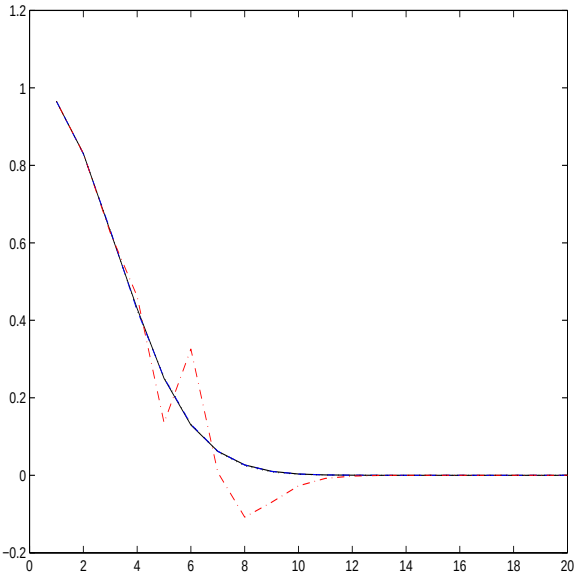

FIG. 3.8. *Example 3.3: Computed approximate solutions for $\eta = 4$: x_4^δ (black continuous graph), x_6^δ (blue dashed graph), and x_7^δ (red dash-dotted graph). The error-free solution $\hat{x}$ is displayed by a black dotted graph. All but the graph for x_7^δ are too close to distinguish.*

L-curve	vertex algor.	vertex index	error
residual	triangle	3	$7.2 \cdot 10^{-2}$
residual	pruning	3	$7.2 \cdot 10^{-2}$
standard	triangle	5	$2.7 \cdot 10^{-1}$
standard	pruning	5	$2.7 \cdot 10^{-1}$

TABLE 3.7
Example 3.3: $\eta = 2$, $\hat{k} = 4$, $\|x_4^\delta - \hat{x}\| = 4.4 \cdot 10^{-2}$.

vertex of the standard L-curve. We therefore believe that the residual L-curve can be a valuable aid for determining approximate solutions of ill-posed problems in situations when not enough information of the solution and the error in the right-hand side is available to allow the application of other techniques with a better theoretical justification, such as the discrepancy principle.

In most of our examples, the triangle and pruning methods determined the same vertex, but occasionally the triangle method performed better than the pruning method; see Tables 3.1 and 3.3.

The residual L-curve also can be applied in conjunction with other solution methods than TSVD, such as with iterative methods, extrapolation methods, and multi-parameter methods; see, e.g., Brezinski et al. [3, 4] for discussions of the latter kinds of methods. We are presently investigating the application of the residual L-curve with some of these methods.

Acknowledgments. We would like to thank Valia Guerra for making MATLAB code for the triangle method available and Kazim Khan for discussions.

REFERENCES

L-curve	vertex algor.	vertex index	error
residual	triangle	4	$8.4 \cdot 10^{-3}$
residual	pruning	4	$8.4 \cdot 10^{-3}$
standard	triangle	6	$3.4 \cdot 10^{-2}$
standard	pruning	6	$3.4 \cdot 10^{-2}$

TABLE 3.8
Example 3.3: $\eta = 3$, $\hat{k} = 4$.

L-curve	vertex algor.	vertex index	error
residual	triangle	4	$7.1 \cdot 10^{-3}$
residual	pruning	4	$7.1 \cdot 10^{-3}$
standard	triangle	7	$2.8 \cdot 10^{-1}$
standard	pruning	7	$2.8 \cdot 10^{-1}$

TABLE 3.9
Example 3.3: $\eta = 4$, $\hat{k} = 6$, $\|x_6^\delta - \hat{x}\| = 4.0 \cdot 10^{-3}$.

[1] A. B. Bakushinskii, Remarks on choosing a regularization parameter using quasi-optimality and ratio criterion, USSR Comp. Math. Math. Phys., 24,4 (1984), pp. 181–182.

[2] Å. Björck, Numerical Methods for Least Squares Problems, SIAM, Philadelphia, 1996.

[3] C. Brezinski, M. Redivo-Zaglia, G. Rodriguez, and S. Seatzu, Extrapolation techniques for ill-conditioned linear systems, Numer. Math., 81 (1998), pp. 1–29.

