

Silvia Gazzola* and Paolo Novati

Some transpose-free CG-like solvers for nonsymmetric ill-posed problems

<https://doi.org/10.1515/jnma-2018-0107>

Received October 09, 2018; revised May 31, 2019; accepted November 20, 2019

Abstract: This paper introduces and analyzes an original class of Krylov subspace methods that provide an efficient alternative to many well-known conjugate-gradient-like (CG-like) Krylov solvers for square nonsymmetric linear systems arising from discretizations of inverse ill-posed problems. The main idea underlying the new methods is to consider some rank-deficient approximations of the transpose of the system matrix, obtained by running the (transpose-free) Arnoldi algorithm, and then apply some Krylov solvers to a formally right-preconditioned system of equations. Theoretical insight is given, and many numerical tests show that the new solvers outperform classical Arnoldi-based or CG-like methods in a variety of situations.

Keywords: Krylov subspaces, iterative regularization methods, transpose-free, CGLS, GMRES

Classification: 65F10, 65F22

1 Introduction

Let us consider a linear system of the form

$$Ax = b, \quad A \in \mathbb{R}^{N \times N} \quad (1.1)$$

coming from a suitable discretization of an inverse ill-posed problem. In this setting, the matrix A typically has ill-determined rank, i.e., when considering the singular value decomposition (SVD) of A , given by $A = U\Sigma V^T$, with $\Sigma = \text{diag}(\sigma_1, \dots, \sigma_N)$, the singular values $\sigma_i \geq \sigma_{i+1} > 0$, $i = 1, \dots, N - 1$, quickly decay and cluster at zero with no evident gap between two consecutive ones to indicate numerical rank. In particular, A is ill-conditioned. Moreover, the right-hand side vector in (1.1) is typically affected by some unknown noise e , i.e., $b = b^{\text{ex}} + e$, where b^{ex} is the unknown exact version of b . Our goal is to compute a meaningful approximation of the solution x^{ex} of the unknown noise-free linear system $Ax^{\text{ex}} = b^{\text{ex}}$ and, because of the ill-conditioning of A and the presence of the noise e , some kind of regularization should be applied to the available system (1.1) (see [14] for an overview). The truncated SVD (TSVD) is a well-established regularization method, which consists in replacing (1.1) by the least square problem $\min_{x \in \mathbb{R}^N} \|A_m x - b\|$, where $\|\cdot\|$ denotes the vectorial 2-norm (or, in the following, the induced matrix 2-norm) and

$$A_m = U_m^A \Sigma_m^A (V_m^A)^T, \quad U_m^A \in \mathbb{R}^{N \times m}, \quad \Sigma_m^A \in \mathbb{R}^{m \times m}, \quad V_m^A \in \mathbb{R}^{N \times m} \quad (1.2)$$

is the best rank- m approximation of A in the matrix 2-norm. Here U_m^A and V_m^A are obtained by taking the first m left and right singular vectors of A , respectively (i.e., the first m columns of U and V , respectively), and Σ_m^A is the diagonal matrix of the first m singular values of A . Since the (T)SVD is computationally expensive, it is not suitable to regularize large-scale and unstructured problems. Therefore, in this paper, we are particularly interested in iterative regularization methods, which compute an approximation of x^{ex} by leveraging the so-called ‘semi-convergence’ phenomenon, so that regularization is achieved by an early termination of the iterations. Iterative regularization methods typically require one matrix-vector product with A and/or A^T at each iteration, and therefore they can also be employed when the coefficient matrix A and/or A^T is not explic-

*Corresponding author: **Silvia Gazzola**, Department of Mathematical Sciences, University of Bath, UK. Email: s.gazzola@bath.ac.uk

Paolo Novati, Department of Mathematics and Geosciences, University of Trieste, Italy.

itly available. Many Krylov subspace methods are efficient iterative regularization methods, as they typically show good accuracy together with a fast initial convergence (see [8] and the references therein).

While the theoretical regularizing properties and the performances of some CG-like solvers, such as CG, CGLS, and CGNE, are well-understood (see, for instance, [11, 12, 14, 16]), the same is not true for the methods based on the Arnoldi algorithm. The authors of [3] prove that, under some assumptions on b^{ex} , GMRES equipped with a stopping rule based on the discrepancy principle (i.e, the residual) is a regularization method in a classical sense, meaning that x_m tends to x^{ex} as the noise e tends to 0. However, well-established arguments (see [16] and the references therein) suggests that GMRES (or even its range-restricted variant [2]) might fail if the problem is highly non-normal because the dominant right singular vectors of A , i.e., the first columns of V , are severely mixed in the GMRES approximate solutions; moreover, the approximation subspace generated by GMRES may fail to reproduce some relevant components of the solution x^{ex} (which may not be expressed as linear combination of the dominant left singular vectors of A). These situations arise, for instance, when considering image deblurring problems characterized by a highly non-symmetric blur.

We should stress that, among all the Krylov methods mentioned so far, GMRES is the only one that can handle a nonsymmetric linear system (1.1) when A^T is unavailable, or when its action is impractical to compute. For this reason, it is important to investigate ways of overcoming the shortcomings of GMRES, e.g., by defining a more appropriate approximation subspace for GMRES. A common way of achieving this is to incorporate some sort of ‘preconditioning’. For instance, the authors of [15] propose to incorporate into GMRES a ‘smoothing-norm preconditioner’, which can enforce some additional regularity into the solution (achieving an effect similar to Tikhonov regularization in general form). The authors of [5] propose to incorporate into GMRES a ‘reblurring preconditioner’ A' , which approximates A^T and is tailored for particular image deblurring problems: by doing so, the original system (1.1) is replaced by an equivalent one, whose coefficient matrix AA' (or $A'A$) well approximates AA^T (or $A^T A$, respectively), so that the problem is somewhat symmetrized. A similar approach is considered in [9], where several right preconditioners are devised and tested for a variety of applications. The use of an approximate A^T has also recently been considered in [6] for algebraic iterative reconstruction methods with applications to computed tomography. We emphasize that, here and in the following, the term ‘preconditioner’ is not used in a classical sense: indeed, these ‘preconditioners’ do not accelerate the ‘convergence’ of GMRES, but rather enforce some desirable properties into the solution subspace.

