BIT Numer Math (2009) 49: 669–696
DOI 10.1007/s10543-009-0239-7

BIT

The regularizing effect of the Golub-Kahan iterative bidiagonalization and revealing the noise level in the data

Iveta Hnětynková · Martin Plešinger · Zdeněk Strakoš

Received: 31 October 2008 / Accepted: 27 August 2009 / Published online: 22 September 2009

Abstract Regularization techniques based on the Golub-Kahan iterative bidiagonalization belong among popular approaches for solving large ill-posed problems. First, the original problem is projected onto a lower dimensional subspace using the bidiagonalization algorithm, which by itself represents a form of regularization by projection. The projected problem, however, inherits a part of the ill-posedness of the original problem, and therefore some form of inner regularization must be applied. Stopping criteria for the whole process are then based on the regularization of the projected (small) problem.

In this paper we consider an ill-posed problem with a noisy right-hand side (observation vector), where the noise level is unknown. We show how the information

Communicated by Lars Eldén.

The work of the first author was supported by the research project MSM0021620839 financed by MŠMT. The work of the second and the third author was supported by the GAAS grant IAA100300802, and by the Institutional Research Plan AV0Z10300504.

I. Hnětynková (✉) · M. Plešinger · Z. Strakoš
Institute of Computer Science, Academy of Sciences, Pod Vodárenskou věží 2, Prague 8, Czech Republic
e-mail: hnetynkova@cs.cas.cz

M. Plešinger
e-mail: martin.plesinger@tul.cz

Z. Strakoš
e-mail: strakos@cs.cas.cz

I. Hnětynková · Z. Strakoš
Faculty of Mathematics and Physics, Charles University in Prague, Sokolovská 83, Prague 8, Czech Republic

M. Plešinger
Faculty of Mechatronics, Technical University of Liberec, Studentská 2/1402, Liberec, Czech Republic

from the Golub-Kahan iterative bidiagonalization can be used for estimating the noise level. Such information can be useful for constructing efficient stopping criteria in solving ill-posed problems.

Keywords Ill-posed problems · Golub-Kahan iterative bidiagonalization · Lanczos tridiagonalization · Noise revealing

Mathematics Subject Classification (2000) 15A06 · 15A18 · 15A23 · 65F10 · 65F22

1 Introduction

Consider an ill-posed linear algebraic system with the right-hand side b contaminated by *white noise*

$$Ax \approx b, \qquad A \in \mathbb{R}^{n\times n}, \qquad b = b^{\text{exact}} + b^{\text{noise}} \in \mathbb{R}^n, \tag{1.1}$$

with A nonsingular and the goal to numerically approximate the exact solution

$$x \equiv x^{\text{exact}} = A^{-1}b^{\text{exact}}. \tag{1.2}$$

Linear approximation problems of this form arise in a broad class of applications. In many cases the matrix A represents a discretized smoothing operator with the singular values of A decaying gradually without a noticeable gap. Since A is ill-conditioned, the presence of the noise makes the naive solution $x^{\text{naive}} = A^{-1}b$ meaningless. Therefore it is necessary to use regularization techniques for finding an acceptable numerical approximation to (1.2) which reflects a sufficient amount of information contained in the data, while suppressing the devastating influence of the noise. In image processing A typically represents a discretized blurring operator. In other applications A might be of different origin. Throughout the paper we assume A square and nonsingular. The presented methods can be extended to a general rectangular case, which is also confirmed by numerical experiments. However, the analysis contains some subtle points which would further extend the length of the paper.

For the unknown noise component b^{noise} we assume

$$\|b^{\text{noise}}\| \ll \|b^{\text{exact}}\|, \tag{1.3}$$

where $\|v\|$ denotes the standard Euclidean norm of the vector v. We will further assume that multiplication of a vector v by A and A^T results in smoothing which reduces the relative sizes of the high frequency components of v. In particular, in comparison to v, the vectors A^TAv and AA^Tv have significantly reduced high frequency components.

The Golub-Kahan iterative bidiagonalization algorithm [14] is widely used for solving large ill-posed problems. In hybrid methods, see, e.g., [18, Chap. 6.7, pp. 162–164] or [10, 27, 28, 36], the outer bidiagonalization (which itself represents

a regularization of the original large problem by projection) is combined with an inner regularization of the projected small problem. The bidiagonalization is stopped when the regularized solution of the projected problem matches some selected stopping criteria. They are typically based, amongst others (see [1, 2]) on estimation of the L-curve and it's curvature [4–6], estimation of the distance between the exact and regularized solution [35], the discrepancy principle [32, 33], and cross validation methods [7, 15, 34]. These techniques have been studied and compared in the context of regularization, e.g., in [18, Chap. 7, pp. 175–208] and in [27].

It is worthwhile recalling that the regularization idea was mentioned, although not fully developed, by Golub and Kahan [14]. It was described in detail in relation to using the Golub-Kahan iterative bidiagonalization, under the name damped least squares, by Paige and Saunders in [42]; see also [41] or [48]. Paige and Saunders alluded to some older interesting references in [42]. From the more recent literature we mention [16], where hybrid methods based on bidiagonalization are described as least-squares projection methods, and [51], where bidiagonalization is used to compute low-rank approximations of large sparse matrices. Numerical stability of the bidiagonalization algorithm was studied and new stable variants have been proposed, e.g., in [3, 51], with a simplified analysis of [3] presented in [40].

A new contribution to the theoretical background for hybrid methods has recently been presented in [43–45]. In exact arithmetic, the bidiagonalization provides a fundamental decomposition of the matrix of the data $[b, A]$. When the bidiagonalization stops, it reveals the so called *core problem* represented by the computed bidiagonal matrix. The core problem is minimally dimensioned and it contains the necessary and sufficient information for solving the original orthogonally invariant linear approximation problem $Ax \approx b$. Whenever the bidiagonalization is stopped before reaching the core problem, it gives (in exact arithmetic) its leading left upper part. Consequently, the approximate solution computed at the given step is based on *information which is necessary for solving the original problem* and it is not influenced by any part of the redundant or irrelevant information. From the properties of the core problem it then follows (see [23, 45]) that further possible steps can be considered a refinement of the current approximation. Ill-posed problems from the core problem point of view were studied in [49, 50]. It should be mentioned that the application of the core problem theory to results of finite precision computations needs further investigation.

In this paper we focus on the Golub-Kahan iterative bidiagonalization and investigate how the noise contained in the right-hand side b is propagated to the projected problem. We demonstrate under the given assumptions that *the unknown noise level* can be determined from the information available during the bidiagonalization process. The knowledge of the noise level can further be used in construction of stopping criteria for hybrid methods.

Similar ideas are used in [20, 46, 47] for selection of a value of the regularization parameter for which the residual vector changes from being dominated by the remaining signal to being white-noise like. This leads to a parameter-choice method based on Fourier analysis of residual vectors. In [25] regularization properties of iterative methods GMRES, MINRES, RRGMRES, MR-II and CGLS are studied. It is shown that MINRES, MR-II and CGLS filter out large singular value decomposition (SVD) components of the residual, but this is not always true for GMRES and

RRGMRES, where the SVD components are mixed [25, Sect. 3]. The noise propagation to reconstructed images computed by regularizing iterations is further studied in [19]. Our approach, which uses information from the iterative bidiagonalization algorithm for estimating the level of the noise in the data, offers another view on the noise propagation studied in [19, 20, 25, 46, 47].

The paper is organized as follows. Section 2 gives a brief recollection of the Golub-Kahan iterative bidiagonalization, its relationship to the Lanczos tridiagonalization and to approximation of the distribution function in the corresponding Riemann-Stieltjes integral. Section 3 describes propagation of the noise in the bidiagonalization vectors. Section 4 shows how to estimate at a negligible cost the unknown noise level in the original data. This can lead to construction of stopping criteria for the bidiagonalization process as well as to construction of the regularized solution using many different approaches which can be applied in subsequent steps, cf. [27, Sect. 3.2]. It is important to emphasize that the first four sections of our paper deal with *mathematical properties* of the problem and of the methods, i.e., they assume *exact arithmetic*. Therefore, unless specified otherwise, the presented experiments were performed with double reorthogonalization of the computed sequences of the bidiagonalization vectors, which ensures preserving orthogonality close to machine precision. Up to now, effects of rounding errors in solving ill-posed problems were not, to our knowledge, thoroughly investigated. In the references known to us, loss of orthogonality in solving ill-posed problems is not considered a fundamental issue although there are papers that acknowledge that it can be a potential problem. If the loss of orthogonality occurs, then the reader is referred to some form of reorthogonalization, with no further investigation of the differences between the non-reorthogonalized and reorthogonalized computations. In many cases the implementation details are missing, and it is unclear whether the experiments used reorthogonalization or not. We will show in Sect. 5 that effects of rounding errors in solving ill-posed problems can be substantial and the matter should be investigated further. Concluding remarks summarize the main ideas and formulate open questions.

2 Golub-Kahan iterative bidiagonalization

Given the initial vectors $w_0 \equiv 0, s_1 \equiv b/\beta_1$, where $\beta_1 \equiv \|b\| \neq 0$, the Golub-Kahan iterative bidiagonalization computes for $j = 1, 2, \ldots$

$$\begin{aligned} \alpha_j w_j &= A^T s_j - \beta_j w_{j-1}, \qquad \|w_j\| = 1, \\ \beta_{j+1} s_{j+1} &= A w_j - \alpha_j s_j, \qquad \|s_{j+1}\| = 1 \end{aligned} \tag{2.1}$$

until $\alpha_j = 0$ or $\beta_{j+1} = 0$, or until the process is stopped by reaching the dimensionality of the problem. Let $S_k = [s_1, \ldots, s_k]$ and $W_k = [w_1, \ldots, w_k]$ be the matrices with the left and right bidiagonalization vectors as their (orthonormal) columns, and

$$L_k \equiv \begin{bmatrix} \alpha_1 & & & \\ \beta_2 & \alpha_2 & & \\ & \ddots & \ddots & \\ & & \beta_k & \alpha_k \end{bmatrix}, \qquad L_{k+} \equiv \begin{bmatrix} L_k \\ \beta_{k+1} e_k^T \end{bmatrix}, \tag{2.2}$$

where e_k is the kth vector of the standard Euclidean basis. Then the first k steps of the Golub-Kahan iterative bidiagonalization can be written in the matrix form as

$$A^T S_k = W_k L_k^T, \qquad A W_k = [S_k, s_{k+1}] L_{k+}, \tag{2.3}$$

see [14, 38].

This algorithm is closely related to the Lanczos tridiagonalization of a symmetric matrix [29, 30]. The Lanczos tridiagonalization of the matrix AA^T with the starting vector $s_1 = b/\beta_1$, $\beta_1 = \|b\|$, yields in k steps the symmetric tridiagonal matrix T_k such that

$$AA^T S_k = S_k T_k + \alpha_k \beta_{k+1} s_{k+1} e_k^T, \tag{2.4}$$

and

$$T_k = L_k L_k^T = \begin{bmatrix} \alpha_1^2 & \alpha_1\beta_2 & & \\ \alpha_1\beta_2 & \alpha_2^2+\beta_2^2 & \ddots & \\ & \ddots & \ddots & \alpha_{k-1}\beta_k \\ & & \alpha_{k-1}\beta_k & \alpha_k^2+\beta_k^2 \end{bmatrix},$$

i.e. the matrix L_k from the Golub-Kahan iterative bidiagonalization represents a Cholesky factor of the matrix T_k. For a more detailed description of the outlined relationship we refer to [23] and to the literature given there.