[4] C. Brezinski, M. Redivo-Zaglia, G. Rodriguez, and S. Seatzu, Multi-parameter regularization techniques for ill-conditioned linear systems, Numer. Math., 94 (2003), pp. 203–228.

[5] D. Calvetti, G. H. Golub, and L. Reichel, Estimation of the L-curve via Lanczos bidiagonalization, BIT, 39 (1999), pp. 603–619.

[6] D. Calvetti, L. Reichel, and A. Shuibi, L-curve and curvature bounds for Tikhonov regularization, Numer. Algorithms, 35 (2004), pp. 301–314.

[7] J. L. Castellanos, S. Gómez, and V. Guerra, The triangle method for finding the corner of the L-curve, Appl. Numer. Math., 43 (2002), pp. 359–373.

[8] H. W. Engl, M. Hanke, and A. Neubauer, Regularization of Inverse Problems, Kluwer, Dordrecht, 1996.

[9] P. C. Hansen, Analysis of the discrete ill-posed problems by means of the L-curve, SIAM Review, 34 (1992), pp. 561–580.

[10] P. C. Hansen, Regularization tools: A MATLAB package for analysis and solution of discrete ill-posed problems, Numer. Algorithms, 6 (1994), pp. 1–35. Software is available in Netlib at `http://www.netlib.org`.

[11] P. C. Hansen, Rank-Deficient and Discrete Ill-Posed Problems, SIAM, Philadelphia, 1998.

[12] P. C. Hansen, T. K. Jensen, and G. Rodriguez, An adaptive pruning algorithm for the discrete L-curve criterion, J. Comput. Appl. Math., 198 (2007), pp. 483–492. MATLAB code for the algorithm is available at `http://www2.imm.dtu.dk/~pch`.

[13] P. C. Hansen and D. P. O'Leary, The use of the L-curve in the regularization of discrete ill-posed problems, SIAM J. Sci. Comput., 14 (1993), pp. 1487–1503.

[14] C. L. Lawson and R. J. Hansen, Solving Least Squares Problems, Prentice-Hall, Englewood Cliffs, 1974.

[15] S. Morigi, L. Reichel, F. Sgallari, and F. Zama, Iterative methods for ill-posed problems and semiconvergent sequences, J. Comput. Appl. Math., 193 (2006), pp. 157–167.

[16] D. L. Phillips, A technique for the numerical solution of certain integral equations of the first kind, J. ACM, 9 (1962), pp. 84–97.

[17] C. B. Shaw, Jr., Improvements of the resolution of an instrument by numerical solution of an integral equation, J. Math. Anal. Appl., 37 (1972), pp. 83–112.

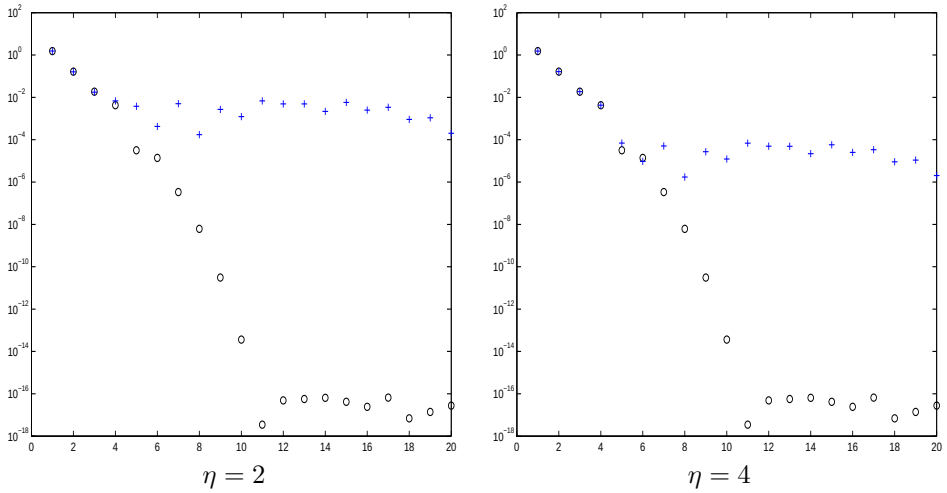

FIG. 3.9. *Example 3.3: Magnitude of noise-free Fourier coefficients* $\hat{c}_j$ *(1.3) marked by* (o) *in black and of the noise-contaminated Fourier coefficients* c_j^δ *(1.8) marked by* $(+)$ *in blue for* $\eta = 2$ *and* $\eta = 4$.