Following one of the strategies employed in [9], this paper studies an efficient and reliable strategy to symmetrize the coefficient matrix of system (1.1). More precisely, after the Krylov subspace

$$\mathcal{K}_m(A, b) = \text{span}\{b, Ab, \dots, A^{m-1}b\}$$

is generated by performing m iterations of the Arnoldi algorithm applied to (1.1), which only involves m matrix-vector products with the matrix A and an orthonormalization procedure, a rank- m matrix $A'_m \in \mathbb{R}^{N \times N}$ is formed by exploiting the quantities computed by the Arnoldi algorithm, in such a way that AA'_m is symmetric semi-positive definite. The original system (1.1) is then replaced by the following ‘preconditioned’, symmetric, rank-deficient problem, to be solved in the least squares sense

$$y_m \in \arg \min_{y \in \mathbb{R}^N} \|AA'_m y - b\|, \quad x_m = A'_m y_m. \quad (1.3)$$

Since in many situations A'_m is a good approximation of A_m^T , where A_m is defined as in (1.2), one can regard the system (1.3) as a rank-deficient symmetric version of (1.1), which is also a good approximation of the normal equations $AA^T y = b$ associated to (1.1), with $x = A^T y$. Furthermore, one can easily see that taking $A'_m = A_m^T$ and solving problem (1.3) is equivalent to computing a TSVD solution. Therefore, problem (1.3) can also be regarded to as a regularized version of (1.1). The least squares problem (1.3) can then be solved directly (by TSVD, as proposed in [9]) or iteratively (by a variety of Krylov subspace methods, as proposed in the present paper). With particular choices of the iterative solver for (1.3), transpose-free CGLS-like and transpose-free CGNE-like methods can be defined, whose accuracy will depend on the choice of m . We stress that the choice of A'_m considered in this paper is not designed for a specific application. We also remark that, as we shall see, once the Arnoldi algorithm is run to compute A'_m , the methods derived here to solve (1.3) have

a negligible computational cost, as the remaining operations can be arranged in such a way that only matrices of order m are involved: therefore, the overall cost of the new methods is essentially the cost of performing m iterations of the Arnoldi algorithm. Moreover, provided that m is sufficiently small, the memory requirements of the new methods are not demanding. Many numerical experiments show that the new methods can achieve accuracies comparable to CGLS and CGNE, without employing A^T .

This paper is organized as follows: Section 2 surveys some known properties of the Arnoldi algorithm, introduces the matrix A'_m appearing in (1.3), and derives some insightful theoretical results. Section 3 describes different algorithmic approaches for the solution of (1.3): some computational details are unfolded, and connections with CGLS and CGNE are explored. Section 4 displays the results of many numerical experiments, which compare the performances of the new class of solvers for (1.3) with traditional Krylov methods for (1.1). Finally, Section 5 draws some concluding remarks.

2 A transpose-free ‘symmetrization’ of the Arnoldi algorithm

The Arnoldi algorithm [23, §6.3] is a process for building an orthonormal basis of the Krylov subspace $\mathcal{K}_m(A, b)$: m steps of the Arnoldi algorithm lead to the following matrix decomposition

$$AW_m = W_{m+1}H_m \quad (2.1)$$

where $W_m = [w_1, \dots, w_m] \in \mathbb{R}^{N \times m}$ has orthonormal columns that span $\mathcal{K}_m(A, b)$, and $H_m \in \mathbb{R}^{(m+1) \times m}$ is an upper Hessenberg matrix. Moreover,

$$w_1 = \frac{b}{\|b\|}, \quad W_{m+1} = [W_m, w_{m+1}] \in \mathbb{R}^{N \times (m+1)}. \quad (2.2)$$

Throughout the paper we assume m to be sufficiently small, so that $\mathcal{K}_m(A, b)$ is of dimension m and decomposition (2.1) exists.

GMRES is arguably the most popular Krylov method based on the Arnoldi algorithm. By initially setting $x_0 = 0$, at the m th step of GMRES (see [23, §6.5]), one updates the decomposition (2.1), and an approximation x_m^{GMR} of the solution of the original linear system is obtained by taking

$$x_m^{\text{GMR}} = W_m s_m^{\text{GMR}}, \quad s_m^{\text{GMR}} = \arg \min_{s \in \mathbb{R}^m} \|H_m s - \|b\|e_1\| \quad (2.3)$$

where e_1 is the first canonical basis vector of \mathbb{R}^{m+1} . Thanks to (2.1), x_m^{GMR} enjoys the optimality property

$$x_m^{\text{GMR}} = \arg \min_{x_m \in \mathcal{K}_m(A, b)} \|Ax_m - b\|. \quad (2.4)$$

Let us now assume that m steps of the Arnoldi algorithm are performed, and let us consider the right-preconditioned system (1.3), where A'_m is the rank- m matrix defined as

$$A'_m = W_m H_m^T W_{m+1}^T = P_m A^{T \in \mathbb{R}^{N \times N}}, \quad P_m = W_m W_m^T \quad (2.5)$$

is the orthogonal projector onto $\mathcal{K}_m(A, b)$; see also [9]. Exploiting once again relation (2.1), one realizes that

$$AA'_m = AP_m A^T = AW_m H_m^T W_{m+1}^T = W_{m+1} H_m H_m^T W_{m+1}^T = C_m C_m^T \quad (2.6)$$

where $C_m = W_{m+1} H_m \in \mathbb{R}^{N \times m}$. Therefore, the least squares problem (1.3) can be reformulated as

$$y_m \in \arg \min_{y \in \mathbb{R}^N} \|AA'_m y - b\| = \arg \min_{y \in \mathbb{R}^N} \|C_m C_m^T y - b\|, \quad x_m = A'_m y_m. \quad (2.7)$$

Directly from definition (2.5), and recalling that $\text{range}(W_m) = \mathcal{K}_m(A, b)$, one can immediately see that $x_m \in \mathcal{K}_m(A, b)$, as computed in (2.7). Therefore, by (2.4), one has

$$\|Ax_m^{\text{GMR}} - b\| \leq \|Ax_m - b\|. \quad (2.8)$$

The following proposition sheds light on the links between the solutions of problems (1.1) and (2.7).

Proposition 2.1. Let $y_m \in \mathbb{R}^N$ be a solution of $C_m C_m^T y = b$. Then $x_m = W_m s_m \in \mathbb{R}^N$ solves $Ax = b$, where $s_m = H_m^T W_{m+1}^T y_m \in \mathbb{R}^m$. Conversely, let $x_m = W_m s_m$ be the solution of (1.1), where $s_m \in \mathbb{R}^m$. Then the system

$$H_m H_m^T t = \|b\| e_1 \quad (2.9)$$

has a solution $t_m \in \mathbb{R}^{m+1}$, and $y_m = W_{m+1} t_m \in \mathbb{R}^N$ is the minimal norm solution of $C_m C_m^T y = b$.