Using the results from [22, Sect. 14], the Lanczos tridiagonalization of a given matrix B generates at each step k a non-decreasing piecewise constant distribution function $\omega^{(k)}$, with the nodes being the (distinct) eigenvalues of the Lanczos matrix T_k and the weights $\omega_j^{(k)}$ being the squared first components of the corresponding normalized eigenvectors. The distribution functions $\omega^{(k)}$, $k = 1, 2, \ldots$ represent approximations to the distribution function ω, a non-decreasing piecewise constant function with the nodes being the distinct eigenvalues $\lambda_1, \ldots, \lambda_t$ of the matrix B and the weights ω_j being the squared components of the normalized starting vector in the direction of the jth invariant subspace of B, $j = 1, \ldots, t$, for more details see, e.g., [31, Sect. 2.2], [11].

Consider the singular value decomposition (SVD) of the bidiagonal matrix of the projected problem

$$L_k = P_k \Theta_k Q_k^T, \tag{2.5}$$

where $P_k^{-1} = P_k^T$, $Q_k^{-1} = Q_k^T$, and Θ_k is a diagonal matrix with singular values $\theta_1^{(k)}, \ldots, \theta_k^{(k)}$ of L_k on its diagonal ordered in *increasing order*. Please note that here the increasing order of the singular values follows the standard ordering of the Ritz values in the Lanczos method. From (2.4) it follows that if $B = AA^T$ and the Lanczos tridiagonalization starts with the vector $s_1 = b/\beta_1$, $\beta_1 = \|b\|$, then

$$T_k \equiv L_k L_k^T = P_k \Theta_k^2 P_k^T$$

is the spectral decomposition of T_k, $(\theta_\ell^{(k)})^2$ are its eigenvalues and $p_\ell^{(k)} \equiv P_k e_\ell$ its eigenvectors, $\ell = 1, \ldots, k$. Furthermore, consider the SVD of A,

$$A = U \Sigma V^T = \sum_{j=1}^{n} \sigma_j u_j v_j^T, \tag{2.6}$$

where $U = [u_1, \ldots, u_n]$, $V = [v_1, \ldots, v_n]$, $U^{-1} = U^T$, $V^{-1} = V^T$, and Σ is a diagonal matrix with the singular values $\sigma_1, \ldots, \sigma_n$ of A on its diagonal ordered in nonincreasing order, $\sigma_n > 0$. Then

$$AA^T = U \Sigma^2 U^T$$

is the spectral decomposition of the matrix AA^T, σ_j^2 are its nonzero eigenvalues and u_j the corresponding eigenvectors, $j = 1, \ldots, n$. Summarizing, the Lanczos tridiagonalization (2.4) generates at each step k the distribution function $\omega^{(k)}$ with the nodes $(\theta_\ell^{(k)})^2$ and the weights $|(p_\ell^{(k)}, e_1)|^2$ that approximates the distribution function ω with the nodes $\sigma_n^2, \ldots, \sigma_1^2$ and the weights $|(b/\beta_1, u_n)|^2, \ldots, |(b/\beta_1, u_1)|^2$.

The fact that the weights of ω are determined by the projections of the normalized right-hand side vector b onto the left singular subspaces of A is especially important. As we will see in Sect. 3, until the level determined by the noise is reached, the absolute value of each of these projections is related to the size of the corresponding singular value. This is because of the discrete Picard condition (see Sect. 3).

Mathematically, the bidiagonalization algorithm yields two sequences of subproblems. Consider an approximation to the solution of (1.1) in the subspace generated by the vectors $w_1, \ldots, w_k$, i.e. $x_k = W_k y_k$, so that in seeking $AW_k y_k \approx b$,

$$r_k = b - AW_k y_k = S_{k+1}(e_1\beta_1 - L_{k+} y_k) = S_k(e_1\beta_1 - L_k y_k) - s_{k+1}\beta_{k+1} e_k^T y_k.$$

If the approximation is computed by ensuring that the residual r_k is orthogonal to S_k, then we obtain

$$L_k y_k^{\mathrm{CGME}} = e_1\beta_1, \qquad L_k \in \mathbb{R}^{k\times k}, \tag{2.7}$$

which corresponds to the CGME method; see, e.g., [16]. If the approximation is computed by minimizing the norm of the residual r_k, then we get

$$\| L_{k+} y_k^{\mathrm{LSQR}} - e_1\beta_1 \| = \min_y \| L_{k+} y - e_1\beta_1 \|, \quad L_{k+} \in \mathbb{R}^{(k+1)\times k}, \tag{2.8}$$

which corresponds to the CGLS or LSQR method; see [22, 41, 42]. In both cases, $[\beta_1 e_1, L_k]$ respectively $[\beta_1 e_1, L_{k+}]$ approximate the core problem within $Ax \approx b$; see [23, 24, 43]. In short, the bidiagonalization *concentrates the useful information from the data* $[b, A]$ *to the leading principal bidiagonal block*. In noisy ill-posed problems (1.1), however, the subproblems (2.7), (2.8) can also be polluted by the noise.

3 Propagation of the noise in the Golub-Kahan iterative bidiagonalization

Noise propagation in the Golub-Kahan iterative bidiagonalization is illustrated for the problem `shaw` from the Regularization Toolbox [17]. The matrix A, the exact right-hand side b^{exact} and the exact solution x were determined by `[A,b_exact,x] = shaw(400)`. We add the noise b^{noise} as a random vector using the MATLAB function `randn(400,1)`, scaled in order to obtain different *noise levels* δ_{noise},

$$\delta_{\text{noise}} \equiv \frac{\|b^{\text{noise}}\|}{\|b^{\text{exact}}\|}. \tag{3.1}$$

Figure 1(a) shows the singular values σ_j of the matrix A (solid line) computed by the MATLAB function `svd`, and the absolute values of the projections $|(b, u_j)|$ of the noisy right-hand side b onto the left singular vector subspaces of the matrix A. We use the noise levels $\delta_{\text{noise}} = 10^{-14}$, 10^{-8}, and 10^{-4}. Until the noise level is reached, the absolute values of the right-hand side projections onto the left singular vector subspaces are close to the corresponding singular values. This is given by the fact that, for this problem, $|(b^{\text{exact}}, u_j)|$, $j = 1, 2, \ldots$ satisfy the discrete Picard condition: on average, they decay faster than the singular values of A; see [25], [20, Sect. 4] and Fig. 1(b). For the subspaces corresponding to small singular values, the projections of b are completely dominated by the noise, and the discrete Picard condition for b is thus drastically violated.

Consider the vectors s_k, w_k generated by the bidiagonalization algorithm (2.1) described in the previous section. The starting vector $s_1 = b/\|b\|$ is the normalized noisy right-hand side and therefore it is *contaminated by the noise*. The vector s_2 is obtained from s_1 as follows. First, application of the smoothing operator AA^T

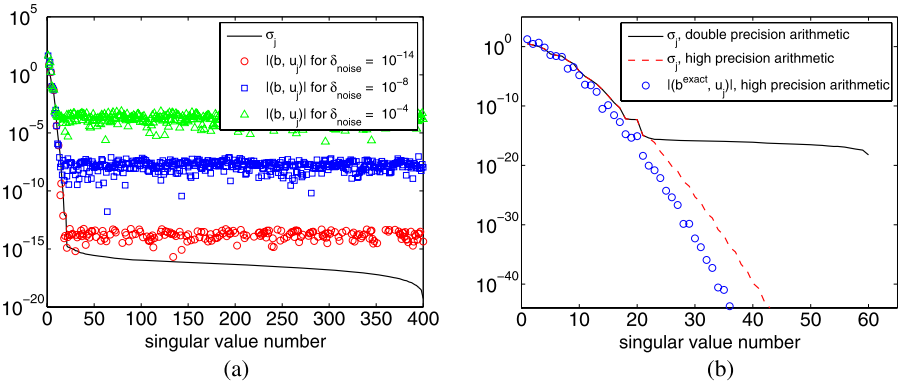

Fig. 1 The singular values σ_j and the absolute values of the projections of the noisy right-hand side b onto the left singular vectors of A for the problem `shaw(400)` and the noise levels $\delta_{\text{noise}} = 10^{-14}$, 10^{-8}, and 10^{-4} (**a**). The singular values of the matrix A computed by the statement `[A,b,x]=shaw(60)` using the standard routine from the Regularization Toolbox [17] (*solid line*) compared with the singular values of the matrix A computed by the modified routine `shaw_vpa`, cf. http://www.cs.cas.cz/krylov, section 'Software', (*dash-dotted line*) using high precision arithmetic guaranteeing 128 decimal digits (**b**). The discrete Picard condition is illustrated by plotting $|(b^{\text{exact}}, u_j)|$, where b^{exact} is computed using the routine `shaw_vpa`

(see (2.4)) smooths out the high frequency components in both b^{exact} and b^{noise}; see also [25, Sect. 3.1], [19, Sect. 4.3]. The subsequent orthogonalization of $AA^T s_1$ against s_1 represents a linear combination of s_1 contaminated by the noise and $AA^T s_1$ *which has been smoothed*. Therefore the contamination of s_1 by the *high frequency part* of the noise is transferred, with multiplication by a scalar coefficient, to s_2. Adding a multiple of $AA^T s_1$ eliminates *a portion of the smooth part* of s_1. Therefore the relative level of the high frequency part of the noise can be expected to be higher in s_2 (which is orthogonal to s_1) than in s_1. Analogous high frequency noise transfer takes place for any k with the vector s_{k+1} obtained from $AA^T s_k$ through the orthogonalization against the vectors s_{k-1} and s_k with a subsequent normalization.

The rest of Sect. 3 is structured into two complementary subsections. First we use the relationship between the Golub-Kahan iterative bidiagonalization and the Lanczos tridiagonalization described in Sect. 2. The noise amplification is described as an effect of damping (filtering out) the smooth (low frequency) components due to convergence of Ritz values to large eigenvalues. Then we show that smoothing and orthogonalization in the Golub-Kahan iterative bidiagonalization can be interpreted as a step-by-step elimination of the dominant smooth (low frequency) components, and it therefore leads to revealing of the high frequency noise.

3.1 Lanczos tridiagonalization and filtering out the low frequency components

The white noise amplification can be described in the frequency domain by computing the spectral coefficients of s_k with respect to the (noise-free) orthonormal left singular vectors of the matrix A. Using (2.6), (2.4) gives the matrix identity for the *spectral components*

$$\Sigma^2 (U^T S_k) = (U^T S_k)(L_k L_k^T) + \alpha_k \beta_{k+1} (U^T s_{k+1}) e_k^T ,$$

with the last column of the last term equal to

$$\alpha_k \beta_{k+1} (U^T s_{k+1}) = \Sigma^2 (U^T s_k) - (\alpha_k^2 + \beta_k^2)(U^T s_k) - \alpha_{k-1} \beta_k (U^T s_{k-1}). \qquad (3.2)$$

Consider for example `shaw(400)` for which the vector s_1 has dominant components in the directions of several left singular vectors representing low frequencies, with a noticeable maximum in u_1. This is illustrated in Figs. 1 and 2 which shows the left singular vectors of A corresponding to $\sigma_1, \sigma_2, \ldots, \sigma_{10}$. An analogous assumption can be used for discrete ill-posed problems in general, see [20, Sect. 2]. At the first step

$$\alpha_1 \beta_2 (U^T s_2) = \Sigma^2 (U^T s_1) - \alpha_1^2 (U^T s_1) = (\Sigma^2 - \alpha_1^2 I) U^T s_1,$$

and, apart from multiplication by $\alpha_1 \beta_2$, the spectral components $U^T s_2$ are given by the spectral components of $U^T s_1$ scaled by $\Sigma^2 - \alpha_1^2 I$. The scaling acts as filtering which damps the dominant lowest frequency component in the direction of u_1. At the same time, the high frequency components are multiplied by the factors $\sigma_j^2/(\alpha_1 \beta_2) - \alpha_1/\beta_2$, where the first term is negligible for large j, and the absolute value of the second term is likely to be significantly larger than one (see the argument below). As a consequence, the relative level of high frequency noise is likely to be much larger