Proof. The first part obviously follows from (2.6), as

$$b = C_m C_m^T y_m = A A_m^T y_m = A W_m H_m^T W_{m+1}^T y_m = A W_m s_m = A x_m.$$

To prove the second part, one should first consider the Arnoldi decomposition (2.1), so that

$$b = A x_m = A W_m s_m = W_{m+1} H_m s_m$$

and, thanks to the first equality in (2.2),

$$H_m s_m = \|b\| e_1. \quad (2.10)$$

Now, consider the economy-size SVD of H_m , given by

$$H_m = U_m \Sigma_m V_m^T, \quad U_m \in \mathbb{R}^{(m+1) \times m}, \quad \Sigma_m \in \mathbb{R}^{m \times m}, \quad V_m \in \mathbb{R}^{m \times m} \quad (2.11)$$

and the associated full-size SVD, given by $H_m = U_m^f \Sigma_m^f V_m^T$, where

$$U_m^f = [U_m, u_{m+1}] \in \mathbb{R}^{(m+1) \times (m+1)}, \quad \Sigma_m^f = \begin{bmatrix} \Sigma_m \\ 0 \end{bmatrix} \in \mathbb{R}^{(m+1) \times m}.$$

Note that (2.10) holds if and only if $U_m^f \Sigma_m^f V_m^T s_m = U_m^f (U_m^f)^T (\|b\| e_1)$, which is equivalent to asking the last component of $(U_m^f)^T e_1$ (i.e., $u_{m+1}^T e_1$) to be zero. Then there exists a solution $t_m \in \mathbb{R}^{m+1}$ of (2.9), as

$$U_m^f \Sigma_m^f (\Sigma_m^f)^T \underbrace{(U_m^f)^T t}_{\hat{t} \in \mathbb{R}^{m+1}} = \|b\| e_1 \quad \text{implies} \quad \begin{bmatrix} \Sigma_m^2 & \\ & 0 \end{bmatrix} \hat{t} = \begin{bmatrix} U_m^T (\|b\| e_1) \\ 0 \end{bmatrix}.$$

At this point, each y such that $W_{m+1}^T y = t_m$ satisfies

$$H_m H_m^T W_{m+1}^T y = \|b\| e_1.$$

By multiplying both terms by W_{m+1} from the left, and by exploiting the first equality in (2.2), one obtains $W_{m+1} H_m H_m^T W_{m+1}^T y = b$, which, thanks to (2.6), can be rewritten as

$$C_m C_m^T y = b. \quad (2.12)$$

Therefore, the minimum norm solution y_m of (2.12) satisfies $W_{m+1}^T y = t_m$, and is obtained by computing $\tilde{t}_m \in \mathbb{R}^{m+1}$ such that $W_{m+1}^T W_{m+1} \tilde{t}_m = t_m$ and taking $y_m = W_{m+1} \tilde{t}_m$. Since $W_{m+1}^T W_{m+1} = I$, $\tilde{t}_m = t_m$, and $y_m = W_{m+1} t_m$. \square

Proposition 2.1 essentially states that solving (1.1) by an Arnoldi-based method is equivalent to solving (2.7). More specifically, whenever the solution of (1.1) can be computed by performing m steps of a solver based on the Arnoldi algorithm (such as GMRES), a minimal norm solution of (2.7) can be recovered by solving the projected symmetric semi-positive definite system (2.9). However, as explained in Section 1, when dealing with ill-posed systems one is not interested in fully solving (1.1) and (2.7), and an iterative solver should be stopped reasonably early. Because of this, in the next section we will derive a variety of approaches for regularizing problem (2.7). We also remark that, as emphasized in [16] and recalled in Section 1, the performance of GMRES as a regularization method can sometimes be unsatisfactory because the approximation subspace for the solution is unsuitable. Since the approximation subspaces for GMRES and for any method applied to (2.7) coincide because, in any case, $x_m \in \mathcal{K}_m(A, b)$, one may suspect the approximate solutions of (2.7) to be

affected by the same issue. As we shall see in the next section, the SVD mixing is somewhat damped in (2.7), depending on the chosen solver.

In the remaining part of this section we provide some motivations underlying the choice of (2.5), which are connected to the regularizing properties of the Arnoldi algorithm. Define

$$\widehat{U}_m = W_{m+1}U_m = [\widehat{u}_1, \dots, \widehat{u}_m] \in \mathbb{R}^{N \times m}, \quad \widehat{V}_m = W_mV_m = [\widehat{v}_1, \dots, \widehat{v}_m] \in \mathbb{R}^{N \times m} \quad (2.13)$$

where U_m and V_m are the matrices of the left and right singular vectors of H_m (2.11), respectively, and define

$$\widehat{A}_m = W_{m+1}H_mW_m^T = \widehat{U}_m\Sigma_m\widehat{V}_m^T \quad (\text{note that } \widehat{A}_m = (A'_m)^T). \quad (2.14)$$

One can easily show that the (T)SVD of \widehat{A}_m is given by $\widehat{U}_m\Sigma_m\widehat{V}_m^T$, and that the Moore-Penrose pseudo-inverse \widehat{A}_m^\dagger of \widehat{A}_m is the regularized inverse (as defined in [14, §4.4]) associated to the m th iteration of GMRES. Indeed, by exploiting relation (2.3), the first equality in (2.2), and the TSVD (2.14), one can write

$$x_m^{\text{GMR}} = W_m s_m^{\text{GMR}} = W_m H_m^\dagger (\|b\|e_1) = W_m H_m^\dagger W_{m+1}^T b = \widehat{V}_m \Sigma_m^{-1} \widehat{U}_m^T b = \widehat{A}_m^\dagger b.$$

In order for a (generic) regularization method to be successful, the regularized matrix should contain information about the dominant singular values of the original matrix A , and filter out the influence of the small ones. If \widehat{A}_m is a good regularized approximation of A , then using $\widehat{A}_m^T = A'_m$ to approximate A^T is meaningful.

If A is severely ill-conditioned, the authors of [8, 20] numerically show that \widehat{A}_m quickly inherits the spectral properties of A . In particular, the following relations hold for $k = 1, \dots, m$:

$$\begin{aligned} A\widehat{v}_k - \sigma_k^{(m)}\widehat{u}_k &= 0 \\ W_m^T(A^T\widehat{u}_k - \sigma_k^{(m)}\widehat{v}_k) &= 0 \end{aligned} \quad (2.15)$$

where $\sigma_k^{(m)}$, $k \leq m$, is the k th singular value of H_m . Moreover, working in a continuous setting and under the hypothesis that A is a Hilbert-Schmidt operator of infinite rank whose singular values form an ℓ_2 sequence (see [22, Chapter 2] for a background), in [19] it has been shown that

$$\|A^T\widehat{u}_k - \sigma_k^{(m)}\widehat{v}_k\| \rightarrow 0 \quad \text{as } m \rightarrow \infty \quad (2.16)$$

where the convergence rate is closely connected to the decay rate of the singular values of A . This property is inherited by the discrete case whenever A is a suitable discretization of a Hilbert-Schmidt operator. Note that this class of operators includes Fredholm integral operators of the first kind with L_2 kernels. As a consequence of (2.15) and (2.16), in many relevant situations the dominant singular values of A are well approximated by the singular values of H_m (see [8] for many numerical examples). Therefore, \widehat{A}_m as defined in (2.14) may represent a good regularized approximation of A for a variety of problems.

3 Solving the ‘preconditioned’ problems

This section proposes two different iterative techniques to solve the rank-deficient symmetric least squares problem (2.7), and therefore to compute a regularized solution of (1.1). Thanks to the definition of C_m , decomposition (2.1), and Proposition 2.1, one can rewrite (2.7) as

$$y_m \in \arg \min_{y \in \mathbb{R}^N} \|AA'_m y - b\| = W_{m+1} \arg \min_{t \in \mathbb{R}^{m+1}} \|H_m H_m^T t - \|b\|e_1\|. \quad (3.1)$$

By using the above reformulation, it is clear that solving system (2.7) does not require a significant computational overload with respect to solving system (1.1) by any standard Arnoldi-based method (such as GMRES). Indeed, once m iterations of the Arnoldi algorithm have been performed, with $m \ll N$, all the additional computations for solving (3.1) are executed in dimension m , so that the computational cost of any algorithm for (3.1) is dominated by the cost of the Arnoldi algorithm. Moreover, the rank- m preconditioner A'_m (2.5) can be

stored in factored form, in order to recover x_m (2.7). Denoting by t_m any solution of the projected least-squares problem at the right-hand side of (3.1), and letting $y_m = W_{m+1}t_m$ be as in (2.7), the residual associated to (1.1) can be conveniently monitored in reduced dimension as

$$\|b - Ax_m\| = \|b - AA'_m y_m\| = \|b - C_m C_m^T y_m\| = \|\|b\|e_1 - H_m H_m^T t_m\|.$$

Since the starting vector b of Krylov subspaces generated by the Arnoldi algorithm (2.1), (2.2) is affected by some noise, noisy components are retained in H_m and W_m , so that the vector t_m in (3.1) should be computed by applying some regularization to the (noisy) projected problem

$$\min_{t \in \mathbb{R}^{m+1}} \|H_m H_m^T t - \|b\|e_1\|. \quad (3.2)$$

The noise propagation may be somehow damped by working with a range-restricted approach that consists in using Ab instead of b as starting vector for the Arnoldi process [2], and the theory developed in the present paper can be easily rearranged to work in this setting.

Direct methods such as Tikhonov regularization or TSVD can be easily applied to (3.2), the latter being particularly meaningful, as suggested in [9], because $H_m H_m^T$ is rank-deficient. However, in this paper, we are interested in using an iterative approach for solving (2.7) or (3.2), once the dimension m has been fixed.

3.1 A transpose-free CGLS-like method

Consider computing an approximation $y_{m,k}$ of y_m in (2.7) by applying k iterations of the MINRES method, with starting vector $x_0 = 0$. This is equivalent to requiring

$$y_{m,k} \in \mathcal{K}_k(AA_{m'}, b), \quad b - AA_{m'} y_{m,k} \perp (AA_{m'})\mathcal{K}_k(AA_{m'}, b), \quad k \leq m. \quad (3.3)$$

The first condition in (3.3), together with (2.1) and the above relation, implies

$$x_{m,k} = W_m H_m^T W_{m+1}^T y_{m,k} = W_m W_m^T A^T y_{m,k} = P_m A^T y_{m,k} \quad (3.4)$$

so that

$$x_{m,k} \in P_m A^T \mathcal{K}_k(AP_m A^T, b) = \mathcal{K}_k(P_m A^T A, P_m A^T b).$$

Similarly, the second condition in (3.3) implies

$$b - AP_m A^T y_{m,k} \perp AP_m A^T \mathcal{K}_k(AP_m A^T, b)$$

and, thanks to (3.4), it can be equivalently rewritten as

$$b - Ax_{m,k} \perp AP_m A^T \mathcal{K}_k(AP_m A^T, b) = A \mathcal{K}_k(P_m A^T A, P_m A^T b).$$

We can summarize the above arguments in the following proposition.

Proposition 3.1. For any given $m \geq 1$ the sequence $\{x_{m,k}\}_{k \leq m}$ obtained by applying k steps of the MINRES method to problem (2.7) is the result of a Krylov method defined by

$$x_{m,k} \in \mathcal{K}_k(P_m A^T A, P_m A^T b), \quad b - Ax_{m,k} \perp A \mathcal{K}_k(P_m A^T A, P_m A^T b). \quad (3.5)$$

The above proposition has two important consequences. Firstly, thanks to a well-known characterization of projection methods (see [23, §5.2]), the residual $b - Ax_{m,k}$ in (3.5) has minimal norm among all the residuals $b - A\hat{x}_{m,k}$, with $\hat{x}_{m,k} \in \mathcal{K}_k(P_m A^T A, P_m A^T b)$. Secondly, recall that, through an implicit construction of the Krylov subspaces $\mathcal{K}_k(A^T A, A^T b)$, CGLS generates a sequence of approximate solutions $\{x_k^{\text{CGLS}}\}_{k \geq 1}$ of (1.1) such that

$$x_k^{\text{CGLS}} \in \mathcal{K}_k(A^T A, A^T b), \quad b - Ax_k^{\text{CGLS}} \perp A \mathcal{K}_k(A^T A, A^T b). \quad (3.6)$$

Instead of the approximation subspace $\mathcal{K}_m(A^T A, A^T b)$ considered in (3.6), the method (3.5) implicitly builds a Krylov subspace where the action of A^T is replaced by its projection $P_m A^T$ onto $\mathcal{K}_m(A, b)$. Therefore, if $P_m = I$, conditions (3.5) and (3.6) are equivalent. In this sense, the new method (3.5) can be regarded as a transpose-free variant of a CGLS-like method, and from now on it will be simply referred to as TF-CGLS; correspondingly, the vector $x_{m,k}$ in (3.5) will be denoted as $x_{m,k}^{\text{LS}}$.

Remark 3.1. The clear advantage of TF-CGLS over CGLS is that A^T is not required, since only the action of A is needed to initially generate W_{m+1} and H_m . The additional k MINRES iterations required by TF-CGLS to compute the solution of (3.1) can be performed on the projected problem (3.2) of order $m+1$. Each approximate solution $\{x_{m,k}\}_{k \leq m}$ belongs to $\mathcal{K}_m(A, b)$ (directly by (3.5) and by the definition of P_m in (2.5)). If $\mathcal{K}_m(A, b)$ well captures the features of the solution that we wish to recover, then multiplication by P_m does not spoil the approximation subspace. Provided that a meaningful regularized solution can be recovered by TSVD (i.e., the columns of V_m^A are a good basis for a regularized solution), this is eventually equivalent to requiring that \widehat{A}_m in (2.14) inherits the spectral properties of A (see relations (2.15) and (2.16)).

Remark 3.2. Hybrid regularization methods [21] consider additional direct regularization (such as TSVD) within each iteration of a regularizing iterative method. We claim that TF-CGLS can be somewhat regarded as a hybrid regularization method. Indeed, considering the first relation in (3.3), one can straightforwardly rewrite

$$y_{m,k} \in \mathcal{K}_k(AA'_m, b) = W_{m+1}\mathcal{K}_k(H_m H_m^T, \|b\|e_1)$$

so that

$$x_{m,k} \in (W_m H_m^T W_{m+1}^T) W_{m+1} \mathcal{K}_k(H_m H_m^T, \|b\|e_1) = W_m \mathcal{K}_k(H_m^T H_m, H_m^T \|b\|e_1)$$

or, equivalently,

$$x_{m,k} = W_m t_k, \quad t_k \in \mathcal{K}_k(H_m^T H_m, H_m^T \|b\|e_1). \quad (3.7)$$