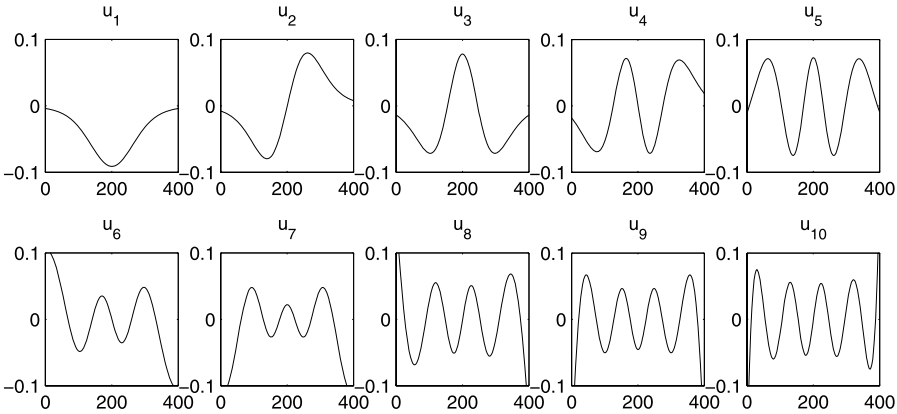

Fig. 2 The left singular vectors of A corresponding to $\sigma_1, \sigma_2, \ldots, \sigma_{10}$ for the problem `shaw(400)`

in s_2 than in s_1. At the general step, the Lanczos recurrence (3.2) can be rewritten in terms of the (Lanczos) polynomial in the diagonal matrix Σ^2,

$$U^T s_{k+1} = \varphi_k(\Sigma^2) U^T s_1, \tag{3.3}$$

where $\varphi_k(\lambda)$ is the kth orthonormal polynomial with respect to the Riemann-Stieltjes integral defined by the distribution function ω with the points of increase $\sigma_n^2, \ldots, \sigma_1^2$ (please recall that we assume, for simplicity of exposition, $\sigma_n > 0$) and the weights $|(s_1, u_n)|^2, \ldots, |(s_1, u_1)|^2$ respectively; see Sect. 2.[1] The roots of $\varphi_k(\lambda)$ are given by $(\theta_\ell^{(k)})^2$, $\ell = 1, \ldots, k$ (the Ritz values). Because of the dominance of the weights corresponding to the large nodes of ω, the large Ritz values $(\theta_k^{(k)})^2, (\theta_{k-1}^{(k)})^2, \ldots$ closely approximate $\sigma_1^2, \sigma_2^2, \ldots$; see Fig. 3. This creates the damping effect of $\varphi_k(\lambda)$ in the direction of the smooth components $(s_1, u_1)u_1$, $(s_1, u_2)u_2, \ldots$ of s_1. The constant term

$$\varphi_k(0) = \prod_{j=1}^{k} \frac{\alpha_j}{\beta_{j+1}} \tag{3.4}$$

then causes the relative amplification of the high frequency noise components present in s_1.

Summarizing, assume that the components of s_1 in the directions of the left singular vectors $u_1, u_2, u_3, \ldots$ decay faster than the associated singular values. Then, as a consequence of the orthogonalization process (3.2), $s_2, s_3, \ldots$ will have dominant components in the directions of the left singular vectors $u_2, u_3, \ldots$ respectively, with the relative levels of the high frequency noise components gradually increasing. Eventually, for some $k \equiv k_{\text{noise}}$, the vector $s_{k_{\text{noise}}+1}$ will have comparable com-

[1] It is worth to mention that the polynomials orthogonal with respect to a given Riemann-Stieltjes integral are called *Stieltjes polynomials*, and the associated three-term recurrence, given in its matrix form by Lanczos, was described by Stieltjes and others in the 19th century, cf. [13, Sect. 1.4, p. 80].

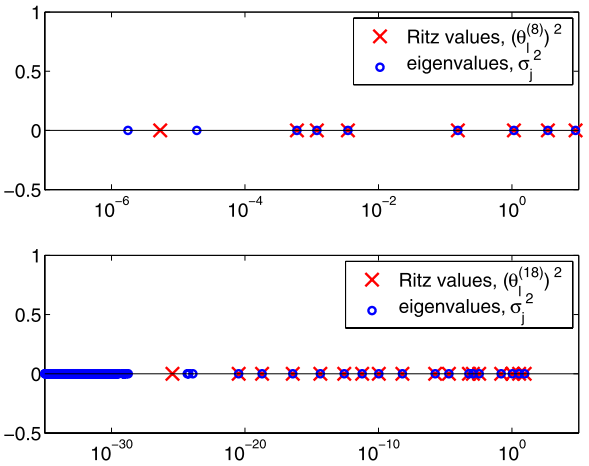

Fig. 3 The eigenvalues of AA^T (*circles*), and the Ritz values (*crosses*) corresponding to the 8th (*top plot*) and 18th (*bottom plot*) iterations respectively, for the problem `shaw(400)`. (All the *circles* in the *top plot* reappear as the nine rightmost circles in the bottom plot)

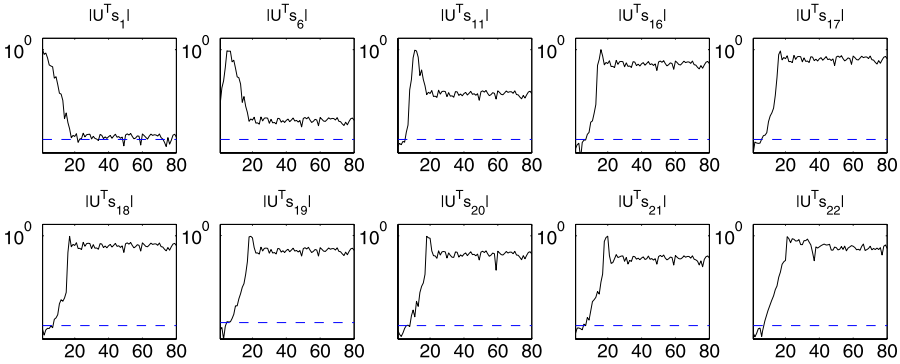

Fig. 4 The absolute values of the first 80 spectral components of the vectors s_k computed using the double reorthogonalized Golub-Kahan iterative bidiagonalization for the problem `shaw(400)` with the noise level $\delta_{\text{noise}} = 10^{-14}$. The *dashed line* represents the machine precision ε_M

ponents in practically all subspaces generated by singular vectors corresponding to $\sigma^2_{k_{\text{noise}}+1}, \sigma^2_{k_{\text{noise}}+2}, \ldots$, and this will reveal the noise. Figure 4 shows the absolute values of the components of several vectors $U^T s_k$ (the spectral components) for the problem `shaw(400)` with the noise level 10^{-14}. The absolute values of the components in the vector $U^T s_{18}$ corresponding to $\sigma_{18}, \sigma_{19}, \ldots$ are comparable, with no dominant maximum. Consequently, the noise is revealed in s_{18} for $k = k_{\text{noise}} = 17$. In the next step the high frequency components in s_{19} slightly decrease, and this is captured in Fig. 5. Components corresponding to low frequencies gradually decrease to the machine precision level ε_M (plotted by the dashed line). Figure 5 nicely illustrates why $k_{\text{noise}} = 17$ is the appropriate choice of the noise revealing iteration, i.e. why the noise is fully revealed in s_{18}.

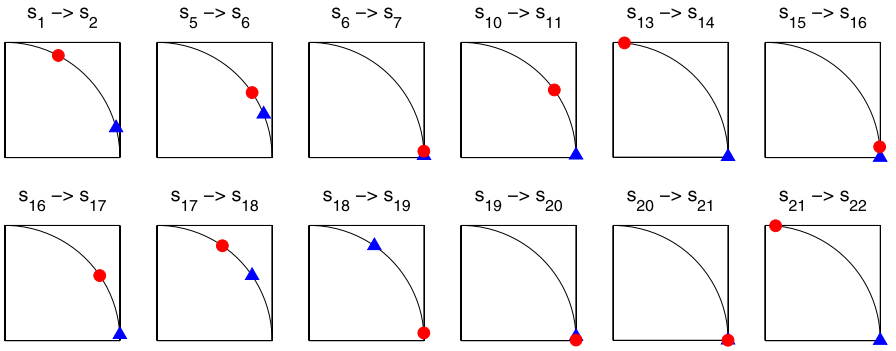

Fig. 5 Each graph shows for a given k the norms of the components of two consecutive vectors s_k (*the triangle*) and s_{k+1} (*the circle*) in the subspace span$\{u_1, \ldots, u_{k+1}\}$, also called the signal subspace (*the horizontal axis*), and the subspace span$\{u_{k+2}, \ldots, u_n\}$, also called the noise subspace (*the vertical axis*), for the problem `shaw(400)` with the noise level $\delta_{\text{noise}} = 10^{-14}$. Since the vectors s_j are normalized, the points are on the unit circle. One can see that until the noise level is reached, each step of the Golub-Kahan iterative bidiagonalization algorithm moves the subsequent vector counterclockwise, i.e. each s_{k+1} is closer to the noise subspace (*the vertical axis*) than s_k. The vector s_{19}, on the other hand, moves significantly clockwise due to the fact that the noise is partially projected out, and therefore the noise is revealed in s_{18}. This justifies the choice of the noise revealing iteration $k_{\text{noise}} = 17$. We are grateful to an anonymous referee for suggesting this visualization of the noise propagation process

3.2 Smoothing and orthogonalization in the Golub-Kahan iterative bidiagonalization

In order to further illustrate the (high frequency) white noise amplification, we consider the decomposition of s_1 into the exact component s_1^{exact} and the noise component s_1^{noise}, $s_1 = s_1^{\text{exact}} + s_1^{\text{noise}}$. Then the second equation in (2.1) gives

$$\beta_2 s_2 = A w_1 - \alpha_1 (s_1^{\text{exact}} + s_1^{\text{noise}}),$$

where Aw_1 is smooth with the low frequency components of the noise negligible relatively to the low frequency components of the exact data. This justifies the following *definition* of s_2^{exact} and s_2^{noise},

$$\begin{aligned} \beta_2 s_2^{\text{exact}} &\equiv A w_1 - \alpha_1 s_1^{\text{exact}}, \\ \beta_2 s_2^{\text{noise}} &\equiv A w_1 - \alpha_1 s_1 - \beta_2 s_2^{\text{exact}} = -\alpha_1 s_1^{\text{noise}}, \end{aligned}$$

giving

$$\beta_2 (s_2^{\text{exact}} + s_2^{\text{noise}}) = A w_1 - \alpha_1 (s_1^{\text{exact}} + s_1^{\text{noise}}).$$

Analogously, for $k = 2, 3, \ldots$

$$\beta_{k+1} s_{k+1}^{\text{exact}} \equiv A w_k - \alpha_k s_k^{\text{exact}}, \tag{3.5}$$

$$\beta_{k+1} s_{k+1}^{\text{noise}} \equiv -\alpha_k s_k^{\text{noise}}, \tag{3.6}$$

$$s_{k+1} = s_{k+1}^{\text{exact}} + s_{k+1}^{\text{noise}}, \qquad \beta_{k+1} s_{k+1} = A w_k - \alpha_k s_k. \tag{3.7}$$

It should be understood that s_k^{exact} and s_k^{noise} do not represent *true* exact data and noise components of s_k, respectively. If the multiplication by AA^T represents a significant smoothing of the high frequency components, then one can however expect that while $\|s_k^{\text{exact}}\| \gg \|s_k^{\text{noise}}\|$, $\|s_k^{\text{exact}}\|$ is close to the norm of the true data component and $\|s_k^{\text{noise}}\|$ is close to the norm of the true noise component of s_k. Using (3.6),

$$s_{k+1}^{\text{noise}} \equiv -\frac{\alpha_k}{\beta_{k+1}} s_k^{\text{noise}} \equiv (-1)^k \rho_k^{-1} s_1^{\text{noise}} \tag{3.8}$$

where the cumulative amplification ratio

$$\rho_k^{-1} \equiv \prod_{j=1}^{k} \frac{\alpha_j}{\beta_{j+1}} = \varphi_k(0), \tag{3.9}$$

see (3.4), on average (rapidly) grows as k increases.