Analogously, considering the second relation in (3.3) and exploiting (2.1), one gets

$$b - AA'_m y_{m,k} \perp (AA'_m) \mathcal{K}_k(AA'_m, b) = W_{m+1} H_m \mathcal{K}_k(H_m^T H_m, H_m^T \|b\|e_1)$$

so that

$$\|b\|e_1 - H_m H_m^T W_{m+1}^T y_{m,k} \perp \mathcal{K}_k(H_m H_m^T, H_m H_m^T \|b\|e_1).$$

Recalling that $W_m^T x_{m,k} = H_m^T W_{m+1}^T y_{m,k}$ (directly from (2.7)) and the definition of t_k in (3.7), one gets

$$\|b\|e_1 - H_m t_k \perp \mathcal{K}_k(H_m H_m^T, H_m H_m^T \|b\|e_1).$$

Therefore, the vector t_k is obtained by applying k steps of the CGLS method to the projected LS problem (2.3) associated to the GMRES method (see the characterization (3.6)). In other words, after performing m steps of the Arnoldi algorithm to build W_m (exactly as GMRES does), the CGLS method is employed to solve the projected LS problem in (2.3), with m fixed. Therefore, in a sequential way, one applies another iterative regularization method (i.e., additional iterative regularization) within a fixed iteration of an iterative regularization method.

Remark 3.3. As briefly mentioned in Sections 1 and 2, regularization methods based on the Arnoldi algorithm (2.1) may sometimes be ineffective because the components of the right singular vectors of A are mixed in the approximation subspace $\mathcal{K}_m(A, b)$. Since an approximate solution $x_m \in \mathcal{K}_m(A, b)$ is such that $V^T x_m \in \mathcal{K}_m(V^T U \Sigma, V^T b)$, the mixing is caused by the presence of the non-diagonal matrix $V^T U \Sigma$; see [16] for more details. Quantitatively, this phenomenon (and therefore the regularization properties of the considered methods) can be characterized by computing the distance of two relevant subspaces, namely:

$$\text{dist}(\text{span}(V_k^A), \text{span}(\widehat{W}_k)) = \|\widehat{W}_k^T (V_k^A)^\perp\|; \quad (3.8)$$

see [10, Chapter 2]. Here and in the following we consider the subspaces spanned by the columns of V_k^A (i.e., the first k right singular vectors of A ; see (1.2)) and \widehat{W}_k (i.e., k orthonormal vectors selected after linearly transforming the columns of W_m ; see below). Here we display a numerical example clearly showing that, while severe SVD mixing affects the basis vectors of the GMRES solution, the SVD components are somewhat unmixed in the TF-CGLS basis vectors, whose behavior is comparable to the CGLS ones. The same holds for the hybrid GMRES-TSVD basis vectors (where the projected problem (2.3) is regularized through TSVD). Analogously to [16], we consider the test problem `i_laplace(100)` from [13], and we add Gaussian white noise e




Fig. 1: Components of the first, third, and fifth columns of \widehat{W}_5 with respect to the right singular vector basis for the GMRES, TF-CGLS, and CGLS methods applied to the `i_laplace(100)` test problem. Lower rightmost frame: filter factors for the TF-CGLS (5 CGLS iterations) and GMRES-TSVD (5 components) methods.

to the data vector b , in such a way that the noise level $\widehat{\varepsilon} = \|e\|/\|b^{\text{ex}}\|$ is $5 \cdot 10^{-4}$. We consider, as an example, approximation subspaces of dimension 5, spanned by the orthonormal columns of a matrix $\widehat{W}_5 \in \mathbb{R}^{100 \times 5}$, associated to the GMRES, TF-CGLS, GMRES-TSVD, and CGLS methods. More specifically:

- GMRES: we first run 5 steps of the Arnoldi algorithm to generate $W_5^{\text{GMR}} \in \mathbb{R}^{100 \times 5}$ and $H_5 \in \mathbb{R}^{6 \times 5}$ as in (2.1), and we then compute the SVD of H_5 (2.11), whose right singular vector matrix is denoted by $V_5^{H_5} \in \mathbb{R}^{5 \times 5}$. We take $\widehat{W}_5 = W_5^{\text{GMR}} V_5^{H_5}$. For this example we have $\text{dist}(\text{span}(V_k^A), \text{span}(\widehat{W}_5)) = 9.9933 \cdot 10^{-1}$.
- TF-CGLS and GMRES-TSVD: we first run 40 steps of the Arnoldi algorithm to generate $W_{40}^{\text{GMR}} \in \mathbb{R}^{100 \times 40}$ and $H_{40} \in \mathbb{R}^{41 \times 40}$ as in (2.1); we then compute the SVD of H_{40} (2.11), and we consider truncation after 5 components. We denote the truncated right singular vector matrix by $V_5^{H_{40}} \in \mathbb{R}^{40 \times 5}$. We take $\widehat{W}_5 = W_{40}^{\text{GMR}} V_5^{H_{40}}$. This corresponds to taking only the first 5 basis vectors in the TF-CGLS approximate solution. For this example we have $\text{dist}(\text{span}(V_k^A), \text{span}(\widehat{W}_5)) = 2.1357 \cdot 10^{-5}$.
- CGLS: we first run 5 steps of the Arnoldi algorithm applied to $A^T A$, with starting vector $A^T b$ (though, in practice, this procedure is unadvisable and one should employ the Golub-Kahan bidiagonalization algorithm; see [23, §8.3]). In this way we generate $W_5^{\text{CG}} \in \mathbb{R}^{100 \times 5}$ with orthonormal columns, and $T_5 \in \mathbb{R}^{6 \times 5}$ tridiagonal. We then compute the SVD of T_5 , whose right singular vector matrix is denoted by $V_5^{T_5} \in \mathbb{R}^{5 \times 5}$. We take $\widehat{W}_5 = W_5^{\text{CG}} V_5^{T_5}$. For this example we have $\text{dist}(\text{span}(V_k^A), \text{span}(\widehat{W}_5)) = 7.4933 \cdot 10^{-2}$.

Figure 1 shows the absolute value of the first, third, and fifth column of \widehat{W}_5 expressed in terms of the right singular values of A , i.e., $V^T \widehat{W}_5$, for the GMRES, TF-CGLS, and CGLS methods.

It can be easily seen that, while the GMRES basis vectors have significant components along all the right singular vectors (i.e., severe ‘SVD mixing’ happens), the same is not true for TF-CGLS. Though the components of the normalized TF-CGLS basis vectors are often larger than the normalized CGLS ones, the components corresponding to the first singular values of A are clearly dominant (moreover: the component of i th basis vector along the i th right singular vector v_i seems to be the leading one). This phenomenon happens because the TF-CGLS (and GMRES-TSVD) approximate solution belong to the Krylov subspace $\mathcal{K}_{40}(A, b)$, which is much larger than the Krylov subspace $\mathcal{K}_5(A, b)$ used in standard GMRES: therefore $\mathcal{K}_{40}(A, b)$ contains much more spectral information on A than $\mathcal{K}_5(A, b)$, which is then appropriately filtered. Indeed, as explained in

Remark 3.2, once 40 Arnoldi steps have been performed, both TF-CGLS and GMRES-TSVD apply additional regularization (or filtering) on the projected least squares problem (2.3), so that

$$x_{40,5} = W_{40}^{\text{GMR}} V_{40}^{H_{40}} \Phi^{H_{40}} (\Sigma_{40}^{H_{40}})^{-1} (U_{40}^{H_{40}})^T (\|b\|e_1)$$

where $U_{40}^{H_{40}}$, $\Sigma_{40}^{H_{40}}$, and $V_{40}^{H_{40}}$ are the matrices appearing in the economy-size SVD of H_{40} , and $\Phi^{H_{40}}$ is a diagonal filtering matrix, whose elements are

$$\Phi_{i,i}^{H_{40}} = p_5(\sigma_i^{(40)}), \quad \Phi_{i,i}^{H_{40}} = \begin{cases} 1, & i = 1, \dots, 5 \\ 0 & \text{otherwise} \end{cases}$$

for TF-CGLS and GMRES-TSVD, respectively. Here p_5 is the polynomial of degree at most 4 associated to 5 CGLS iterations for the projected LS problem in (2.3). These filter factors are displayed in the lower rightmost frame of Fig. 1. Starting from an extended Krylov subspace, and being able to filter out the dominant singular components of the projected quantities in (2.3), both TF-CGLS and GMRES-TSVD build a solution subspace where the original SVD components of A are not as mixed as in the standard GMRES one.

3.2 A transpose-free CGNE-like method

Now consider computing an approximation $y_{m,k}$ of y_m in (2.7) by applying k iterations of the CG method, with starting vector $x_0 = 0$. This means that

$$y_{m,k} \in \mathcal{K}_k(AA'_m, b), \quad b - AA'_m y_{m,k} \perp \mathcal{K}_k(AA'_m, b), \quad k \leq m. \quad (3.9)$$

As done in (3.4), we can write $x_{m,k} = P_m A^T y_{m,k}$, so that the first condition in (3.9) can be rewritten as

$$x_{m,k} \in P_m A^T \mathcal{K}_k(AP_m A^T, b).$$

Moreover, the second condition in (3.9) leads to

$$\begin{aligned} Ax - AP_m A^T y_{m,k} &\perp \mathcal{K}_k(AP_m A^T, b) \\ x - P_m A^T y_{m,k} &\perp A^T \mathcal{K}_k(AP_m A^T, b) \\ x - x_{m,k} &\perp A^T \mathcal{K}_k(AP_m A^T, b). \end{aligned}$$

We can summarize the above arguments in the following proposition.

Proposition 3.2. For any given $m \geq 1$ the sequence $\{x_{m,k}\}_{k \leq m}$ obtained by applying k steps of the CG method to problem (2.7) is the result of a Krylov method defined by

$$x_{m,k} \in P_m A^T \mathcal{K}_k(AP_m A^T, b), \quad x - x_{m,k} \perp A^T \mathcal{K}_k(AP_m A^T, b). \quad (3.10)$$

The above proposition allows us to see how this approach relates to the well-known CGNE method, whose approximate solutions satisfy

$$x_k^{\text{CGNE}} \in A^T \mathcal{K}_k(AA^T, b), \quad x - x_k^{\text{CGNE}} \perp A^T \mathcal{K}_k(AA^T, b)$$

and are computed through an implicit construction of the Krylov subspaces $A^T \mathcal{K}_k(AA^T, b) = \mathcal{K}_k(A^T A, A^T b)$. Using similar arguments to the ones in Section 3.1, the new method (3.10) can be regarded as a transpose-free variant of a CGNE-like method, and from now on it will be simply referred to as TF-CGNE; correspondingly, the vector $x_{m,k}$ in (3.10) will be denoted as $x_{m,k}^{\text{NE}}$. Statements analogous to the ones explained in Remark 3.1 also hold for the TF-CGNE case.

We conclude this section by mentioning that, although CGNE is an iterative regularization method, in practice it may perform very badly. Indeed, if system (1.1) is inconsistent, CGNE does not even converge to $A^{-1}b$ (see [11, Chapter 4]). This means that, if the unperturbed system $Ax^{\text{ex}} = b^{\text{ex}}$ is consistent, only small perturbations e of b^{ex} are allowed, in such a way that b still belongs to the range of A . The same behavior is experimentally observed when performing the TF-CGNE method (see the numerical experiments in Section 4). Therefore, even if TF-CGNE is a potential alternative to TF-CGLS, the latter is to be preferred when dealing with ill-posed problems.

3.3 Setting the regularization parameters

The transpose-free CG-like methods described in Sections 3.1 and 3.2 (here briefly denoted by TF-CG) are, indeed, multi-parameter iterative methods, whose success depends on an accurate tuning of both the scalars m and k . It should be also remarked that the parameters m and k act sequentially, i.e., the former is used to compute A'_m in (2.5) and has to be set in advance of the TF-CG iterations; the latter is the number of TF-CG iterations; this is the main difference between some hybrid methods and the TF-CG-like methods. One way to fix m is to stop the preliminary iterations when

$$h_{m+1,m} < \tau \quad (3.11)$$

where $\tau > 0$ is a specified threshold. In this way, one stops when the Krylov subspace $\mathcal{K}_m(A, b)$ is not significantly expanded; see [8, 9, 20] for full motivations.

In principle, another natural approach to set m can be devised by monitoring the values of the quantity

$$\zeta_m = \|A^T A - P_m A^T A\|. \quad (3.12)$$

The smaller ζ_m , the nearer $A^T A$ to $P_m A^T A$, i.e., the more accurate the transpose-free approximation of $A^T A$. Since the approximate solutions x_m computed by TF-CG methods belong to the subspace $\mathcal{K}_m(P_m A^T A, P_m A^T b)$ (see the first relation in (3.5) and (3.10)), a small ζ_m also implies that the generated approximation subspaces are close to $\mathcal{K}_m(A^T A, A^T b)$. However, since one of the main motivations behind TF-CG methods is the lack of knowledge of A^T for some large-scale problems, the quantities ζ_m in (3.12) cannot be computed in practice. Therefore, after some simple derivations one can provide the following upper bound:

$$\zeta_m = \|(I - P_m)A^T A\| \leq \|A\| \cdot \|(I - P_m)A^T\| = \sigma_1 \|A(I - P_m)\|.$$