In order to justify the expectation that ρ_k^{-1} on average (rapidly) grows with increasing k, we rewrite the Golub-Kahan bidiagonalization (2.1) for the spectral components $U^T s_j$ and $V^T w_j$,

$$\alpha_1 (V^T w_1) = \Sigma (U^T s_1), \tag{3.10}$$

$$\beta_2 (U^T s_2) = \Sigma (V^T w_1) - \alpha_1 (U^T s_1), \tag{3.11}$$

and for $k = 2, 3, \ldots$

$$\alpha_k (V^T w_k) = \Sigma (U^T s_k) - \beta_k (V^T w_{k-1}), \tag{3.12}$$

$$\beta_{k+1} (U^T s_{k+1}) = \Sigma (V^T w_k) - \alpha_k (U^T s_k). \tag{3.13}$$

From (3.10) we see that $(V^T w_1)$ is dominated by the same components as $(U^T s_1)$, with the dominance even enlarged as a consequence of the scaling by Σ. In (3.11), however, $\Sigma (V^T w_1)$ is orthogonalized against $U^T s_1$ in order to get, after normalization, $U^T s_2$. This requires that the dominance in $\Sigma (V^T w_1)$ and $U^T s_1$ is cancelled out, otherwise the orthogonality $U^T s_2 \perp U^T s_1$ can not hold. If the dominance is significant, one can therefore expect $\beta_2 \ll \alpha_1$. An analogous argument can be applied to the general step (3.12) and (3.13). Since the dominance in $\Sigma (U^T s_k)$ and $(V^T w_{k-1})$ is shifted by one component, one can not expect a significant cancellation, and, as a consequence

$$\alpha_k \approx \beta_k$$

should roughly hold. On the other hand, the vectors $\Sigma (V^T w_k)$ and $U^T s_k$ do exhibit dominance in the direction of the same components. If this dominance is strong enough, then the required orthogonality of s_{k+1} and s_k can not be achieved without a significant cancellation, and

$$\beta_{k+1} \ll \alpha_k$$

can be expected, giving the large cumulative amplification ratio ρ_k^{-1} (small ρ_k) as k progresses. The process is illustrated in Fig. 6.

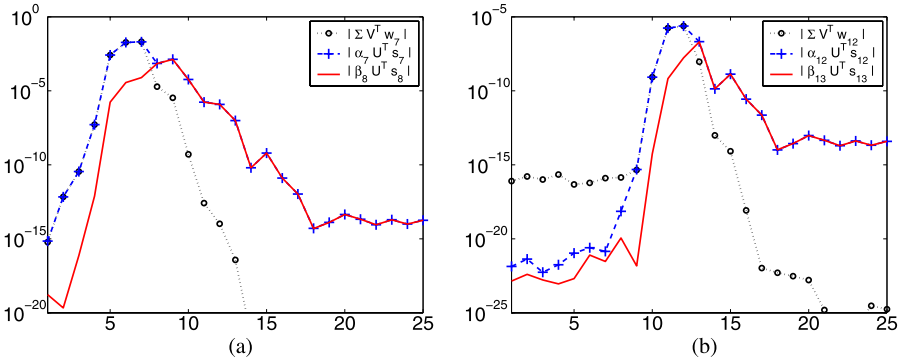

Fig. 6 The absolute values of the first 25 components of the vectors $\Sigma(V^T w_k)$, $\alpha_k(U^T s_k)$, and $\beta_{k+1}(U^T s_{k+1})$ for the problem `shaw(400)` for $k = 7$, $\beta_8/\alpha_7 = 0.0542$ (**a**), and $k = 12$, $\beta_{13}/\alpha_{12} = 0.0677$ (**b**)

Summarizing, starting with $s_1 = b/\beta_1$ contaminated by white noise, the Golub-Kahan iterative bidiagonalization algorithm applied to the discrete ill-posed problem $Ax \approx b$, where A represents *a smoothing operator*, amplifies the relative level of the noise in s_{k+1} as k increases with the cumulative amplification ratio ρ_k^{-1}. Note that the recurrence for the vectors w_k starts with the smoothed vector $w_1 = A^T s_1/\|A^T s_1\|$. Consequently, all vectors w_k are smoothed, and, in comparison to the vectors s_k, they do not contain significant information about the noise.

Remark Here we deal with mathematical properties (assuming exact arithmetic) and we do not need to consider specific implementation details. It should be noted, however, that the transfer of the high frequency noise from s_{k-1} and s_k to s_{k+1} is a local phenomenon which seemingly does not require preserving global orthogonality among the vectors $s_1, \ldots, s_k$. Since in practical implementations the local orthogonality among s_{k-1}, s_k and s_{k+1} is well preserved, cf. [39], one can intuitively expect that the preceding considerations are valid, with a small inaccuracy, also in practical finite precision computations. As we will demonstrate numerically, although such an intuitive argument is to some extent valid, the transfer of noise in finite precision computations is much more complicated; see Sect. 5.

The amplification of white noise in the Golub-Kahan iterative bidiagonalization is illustrated in Figs. 7–12. Figures 7–9 show individual components of several left bidiagonalization vectors s_k and their components s_k^{exact} and s_k^{noise}, computed using (3.5)–(3.7), for the problem `shaw(400)` with the noise level 10^{-14}. The Golub-Kahan iterative bidiagonalizations were performed with double reorthogonalization of the computed vector sequences, which ensures preserving orthogonality close to machine precision. Here we use the low noise level *on purpose* to illustrate the propagation of the noise through many steps of the iterative bidiagonalization. Results with larger noise levels will be presented below. As k increases, we observe the increasing oscillating pattern of s_k. For the vectors s_{17} and s_{18} the basic oscillating pattern (determined by s_k^{exact}) is strongly modulated by the high frequency noise; compare

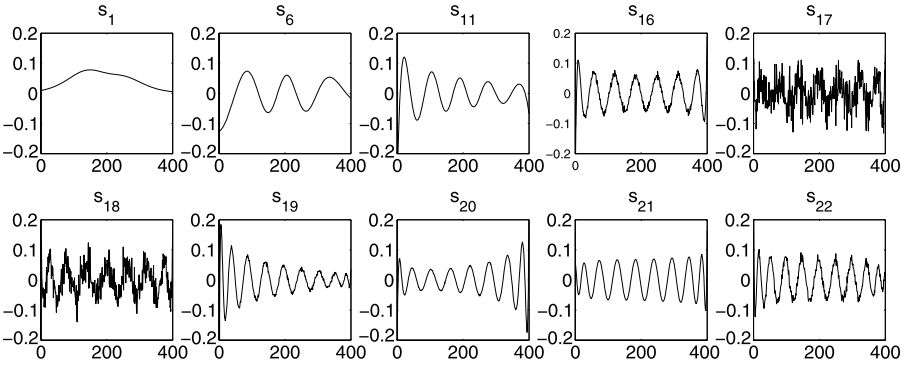

Fig. 7 Individual components of several left bidiagonalization vectors s_k computed using the double reorthogonalized Golub-Kahan iterative bidiagonalization for the problem `shaw(400)` with the noise level $\delta_{\text{noise}} = 10^{-14}$

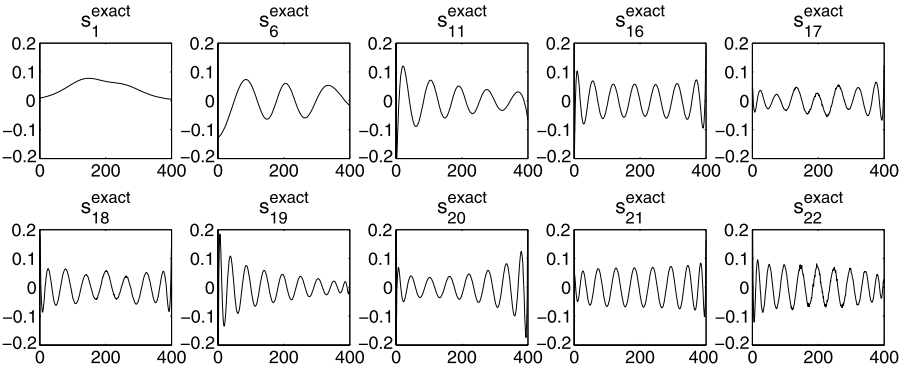

Fig. 8 Individual components of several vectors s_k^{exact} computed by (3.5) and (3.7) using the double reorthogonalized Golub-Kahan iterative bidiagonalization for the problem `shaw(400)` with the noise level $\delta_{\text{noise}} = 10^{-14}$

Figs. 7 and 8. It is interesting to observe that the level of the high frequency noise in the subsequent vector s_{19} is significantly lower than in s_{18}. As a consequence of smoothing and the orthogonality of s_{19} to s_{18} and s_{17}, the high frequency noise is partially projected out, cf. Fig. 5. Figure 10 presents the norms of the components s_k^{exact}, s_k^{noise} and the smallest singular value of the matrix S_k as k increases, which nicely complements the visual description in Figs. 7–9. Please notice that the decomposition of s_k into the components s_k^{exact} and s_k^{noise}, which should be close to the (unknown) exact data and noise components, is no longer relevant for $k \geq k_{\text{noise}}$. Finally, Fig. 11 gives the values of the normalization coefficients α_k, β_{k+1} and of their cumulative ratio ρ_k, and Fig. 12 depicts α_k, β_k and the cumulative ratio

$$\prod_{j=2}^{k} \frac{\beta_j}{\alpha_j}. \tag{3.14}$$

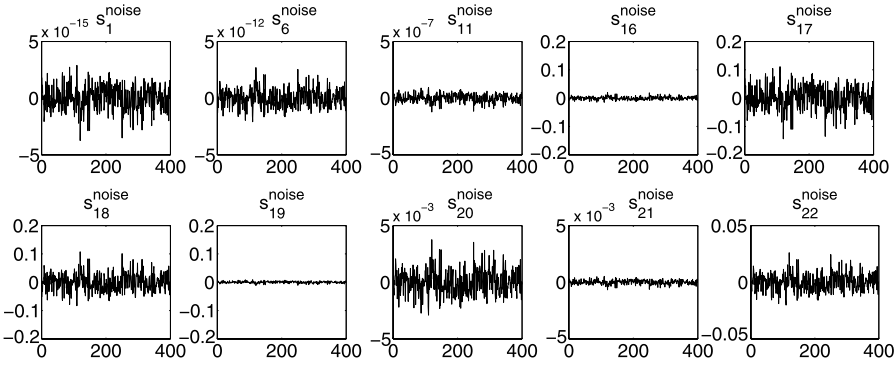

Fig. 9 Individual components of several vectors s_k^{noise} computed by (3.6) and (3.7) using the double reorthogonalized Golub-Kahan iterative bidiagonalization for the problem `shaw(400)` with the noise level $\delta_{\text{noise}} = 10^{-14}$. Note that the scale on the y axis is for different k different

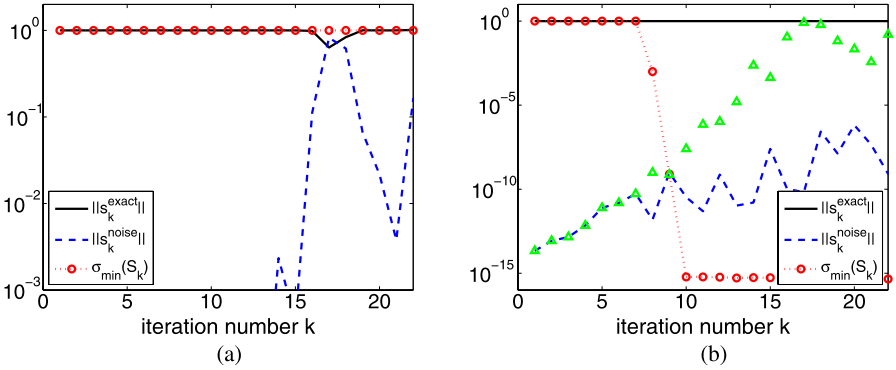

Fig. 10 The norms of $s_k^{\text{exact}}, s_k^{\text{noise}}$ and the smallest singular value of the matrix S_k for the problem `shaw(400)` with the noise level $\delta_{\text{noise}} = 10^{-14}$, computed by the Golub-Kahan iterative bidiagonalization with double reorthogonalization (**a**), and without reorthogonalization (**b**). For comparison, *the triangles* in (**b**) represent the norm of the components s_k^{noise} computed with double reorthogonalization. Please note the different vertical scales in (**a**) and (**b**)

The experimental results in Figs. 10(a), 11(a) and 12 correspond to the Golub-Kahan iterative bidiagonalization with double reorthogonalization, while Figs. 10(b) and 11(b) correspond to the results obtained without reorthogonalization. The differences will be discussed in Sect. 5 below.