Though the above bound does not explicitly involve A^T , A^T is required by algorithms for computing σ_1 . Moreover, when dealing with large-scale problems, both σ_1 and $\|A(I - P_m)\|$ can be expensive to compute. Therefore, one should look for yet other alternative bounds. One can take $\sigma_1^{(m)}$, i.e., the largest singular value of H_m , as an approximation of σ_1 : indeed, thanks to the interlacing property of the singular values (see, for instance, [4, 7]), one can prove that $\sigma_1 \geq \sigma_1^{(\ell+1)} \geq \sigma_1^{(\ell)}$. Many numerical experiments available in the literature show that $\sigma_1^{(m)}$ quickly approaches σ_1 (see also [19]), so that

$$\zeta_m \leq \sigma_1 \|A(I - P_m)\| = (\sigma_1^{(m)} + \varepsilon_m) \|A - W_{m+1} H_m W_m^T\| \quad (3.13)$$

where $\varepsilon_m \rightarrow 0$ as m increases. Replacing σ_1 with $\sigma_1^{(m)}$ may not be meaningful when m is very small, but this is not the case when performing the first cycle of iterations of the Arnoldi algorithm for the TF-CG methods. Note that, to rewrite the second factor of the last equality in the above equation, we have also exploited (2.1). While some numerical experiments available in the literature (see [8]) suggest that the quantity $\|A - W_{m+1} H_m W_m^T\|$ decays similarly to the singular values of A , no theoretical results have been established, yet. Similarly to what happens in the TSVD case, one can consider $\|A - W_{m+1} H_m W_m^T\| \approx \sigma_{m+1}^{(m+1)}$. Even if the above estimate can be quite optimistic (see [9] for a discussion), experimentally it seems reliable to stop the first set of Arnoldi iterations when $\sigma_1^{(m)} \sigma_{m+1}^{(m+1)}$ is sufficiently small, i.e., one should stop as soon as

$$\sigma_1^{(m)} \sigma_{m+1}^{(m+1)} < \tau' \quad (3.14)$$

where $\tau' > 0$ is a specified threshold.

To choose the number k of additional iterations for the TF-CG methods, some standard parameter choice strategies can be used. For instance, if one has a good estimate of the noise level $\hat{\varepsilon}$, the discrepancy principle can be applied, i.e., the iterations can be stopped as soon as

$$\|b - Ax_{m,k}\| = \|b - C_m C_m^T y_{m,k}\| = \left\| \|b\| e_1 - H_m H_m^T z \right\| < \eta \hat{\varepsilon} \|b\| \quad (3.15)$$

where $\eta > 1$ is a safety factor. If $\hat{\varepsilon}$ is not known, one can resort to other classical parameter choice methods such as GCV and the L-curve (see [14, Chapter 7]). The TF-CG methods are summarized in Algorithm 1.

Algorithm 1 TF-CG methods

```

input  $A, b, \tau$  or  $\tau'$ , solver,  $\eta, \hat{\varepsilon}$ 
for  $m = 1, 2, \dots$ , until the stopping criterion (3.11) or (3.14) is satisfied do
  update the Arnoldi decomposition:  $AW_m = W_{m+1}H_m$ 
end for
for  $k = 1, 2, \dots$ , until (3.15) is satisfied do
  if solver is TF-CGLS then
    apply MINRES to the system  $H_m H_m^T t = \|b\|e_1$ , to get  $t_k$  (see Section 3.1)
  else if solver is TF-CGNE then
    apply CG to the system  $H_m H_m^T t = \|b\|e_1$ , to get  $t_k$  (see Section 3.2)
  end if
end for
take  $x_{m,k} = W_m H_m^T t_k$ 

```

4 Numerical experiments

This section shows the performances of the methods summarized in Algorithm 1 on a variety of test problems: comparisons with GMRES and, whenever possible, CGLS and CGNE, will be displayed. The quality of the computed solution is measured by the relative error $\|x - x^{\text{ex}}\|/\|x^{\text{ex}}\|$. A first set of experiments considers moderate-scale problems from [13], while a second set of experiments considers realistic large-scale problems arising in the framework of 2D image deblurring. Unless otherwise stated, the TF-CG-like methods are implemented using modified Gram-Schmidt orthogonalization. All the tests are performed running MATLAB R2016b.

4.1 The first set of experiments

We consider problems with a nonsymmetric coefficient matrix whose right-hand-side vector is affected by Gaussian white noise of level $\hat{\varepsilon} = 10^{-2}$. Since A^T , as well as the SVD of A , are easily available for these problems, the use of the TF-CG-like methods may appear meaningless in this setting: these experiments are nonetheless included to compare the behavior of the TF-CG-like, the CGLS, and the CGNE methods, and to test some theoretical estimates (such as (3.12)–(3.14)). For all the tests, the maximum allowed number of Arnoldi iterations (in the first cycle of iterations in Algorithm 1) is $m_{\max} = 40$, and $\eta = 1.01$. The values $\tau = 10^{-10}$ and $\tau' = 10^{-15}$ are chosen for the stopping criteria in (3.11) and (3.14), respectively.

1. **i_laplace**. Let us consider the severely ill-posed inverse Laplace transform of the function $f(t) = t^2 \exp(-t/2)$, discretized using the codes in [13] with $N = 128$, so that $\|A - A^T\|/\|A\| = 0.6922$. Figure 2 compares the GMRES, CGLS, CGNE, TF-CGLS, and TF-CGNE methods (for different choices of the stopping criterion for the first set of iterations). The stopping criteria (3.11) and (3.14) are satisfied after 14 and 13 Arnoldi iterations, respectively, and the TF-CG-like solutions are not affected by the different stopping criteria. Enlarged markers are used to highlight the iterations satisfying the discrepancy principle (3.15) for all the methods (so that, in the GMRES, CGLS and CGNE cases, the quantities $\|b - Ax_m\|$ are monitored); we note that both CGNE and TF-CGNE do not satisfy the discrepancy principle within the chosen range of iterations. The TF-CGLS method delivers a huge improvement over the standard GMRES method for this problem, and its behavior is very similar to the CGLS one (see Fig. 2a, where relative errors versus number of iterations are plotted). TF-CGLS seems also very robust with respect to ‘semi-convergence’. Figure 2b shows the relative errors versus number of iterations for the GMRES, CGNE, and TF-CGNE methods. Although all the methods tested in both Figs. 2a and 2b perform quite poorly, the performances of ‘minimal error’ methods are clearly much worse than the performances of ‘minimal residual’ methods (this agrees with the analysis performed in [11, Chapter 4]). Despite this, the TF-CGNE method is able to

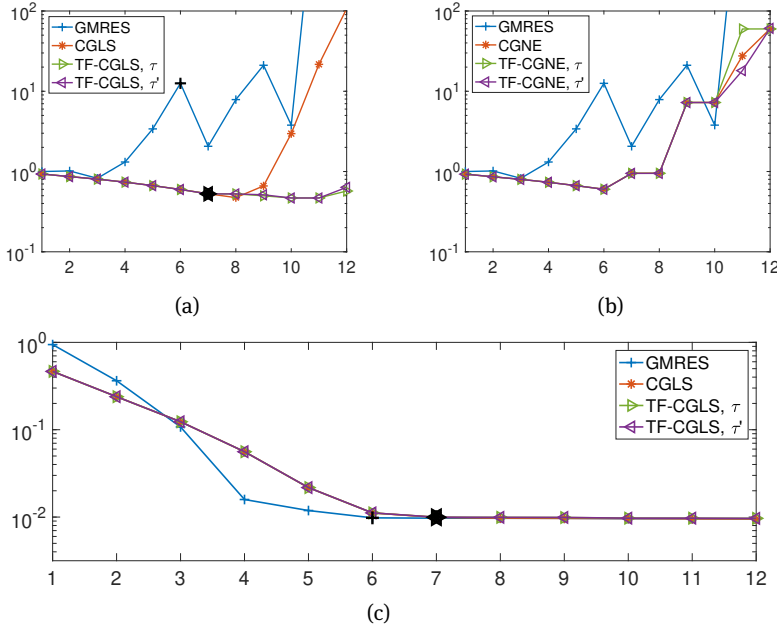


Fig. 2: Test problem $i_laplace$, with $f(t) = t^2 \exp(-t/2)$, $N = 128$, and $\hat{\varepsilon} = 10^{-2}$. (a) Relative errors versus number of iterations for GMRES, CGLS, and TF-CGLS (here and in the remaining frames, two different stopping criteria are considered). (b) Relative errors versus number of iterations for GMRES, CGNE, and TF-CGNE. (c) Relative residuals versus number of iterations for GMRES, CGLS, and TF-CGLS. Bigger black markers highlight the iterations satisfying the discrepancy principle.