The basis consisting of the left singular vectors of A is computationally expensive, and computing the corresponding components of the vectors s_k is practically infeasible. Analogous noise-revealing behavior can be observed in any other suitable orthonormal basis. Consider, e.g., the standard Fourier trigonometric basis

$$f_j(x) \equiv e^{(\frac{2\pi\iota}{n})jx}, \quad \text{for } j = 0, \pm 1, \ldots, \pm n, \tag{3.15}$$

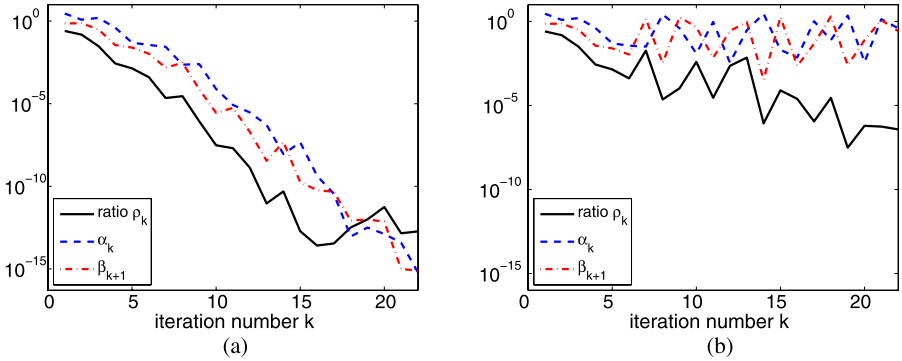

Fig. 11 Normalization coefficients α_k, β_{k+1} and their cumulated ratio ρ_k, see (3.9), for the problem `shaw(400)` with the noise level $\delta_{\text{noise}} = 10^{-14}$, computed by the Golub-Kahan iterative bidiagonalization with double reorthogonalization (**a**), and without reorthogonalization (**b**)

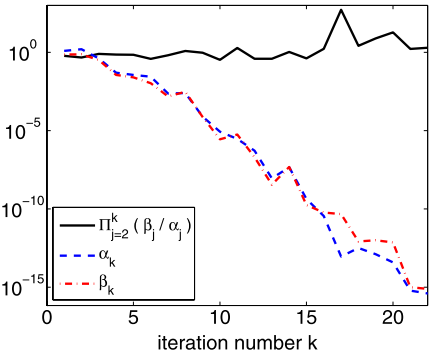

Fig. 12 Normalization coefficients α_k, β_k and their cumulated ratio (3.14) for the problem `shaw(400)` with the noise level $\delta_{\text{noise}} = 10^{-14}$, computed by the Golub-Kahan iterative bidiagonalization with double reorthogonalization

where ι is the imaginary unit. Fourier coefficients in this basis can be computed efficiently using the fast Fourier transform, [8, 9]. Figure 13 shows similar results as Fig. 4 computed for the basis (3.15) using the MATLAB function `fft`.

Similar behavior can be observed for other noise levels, as well as for noise-free problems. In this last case the "high frequency noise" is simply caused by local rounding errors from the finite precision arithmetic. Recall here that, because of the double reorthogonalization, the local rounding errors do not cause loss of orthogonality of the computed bidiagonalization vectors (the loss of orthogonality is kept close to the machine precision level). For more details see Sect. 5.

4 Determination of the noise level in the data and detection of the noise revealing iteration

As mentioned in Sect. 3, Fourier coefficients of the vectors s_k can be computed cheaply using, e.g., MATLAB function `fft`. Then some kind of statistical criteria can be used to identify the iteration where the Fourier coefficients $k+1, \ldots, n$ of the vector s_{k+1} are (on average) comparable, i.e. s_{k+1} resembles (except the first k

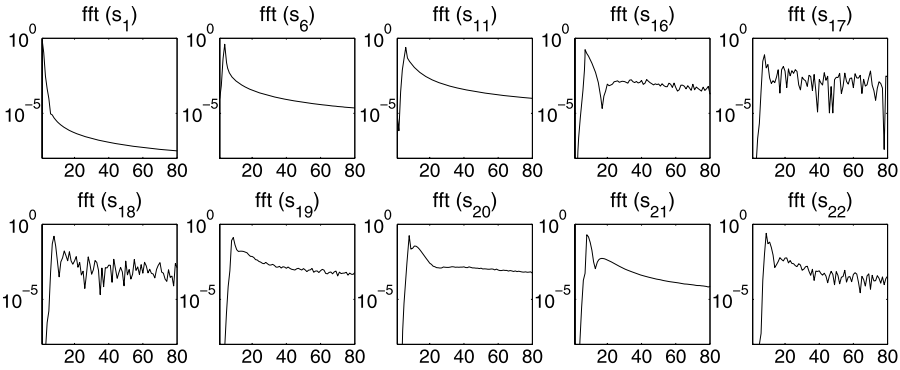

Fig. 13 The absolute values of the first 80 Fourier components of the vectors s_k, computed using the double reorthogonalized Golub-Kahan iterative bidiagonalization, in the trigonometric basis (3.15) for the problem `shaw(400)` with the noise level $\delta_{\text{noise}} = 10^{-14}$

components) white noise. Such statistical tools are used in a different context in [20, 46, 47] to decide whether a given residual vector is white-noise like. Here we propose a different approach which is more straightforward. The connection between the Golub-Kahan iterative bidiagonalization and the Lanczos tridiagonalization described in Sect. 2 suggests a new way of determining the noise level in the data without computing spectral or Fourier coefficients of s_k.

4.1 Automatic determination of the noise level based on approximation of the Riemann-Stieltjes distribution function

As recalled in Sect. 2, the Lanczos tridiagonalization generates a sequence of distribution functions $\omega^{(k)}$ with the nodes $(\theta_\ell^{(k)})^2$ and the weights $|(p_\ell^{(k)}, e_1)|^2$, $\ell = 1, \ldots, k$, that approximate the distribution function ω with the nodes σ_j^2 and the weights $|(b/\beta_1, u_j)|^2$, $j = n, n-1, \ldots, 1$; see [12, 31, 37, 52]. Depending on the noise level, for the smaller nodes of the distribution function ω the weights are completely dominated by noise, i.e., there exists an index J_{noise} such that for $j \geq J_{\text{noise}}$

$$|(b/\beta_1, u_j)|^2 \approx |(b^{\text{noise}}/\beta_1, u_j)|^2,$$

and the weight of the *set of the associated nodes* is given by

$$\delta^2 \equiv \sum_{j=J_{\text{noise}}}^{n} |(b^{\text{noise}}/\beta_1, u_j)|^2, \tag{4.1}$$

for illustration see Fig. 1. Since $\|b^{\text{noise}}\| \ll \|b^{\text{exact}}\|$, we can approximately write

$$\delta_{\text{noise}}^2 = \frac{\|b^{\text{noise}}\|^2}{\|b^{\text{exact}}\|^2} \approx \frac{\|b^{\text{noise}}\|^2}{\|b\|^2} = \sum_{j=1}^{n} |(b^{\text{noise}}/\beta_1, u_j)|^2. \tag{4.2}$$

With discrete ill-posed problems one can assume $J_{\text{noise}} \ll n$, and therefore, *assuming white noise*,

$$\delta^2 \approx \frac{n - J_{\text{noise}}}{n} \delta^2_{\text{noise}} \approx \delta^2_{\text{noise}}. \tag{4.3}$$

Please recall that the large nodes $\sigma_1^2, \sigma_2^2, \ldots$ are well-separated relative to the small ones and their weights on average decrease faster than $\sigma_1^2, \sigma_2^2, \ldots$. Therefore the distribution function $\omega^{(k)}$ approximates the distribution function ω in a special way—the large nodes essentially control the behavior of the early stages of the Lanczos tridiagonalization; see also Sect. 3, in particular Fig. 3. For $k = 1, 2, \ldots$ the large singular values of L_k become close to the largest singular values of A due to the dominance described above. At *any* iteration step the weight corresponding to $(\theta_1^{(k)})^2$ must be larger than the sum of weights of all σ_j^2 smaller than $(\theta_1^{(k)})^2$; see [12, 26], [11, Sect. 5.3]. As k increases, some $(\theta_1^{(k)})^2$ eventually approaches (or becomes smaller than) the node $\sigma^2_{J_{\text{noise}}}$, and its weight $|(p_1^{(k)}, e_1)|^2$ becomes

$$|(p_1^{(k)}, e_1)|^2 \approx \delta^2 \approx \delta^2_{\text{noise}}. \tag{4.4}$$

Using the notation of Sect. 3, this iteration step is equal to $k_{\text{noise}} + 1$. Indeed, $|(p_1^{(k)}, e_1)|$ is proportional to the noise level only *after* all smooth components of s_1 with the norms larger than the noise level are damped at the iteration step k_{noise}. The smallest nodes $(\theta_1^{(k_{\text{noise}}+1)})^2$, $(\theta_1^{(k_{\text{noise}}+2)})^2, \ldots$ strictly decrease due to the strict interlacing property of the Ritz values, but the corresponding weights (which also strictly decrease) remain in the subsequent steps approximately the same until the set of the smallest nodes $\{\sigma_n^2, \ldots, \sigma^2_{J_{\text{noise}}}\}$ is approximated by more than one Ritz value and the weight (4.1) is split, which happens when

$$(\theta_1^{(k)})^2 < \sigma_j^2 \quad \text{for } j \gg J_{\text{noise}}. \tag{4.5}$$

Summarizing, the weight $|(p_1^{(k)}, e_1)|^2$ corresponding to the smallest Ritz value $(\theta_1^{(k)})^2$ is strictly decreasing. At some iteration step $k = k_{\text{noise}} + 1$ it sharply starts to (almost) stagnate on the level close to the squared noise level δ^2_{noise}, see (4.3) and (4.4). In order to determine k_{noise} and to estimate δ_{noise}, it is therefore sufficient to monitor the first component of the left singular vector $|(p_1^{(k)}, e_1)|$ of the bidiagonal matrix L_k that is associated to its *smallest singular value* $\theta_1^{(k)}$; see (4.4). When it starts to stagnate,

$$\delta_{\text{noise}} \approx |(p_1^{(k_{\text{noise}}+1)}, e_1)|. \tag{4.6}$$

The stagnation can be detected by an automated procedure that does not rely on human interaction.

Figure 14 shows the positions of the leftmost points $[(\theta_1^{(k)})^2, |(p_1^{(k)}, e_1)|^2]$, $k = 1, 2, \ldots$ with respect to the distribution function $\omega(\lambda)$ (a) and $|(p_1^{(k)}, e_1)|$ as a function of k (b) for the problem `shaw(400)` with the noise level $\delta_{\text{noise}} = 10^{-14}$; (c) and (d) show the same for the noise level $\delta_{\text{noise}} = 10^{-4}$.

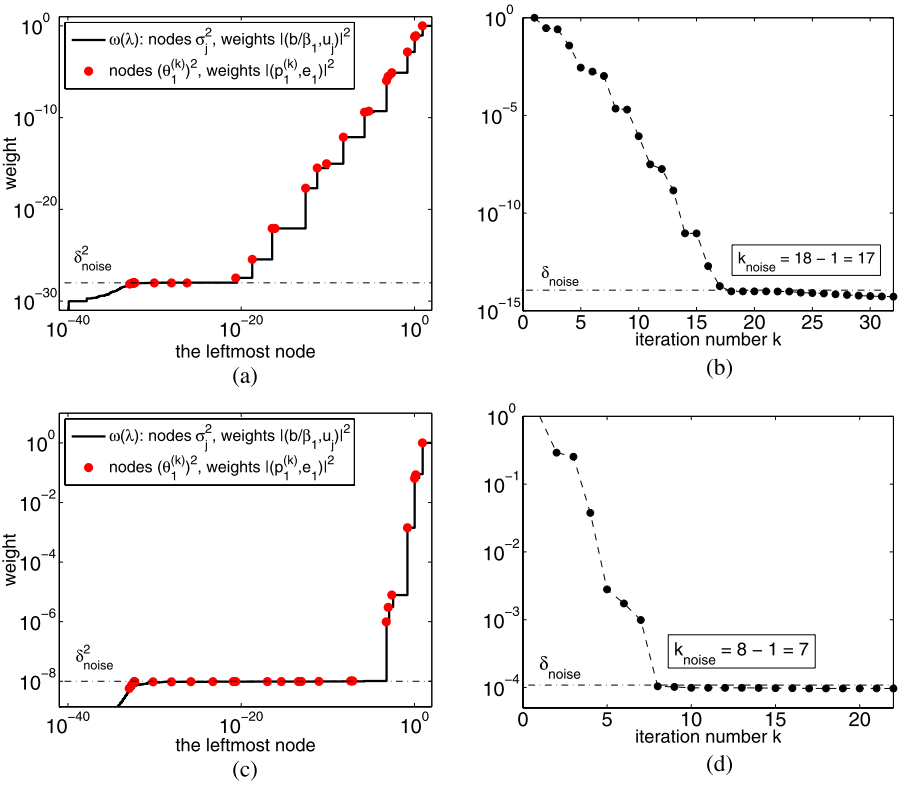

Fig. 14 The positions of the leftmost points $[(\theta_1^{(k)})^2, |(p_1^{(k)}, e_1)|^2]$, $k = 1, 2, \ldots$ with respect to the distribution function $\omega(\lambda)$ (**a**) and the absolute value of the first component $|(p_1^{(k)}, e_1)|$ of the left singular vector of L_k corresponding to its smallest singular value (**b**) for the problem `shaw(400)` with the noise level $\delta_{\text{noise}} = 10^{-14}$; (**c**) and (**d**) show the same for the noise level $\delta_{\text{noise}} = 10^{-4}$. The *horizontal dashed-dotted lines* represent the squared noise level δ^2_{noise} and the noise level δ_{noise}, respectively

4.2 An additional way of estimating the noise level in the data

Knowing the iteration k_{noise} where the noise is revealed in the Golub-Kahan iterative bidiagonalization, one can also estimate the level of the noise in the original data from the bidiagonalization coefficients. Here k_{noise} can be determined as described in the previous section. This section therefore describes an additional way of checking the estimate of δ_{noise} obtained in Sect. 4.1. Using (3.8),

$$s_{k+1}^{\text{noise}} = (-1)^k \rho_k^{-1} \frac{b^{\text{noise}}}{\|b\|},$$

where ρ_k is defined by (3.9), which gives

$$\frac{\|b^{\text{noise}}\|}{\|b\|} = \rho_k \|s_{k+1}^{\text{noise}}\|. \tag{4.7}$$

At the iteration step k_{noise} the noise is revealed, and we can assume that *on average* the components s_k^{exact} and s_k^{noise} are of about the same norm, cf. s_{18}^{exact} and s_{18}^{noise} in Figs. 8 and 9. Therefore

$$1 = \|s_{k_{\text{noise}}+1}\| = \|s^{\text{exact}}_{k_{\text{noise}}+1} + s^{\text{noise}}_{k_{\text{noise}}+1}\| \leq \|s^{\text{exact}}_{k_{\text{noise}}+1}\| + \|s^{\text{noise}}_{k_{\text{noise}}+1}\|$$

gives, considering that $s^{\text{exact}}_{k_{\text{noise}}+1}$ is much smoother than $s^{\text{noise}}_{k_{\text{noise}}+1}$ and thus there is not much cancellation between the individual components,

$$\|s^{\text{noise}}_{k_{\text{noise}}+1}\| \approx \frac{1}{2}.$$

Assuming $\|b^{\text{noise}}\| \ll \|b^{\text{exact}}\|$, the left part of (4.7) can be approximated, as in (4.2), by

$$\frac{\|b^{\text{noise}}\|}{\|b\|} = \frac{\|b^{\text{noise}}\|}{\|b^{\text{exact}} + b^{\text{noise}}\|} \approx \frac{\|b^{\text{noise}}\|}{\|b^{\text{exact}}\|},$$

which finally gives the estimate of the noise level in the original data

$$\delta_{\text{noise}} = \frac{\|b^{\text{noise}}\|}{\|b^{\text{exact}}\|} \approx \frac{1}{2}\rho_{k_{\text{noise}}}. \tag{4.8}$$

Table 1 shows the iterations k_{noise} (second row) and the corresponding estimates of the noise level $|(p_1^{(k_{\text{noise}}+1)}, e_1)|$ (third row) and $\frac{1}{2}\rho_{k_{\text{noise}}}$ (last row) for the problems `shaw(400)` and `ilaplace(100,1)` from the Regularization Toolbox [17] with different noise levels δ_{noise}. The estimates are average values computed using the set of 1000 randomly chosen sample vectors b^{noise}. In our experiments k_{noise} was

Table 1 Noise level in the data (first row), iteration k_{noise} where the noise is revealed (second row), the estimated noise level $|(p_1^{(k_{\text{noise}}+1)}, e_1)|$ (third row) and $\frac{1}{2}\rho_{k_{\text{noise}}}$ (last row), see (4.6) and (4.8) respectively, for problems `shaw(400)` and `ilaplace(100,1)`. The estimates represent average values computed using 1000 randomly chosen vectors b^{noise}

problem	`shaw(400)`				
noise level δ_{noise}	1×10^{-14}	1×10^{-10}	1×10^{-6}	1×10^{-4}	1×10^{-2}
k_{noise}	16	13	9	7	4
$\lvert(p_1^{(k_{\text{noise}}+1)}, e_1)\rvert$	1.80×10^{-14}	8.99×10^{-11}	1.31×10^{-6}	1.01×10^{-4}	1.03×10^{-2}
estimate $\frac{1}{2}\rho_{k_{\text{noise}}}$	8.93×10^{-15}	4.95×10^{-11}	6.55×10^{-7}	5.24×10^{-5}	5.55×10^{-3}
problem	`ilaplace(100,1)`				
noise level δ_{noise}	1×10^{-13}	1×10^{-10}	1×10^{-7}	1×10^{-2}	1×10^{-1}
k_{noise}	22	18.75	15.30	6.02	2
$\lvert(p_1^{(k_{\text{noise}}+1)}, e_1)\rvert$	9.12×10^{-14}	1.24×10^{-10}	1.34×10^{-7}	1.02×10^{-2}	1.11×10^{-1}
estimate $\frac{1}{2}\rho_{k_{\text{noise}}}$	4.77×10^{-14}	6.42×10^{-11}	7.11×10^{-8}	8.98×10^{-3}	5.57×10^{-2}

determined as the first iteration step k for which

$$\frac{|(p_1^{(k+1)}, e_1)|}{|(p_1^{(k+1+step)}, e_1)|} < \left(\frac{|(p_1^{(k)}, e_1)|}{|(p_1^{(k+1)}, e_1)|} \right)^{\zeta}, \tag{4.9}$$

where ζ was set to 0.5 and *step* was set to 3. We emphasize that (4.9) should be considered as *an example* of a possible automated stopping criteria. Its form has not resulted from an extensive research; such work is yet to be done and the results will most probably depend on particular application areas. For the problem `shaw` (4.9) works well except for the noise level $\delta_{\text{noise}} = 1 \times 10^{-14}$ where the automatically determined value $k_{\text{noise}} = 16$ is one less than the value determined in Sect. 3. The error is, however, negligible; see Fig. 14.

4.3 A comment on regularization and stopping criteria

When solving discrete ill-posed problems using hybrid methods based on the Golub-Kahan iterative bidiagonalization, for sufficiently large k, $k \ll n$, the absolute value of the first component of the left singular vector of L_k corresponding to its smallest singular value (almost) stagnates close to the *level of the noise present in the original data*. The beginning of the stagnation determines $|(p_1^{(k_{\text{noise}}+1)}, e_1)| \approx \delta_{\text{noise}}$ and the iteration k_{noise} when the noise level is revealed; see (4.6). Moreover, (4.8) gives the additional (secondary) noise level estimate.

Knowing the noise level, many different approaches can be applied in the subsequent steps, cf. [27, Sect. 3.2] including the straightforward application of the discrepancy principle [32, 33] (see also [18, Chap. 7.2, pp. 179–181]). It remains to determine which of these are effective in this context. Results in this direction will be reported elsewhere.

5 Noise propagation and the loss of orthogonality in the bidiagonalization vectors

To our knowledge, the effects of rounding errors have up to now not been thoroughly investigated in the literature on hybrid methods for solving discrete ill-posed problems, although they are sometimes acknowledged as a potential difficulty. As illustrated on Figs. 10 and 11, loss of orthogonality and subsequently loss of linear independence among the computed bidiagonalization vectors very significantly affect the propagation of the noise in the Golub-Kahan iterative bidiagonalization process. The corresponding Lanczos process without reorthogonalization computes multiple approximations to the well-separated squared large singular values. Consequently, the convergence of the Ritz values in the other parts of the spectrum can be significantly delayed, which affects convergence of the hybrid methods to the regularized solutions. The noise-revealing phenomenon is then complicated by the fact that computation of multiple approximations for the well-separated large singular values is connected with reappearance of the smooth components in the computed left bidiagonalization vectors s_k, which makes the propagation of the noise in the bidiagonalization process rather irregular. This is illustrated in Figs. 15–18, analogous to

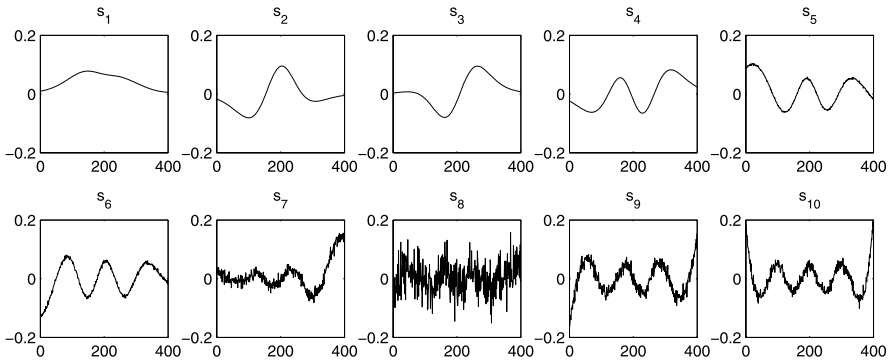

Fig. 15 Individual components of several left bidiagonalization vectors s_k computed by the double reorthogonalized Golub-Kahan iterative bidiagonalization for the problem `shaw(400)` with the noise level $\delta_{\text{noise}} = 10^{-4}$

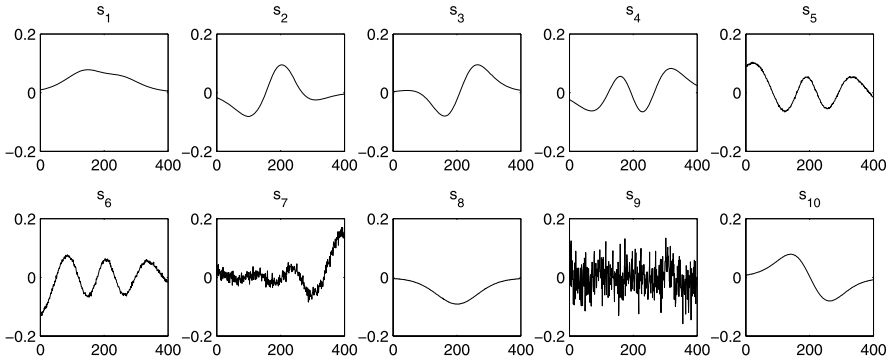

Fig. 16 Individual components of several left bidiagonalization vectors s_k computed by the Golub-Kahan iterative bidiagonalization without reorthogonalization for the problem `shaw(400)` with the noise level $\delta_{\text{noise}} = 10^{-4}$

Figs. 7, 10 and 11. Figures 15 and 16 show individual components of several left bidiagonalization vectors s_k computed by the Golub-Kahan iterative bidiagonalization *with double reorthogonalization and without reorthogonalization*, respectively, for the problem `shaw(400)` with the noise level $\delta_{\text{noise}} = 10^{-4}$. Figure 17 presents the norms of the components s_k^{exact}, s_k^{noise} and the smallest singular value of the matrix S_k as k increases. Finally, Fig. 18 gives the values of the normalization coefficients α_k, β_{k+1} and of their cumulated ratio ρ_k, see (3.9). We see that here the appearance of the smooth left bidiagonalization vector s_8 in the Golub-Kahan iterative bidiagonalization without reorthogonalization *delays* the revealing of the noise level by one iteration. In more difficult examples the loss of orthogonality among the bidiagonalization vectors can significantly delay the noise-revealing process. Nevertheless, as illustrated in Fig. 19 presenting results computed by the Golub-Kahan iterative bidiagonalization without reorthogonalization, the determination of the noise level based on monitoring the absolute value of the first component of the left singular vector

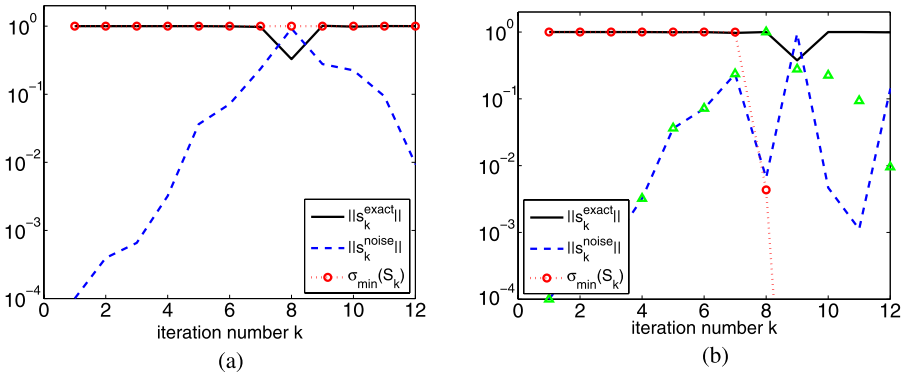

(a) (b)

Fig. 17 The norms of s_k^{exact}, s_k^{noise} and the smallest singular value of the matrix S_k for the problem `shaw(400)` with the noise level $\delta_{\text{noise}} = 10^{-4}$, computed by the Golub-Kahan iterative bidiagonalization with double reorthogonalization (**a**), and without reorthogonalization (**b**). For comparison, *the triangles* in (**b**) represent the norm of the component s_k^{noise} computed with double reorthogonalization

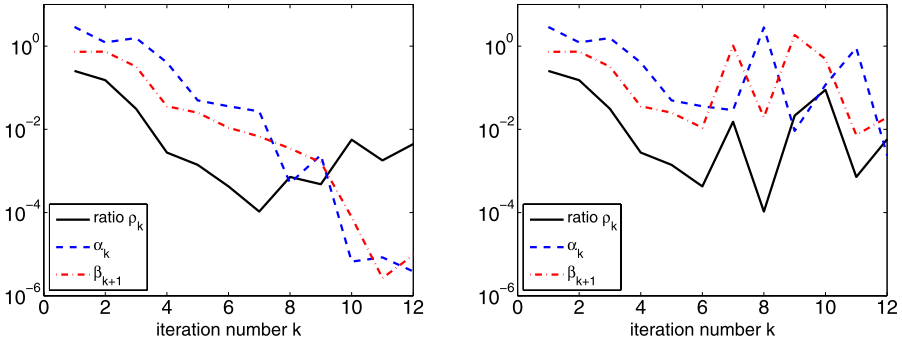

(a) (b)

Fig. 18 Normalization coefficients α_k, β_{k+1} and their cumulated ratio ρ_k, see (3.9), for the problem `shaw(400)` with the noise level $\delta_{\text{noise}} = 10^{-4}$, computed by the Golub-Kahan iterative bidiagonalization with double reorthogonalization (**a**), and without reorthogonalization (**b**)

of L_k corresponding to its smallest singular value still works. The k_{noise} determined from Fig. 19(a) is, however, shifted by one step in comparison to Fig. 16, where $k_{\text{noise}} = 8$.

It is worth noting that the propagation of noise in the Golub-Kahan iterative bidiagonalization for discrete ill-posed problems is different from the propagation of the elementary finite precision rounding errors which are present in each iteration. An investigation of this interesting topic is, however, outside the scope of this paper.

6 Conclusion

This paper considers discrete ill-posed problems represented by the linear system (1.1) with the following properties:

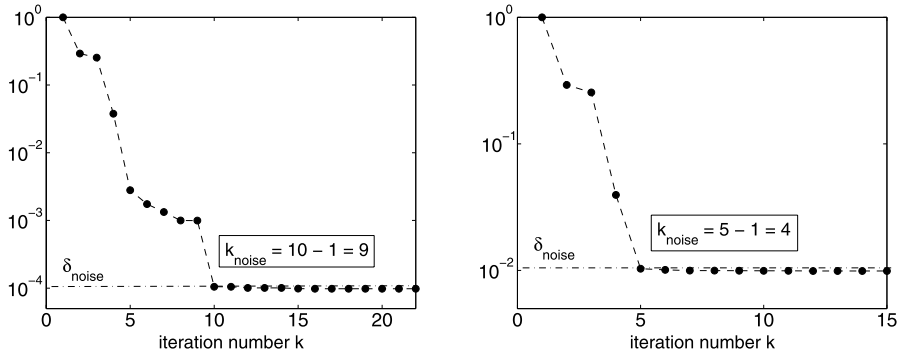

Fig. 19 The absolute value of the first component $|(p_1^{(k)}, e_1)|$ of the left singular vector of L_k corresponding to its smallest singular value for the problem `shaw(400)` with the noise level $\delta_{\text{noise}} = 10^{-4}$ (**a**) and 10^{-2} (**b**). The matrix L_k was computed using the Golub-Kahan iterative bidiagonalization without reorthogonalization. The *horizontal dashed-dotted line* represents the noise level δ_{noise}

- the matrices A, A^T, AA^T have a smoothing property;
- the left singular vectors u_j of A represent increasing frequencies as j increases;
- related to the last point, the system satisfies the discrete Picard condition; on average, the absolute value of the projections of the exact right-hand side b^{exact} to the left singular subspaces of A decays faster than the corresponding singular values;
- the noise component b^{noise} in the right-hand side represents white noise.

We showed that for this class of ill-posed problems it is possible to identify the iteration when the noise present in the data begins to propagate significantly to the projected problem computed by the Golub-Kahan iterative bidiagonalization. The unknown level of the noise in the original data can be estimated at a negligible cost from the absolute value of the first component of the left singular vector of the bidiagonal matrix of the projected system corresponding to its smallest singular value. This estimate is reliable and accurate. It can also be subsequently compared with the secondary estimate which uses the computed bidiagonalization coefficients. The estimated noise level can be used for construction of efficient stopping criteria based on many different approaches. We emphasize that throughout the paper the assumption on the white noise character of b^{noise} is substantial. Possible generalizations to cases with colored noise of varying degrees of dispersion, but dominated by high frequencies, need further investigation.

It is worth recognizing that if any of the assumptions made is not satisfied, then the presented noise revealing techniques may not succeed. Figures 20 and 21 show individual components of several left bidiagonalization vectors s_k computed by the double reorthogonalized Golub-Kahan iterative bidiagonalization, and their components in the basis of the left singular vectors of A, for the problem `phillips(256)` from the Regularization Toolbox [17] with the noise level 10^{-6}. Even though we observe an increase of the high frequency components in the vectors $s_1, s_2, \ldots$, it is not possible to identify the noise revealing iteration k_{noise} so clearly as proposed in Sect. 3. The reason is that in `phillips(256)` the discrete Picard condition is par-

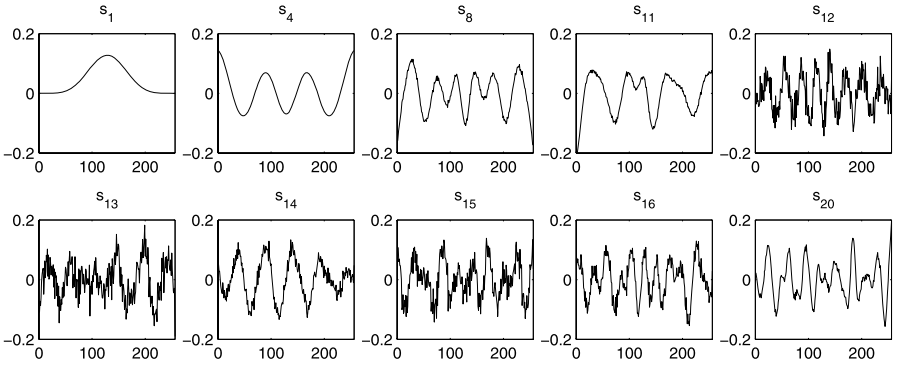

Fig. 20 Individual components of several left bidiagonalization vectors s_k computed using the double reorthogonalized Golub-Kahan iterative bidiagonalization for the problem `phillips(256)` with the noise level $\delta_{\text{noise}} = 10^{-6}$

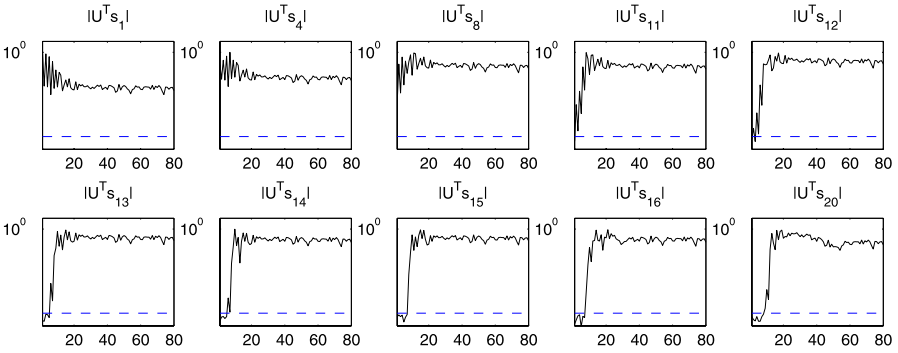

Fig. 21 The absolute values of the first 80 spectral components of the vectors s_k computed using the double reorthogonalized Golub-Kahan iterative bidiagonalization for the problem `phillips(256)` with the noise level $\delta_{\text{noise}} = 10^{-6}$. The *dashed line* represents machine precision ε_M

tially violated due to oscillations of the absolute value of the individual components. The technique presented in [20, 46] based on the Fourier analysis of the residual vectors works well also for this problem.

In this paper we have proposed noise revealing tools which may have potentially wide application for the solution of discrete ill-posed problems. The numerical experiments presented in this paper report only results obtained for the problem `shaw` (and to some extent also `ilaplace` and `phillips`) from the Regularization Toolbox [17]. Further experiments were performed with `ilaplace`, `phillips`, `deriv2`, and, without reorthogonalization, with ODF (here the matrix is rectangular of dimensions 5290×3375; see [21]) and with image deblurring examples Elephant (square problem of dimension 152280) and Barbara (two square problems of dimensions 65536 and 262144), cf. http://www.cs.cas.cz/krylov, section 'Software'. Based on the results, which can be found on the given www page, we believe that the conclusions derived from our observations and their mathematical justification offered

in Sects. 3 and 4 will find wide applications subject to the assumptions given above. We have also formulated several open questions which will be the subject of further investigation. The paper presents a step in understanding the noise revealing and regularizing properties of the Golub-Kahan iterative bidiagonalization. An application of the methods presented here to large scale problems needs further work. Results will be reported elsewhere.

Throughout the paper we have assumed A square and nonsingular. An extension of the results presented in this paper to problems with rectangular and even rank deficient matrices would require an additional assumption on the norm of the component of the right side b in the nullspace of A^T.

Acknowledgements We wish to thank Petr Tichý for his advice concerning numerical experiments and Gerard Meurant for useful comments. We are indebted to Per Christian Hansen for providing his codes and test examples used in additional tests, to Chris Paige for suggestion several corrections and clarifications, and to two anonymous referees for the very insightful remarks and suggestions which has led to significant improvements of our presentation.

Appendix

The terminology "*Golub-Kahan iterative bidiagonalization*" used throughout the paper is worth a short explanation. In their seminal paper [14], Golub and Kahan proposed two approaches for orthogonal bidiagonalization of a given matrix. In the first approach the bidiagonalization is computed by a sequence of Householder transformations from the left and right of A. The second approach, called here the Golub-Kahan iterative bidiagonalization, is introduced in the paper by the following words (here we use, for consistency, our notation for the lower bidiagonalization of A, while in [14] the authors consider the upper bidiagonalization):

> An alternative approach to bidiagonalization of A is to generate the columns of S and W sequentially as is done by the Lanczos algorithm for tridiagonalizing a symmetric matrix. The equation
>
> $$AW = SL \quad \text{and} \quad S^T A = LW^T$$
>
> can be expanded in terms of columns s_i of S and w_i of W to yield ...,

where W, S and L denote the results of the full bidiagonalization; see [14, p. 210]. For computation of the singular values of the bidiagonal matrix their paper refers to the idea of computation of the eigenvalues of the augmented matrix, which leads (through the Lanczos algorithm) to computation of the eigenvalues of the symmetric tridiagonal matrix, with reference to [29, Chap. 3].

Golub and Kahan gave by their reference to Lanczos an example of fairness which should be appreciated and followed. But their iterative bidiagonalization algorithm can not be attributed, according to our opinion, to Lanczos. The iterative bidiagonalization was proposed in [14]. Though the term "Lanczos bidiagonalization" is widespread in a part of literature, we concur with the remaining literature, in particular with [41, 42], that it is appropriate to use the name Golub-Kahan iterative bidiagonalization.

References

1. Björck, Å.: A bidiagonalization algorithm for solving large sparse ill-posed systems of linear equations. BIT **28**, 659–670 (1988)
2. Björck, Å., Grimme, E., Van Dooren, P.: An implicit shift bidiagonalization algorithm for ill-posed systems. BIT **34**, 510–534 (1994)
3. Barlow, J.L., Bosner, N., Drmač, Z.: A new stable bidiagonal reduction algorithm. Linear Algebra Appl. **397**, 35–84 (2005)
4. Calvetti, D., Golub, G.H., Reichel, L.: Estimation of the L-curve via Lanczos bidiagonalization. BIT **39**, 603–619 (1999)
5. Calvetti, D., Morigi, S., Reichel, L., Sgallari, F.: Tikhonov regularization and the L-curve for large discrete ill-posed problems. J. Comput. Appl. Math. **123**, 423–446 (2000)
6. Calvetti, D., Reichel, L., Shuibi, A.: L-curve and curvature bounds for Tikhonov regularization. Numer. Algorithms **35**, 301–314 (2004)
7. Chung, J., Nagy, J.G., O'Leary, D.P.: A weighted GCV method for Lanczos hybrid regularization. Electron. Trans. Numer. Anal. **28**, 149–167 (2008)
8. Cooley, J.W., Tukey, J.W.: An algorithm for the machine computation of the complex Fourier series. Math. Comput. **19**, 297–301 (1965)
9. Duhamel, P., Vetterli, M.: Fast Fourier transforms: A tutorial review and a state of the art. Signal Process. **19**, 259–299 (1990)
10. Fierro, R.D., Golub, G.H., Hansen, P.C., O'Leary, D.P.: Regularization by truncated total least squares. SIAM J. Sci. Statist. Comput. **18**, 1225–1241 (1997)
11. Fischer, B.: Polynomial Based Iteration Methods for Symmetric Linear Systems. Wiley-Teubner Series Advances in Numerical Mathematics. Wiley, New York (1996)
12. Fischer, B., Freund, R.W.: On adaptive weighted polynomial preconditioning for Hermitian positive definite matrices. SIAM J. Sci. Comput. **15**, 408–426 (1994)
13. Gautschi, W.: Gauss-Christoffel Quadrature Formulae. In: E.B. Christoffel, Aachen/Monschau, 1979, pp. 72–147. Birkhäuser, Basel (1981)
14. Golub, G.H., Kahan, W.: Calculating the singular values and pseudo-inverse of a matrix. SIAM J. Numer. Anal. Ser. B **2**, 205–224 (1965)
15. Golub, G.H., Von Matt, U.: Generalized cross-validation for large scale problems. J. Comput. Graph. Stat. **6**, 1–34 (1997)
16. Hanke, M.: On Lanczos based methods for regularization of discrete ill-posed problems. BIT **41**, 1008–1018 (2001)
17. Hansen, P.C.: Regularization Tools—version 3.2 for MATLAB 6.0, a package for analysis and solution of discrete ill-posed problems, http://www2.imm.dtu.dk/~pch/Regutools/index.html
18. Hansen, P.C.: Rank-Deficient and Discrete Ill-Posed Problems, Numerical Aspects of Linear Inversion. SIAM, Philadelphia (1998)
19. Hansen, P.C., Jensen, T.K.: Noise propagation in regularizing iterations for image deblurring. ETNA **31**, 204–220 (2008)
20. Hansen, P.C., Kilmer, M.E., Kjeldsen, R.: Exploiting residual information in the parameter choice for discrete ill-posed problems. BIT **46**, 41–59 (2006)
21. Hansen, P.C., Sørensen, H.O., Sükösd, Z., Poulsen, H.F.: Reconstruction of single-grain orientation distribution functions for crystalline materials. SIAM J. Image Sci. **2**(2), 593–613 (2009)
22. Hestenes, M.R., Stiefel, E.: Methods of conjugate gradients for solving linear systems. J. Res. Nat. Bur. Stand. **49**, 409–435 (1952)
23. Hnětynková, I., Strakoš, Z.: Lanczos tridiagonalization and core problems. Linear Algebra Appl. **421**, 243–251 (2007)
24. Hnětynková, I., Plešinger, M., Strakoš, Z.: Lanczos tridiagonalization, Golub-Kahan bidiagonalization and core problem. PAMM **6**, 717–718 (2006)
25. Jensen, T.K., Hansen, P.C.: Iterative regularization with minimum-residual methods. BIT **47**, 103–120 (2007)
26. Karlin, S., Shapley, L.S.: Geometry of Moment Spaces. American Mathematical Society, Providence (1953)
27. Kilmer, M.E., O'Leary, D.P.: Choosing regularization parameters in iterative methods for ill-posed problems. SIAM J. Matrix Anal. Appl. **22**, 1204–1221 (2001)
28. Kilmer, M.E., Hansen, P.C., Español, M.I.: A projection-based approach to general form Tikhonov regularization. SIAM J. Sci. Comput. **29**, 315–330 (2006)

29. Lanczos, C.: Linear Differential Operators. Van Nostrand, London (1961)
30. Lanczos, C.: An iteration method for the solution of the eigenvalue problem of linear differential and integral operators. J. Res. Nat. Bur. Stand. **45**, 255–282 (1950)
31. Meurant, G., Strakoš, Z.: The Lanczos and conjugate gradient algorithms in finite precision arithmetic. Acta Numer. **15**, 471–542 (2006)
32. Morozov, V.A.: On the solution of functional equations by the method of regularization (in Russian). Sov. Math. Dokl. **7**, 414–417 (1966)
33. Morozov, V.A.: Methods for Solving Incorrectly Posed Problems. Springer, New York (1984)
34. Nguyen, N., Milanfar, P., Golub, G.H.: Efficient generalized cross-validation with applications to parametric image restoration and resolution enhancement. IEEE Trans. Image Proces. **10**, 1299–1308 (2001)
35. O'Leary, D.P.: Near-optimal parameters for Tikhonov and other regularization methods. SIAM J. Sci. Comput. **23**, 1161–1171 (2001)
36. O'Leary, D.P., Simmons, J.A.: A bidiagonalization-regularization procedure for large scale discretizations of ill-posed problems. SIAM J. Sci. Stat. Comput. **2**, 474–489 (1981)
37. O'Leary, D.P., Strakoš, Z., Tichý, P.: On sensitivity of Gauss-Christoffel quadrature. Numer. Math. **107**, 147–174 (2007)
38. Paige, C.C.: Bidiagonalization of matrices and solution of linear equations. SIAM J. Numer. Anal. **11**, 197–209 (1974)
39. Paige, C.C.: Accuracy and effectiveness of the Lanczos algorithm for the symmetric eigenproblem. Linear Algebra Appl. **34**, 235–258 (1980)
40. Paige, C.C.: A useful form of unitary matrix obtained from any sequence of unit 2-norm n-vectors. SIAM J. Matrix Anal. Appl. **31**, 565–583 (2009)
41. Paige, C.C., Saunders, M.A.: LSQR: An algorithm for sparse linear equations and sparse least squares. ACM Trans. Math. Softw. **8**, 43–71 (1982)
42. Paige, C.C., Saunders, M.A.: ALGORITHM 583 LSQR: Sparse Linear equations and least squares problems. ACM Trans. Math. Softw. **8**, 195–209 (1982)
43. Paige, C.C., Strakoš, Z.: Scaled total least squares fundamentals. Numer. Math. **91**, 117–146 (2002)
44. Paige, C.C., Strakoš, Z.: Unifying least squares, total least squares and data least squares. In: Van Huffel, S., Lemmerling, P. (eds.) Total Least Squares and Errors-in-Variables Modeling, pp. 25–34. Kluwer Academic, Dordrecht (2002)
45. Paige, C.C., Strakoš, Z.: Core problem in linear algebraic systems. SIAM J. Matrix Anal. Appl. **27**, 861–875 (2006)
46. Rust, B.W.: Parameter selection for constrained solutions to ill-posed problems. Comput. Sci. Stat. **32**, 333–347 (2000)
47. Rust, B.W., O'Leary, D.P.: Residual periodograms for choosing regularization parameters for ill-posed problems. Inverse Probl. **24** (2008). doi:10.1088/0266-5611/24/3/034005
48. Saunders, M.A.: Computing projections with LSQR. BIT **37**, 96–104 (1997)
49. Sima, D.M., Van Huffel, S.: Using core formulations for ill-posed linear systems. PAMM **5**, 795–796 (2005)
50. Sima, D.M.: Regularization techniques in model fitting and parameter estimation. Ph.D. thesis, Dept. of Electrical Engineering, Katholieke Universiteit Leuven (2006)
51. Simon, H.D., Zha, H.: Low-rank matrix approximation using the Lanczos bidiagonalization process with applications. SIAM J. Sci. Stat. Comput. **21**, 2257–2274 (2000)
52. Strakoš, Z.: Model reduction using the Vorobyev moment problem. Numer. Algorithms **51**, 363–376 (2009)