reproduce quite faithfully the behavior of CGNE (in terms of relative errors). Further tests with CGNE and TF-CGNE will not be performed in the following experiments. Figure 2c displays relative residuals versus number of iterations for the GMRES, the CGLS, and the TF-CGLS methods. We can see that inequality (2.8) surely applies to the TF-CGLS case, i.e., when $x_m = x_{m,k} \in \mathcal{K}_m(A, b)$ (recall (3.4)): more precisely, once m has been set, $\|Ax_m^{\text{GMR}} - b\|$ is smaller than any $\|Ax_{m,k} - b\|$, for $k \leq m$ (but this does not imply any other relation between $\|Ax_\ell^{\text{GMR}} - b\|$, $\ell < m$, and $\|Ax_{m,k} - b\|$). Moreover, regarding the stopping criteria, a word of caution is mandatory: although looking at Fig. 2 it may seem that all the methods stop after 6 or 7 iterations, we should recall that the TF-CGLS iteration count refers to the second cycle in Algorithm 1 (after at most 14 Arnoldi iterations have been performed). Therefore, for this test problem, the computational cost of GMRES is dominated by 6 matrix-vector products with a matrix of size $N \times N$, while CGLS and TF-CGLS require 14 matrix-vector products with a matrix of size $N \times N$. We think that the additional (but still small) number of matrix-vector products required by TF-CGLS with respect to GMRES is tolerable if we consider the improved quality of the solution.

Figure 3a displays the behavior of the quantities (3.12)–(3.14) versus the number of Arnoldi iterations m . One can clearly see that (3.13) is a tight bound for the potentially unknown quantity (3.12). One also realizes that estimate (3.14) is indeed very optimistic, as anticipated in Section 3.3. We also emphasize that the behavior of the sequence $(\zeta_m)_{m \geq 1}$ is not monotonic because loss of orthogonality happens in the columns of the matrix W_m in (2.1) when modified Gram-Schmidt is employed. For this experiment, the TF-CG-like approximations are not very affected by the loss of orthogonality (as the stopping criteria for the first cycle of Arnoldi iterations prescribe to stop after 13 or 14 iterations, i.e., before a severe loss of orthogonality sets in). However, when a larger number of Arnoldi iterations is expected during the first cycle of Algorithm 1, one may consider the more numerically accurate (and more expensive) Householder-Arnoldi implementation, in order to reduce the effect of the loss of orthogonality (see [23, §6.3] for details). Figure 3b displays the solutions achieved by GMRES, CGLS, and TF-CGLS when the discrepancy principle is satisfied. While the CGLS and TF-CGLS solutions resemble the exact one, this is not the case for the GMRES solution, which is heavily unregularized (the values above 4 are truncated in the plot).

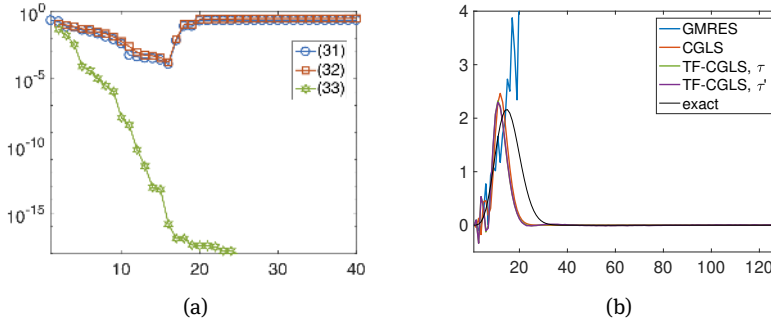


Fig. 3: Test problem `i_laplace`, with $f(t) = t^2 \exp(-t/2)$, $N = 128$, and $\varepsilon = 10^{-2}$. (a) Values of the quantities (3.12), (3.13), and (3.14) versus the number of Arnoldi iterations m . (b) Best approximations of $f(t)$ achieved by the GMRES, CGLS, and TF-CGLS methods.

2. **heat.** We consider a discretization of the inverse heat equation formulated as a Volterra integral equation of the first kind, as provided in [13], with the default settings. We choose $N = 256$, so that $\|A - A^T\|/\|A\| = 1.1249$; this problem can be regarded as numerically rank-deficient, with numerical rank equal to 250. Figure 4a shows the history of the relative errors for GMRES, CGLS, and TF-CGLS (with both the MGS and HH implementations). While CGLS delivers the best approximations, the quality of the TF-CGLS solutions is much better than the GMRES ones (which do not converge to $A^\dagger b$). Moreover, the TF-CGLS solutions obtained by the MGS and the HH implementations are comparable. Figure 4b displays the most accurate approximations obtained by each method: in the GMRES case, this is the zero solution (i.e., the initial guess); the TF-CGLS solution has slightly more oscillations than the CGLS one, and this shortcoming might be partially remedied by including additional (standard form) Tikhonov regularization within the TF-CGLS iteration (in a hybrid-like fashion, see [9]). Finally, Figures 4c and 4d display relative errors versus number of iterations for the TF-CGLS and the GMRES-TSVD methods. We also consider another variant of hybrid GMRES (denoted by GMRES-TSVD*) that expands the GMRES approximation subspace at each iteration and performs TSVD on (2.3) using the discrepancy principle. Figure 4c considers 40 Arnoldi iterations. Figure 4d potentially allows the full 256 Arnoldi iterations, and stopping criterion (3.14) is satisfied after 74 iterations. These plots clearly show that, in some situations, it may be convenient to adopt TF-CGLS rather than the more traditional hybrid GMRES methods, both in terms of accuracy and efficiency (in particular, the cost of computing the TF-CGLS solution is dominated by 40 matrix-vector products with A , while more than 100 matrix-vector products with A are required by GMRES-TSVD* to achieve a solution of similar quality).

Average values of relative errors and number of iterations are reported in Table 1, where the severely ill-conditioned test problem `baart` from [13] is also considered (in this case, $\|A - A^T\|/\|A\| = 0.6035$). The average is computed over 20 runs of each test problem, with different realizations of the random noise vector in the data. We can see that, except for `baart`, the values computed by the stopping criteria (3.12) and (3.14) mainly agree. Moreover, when the discrepancy principle is employed as an overall stopping criterion, the new TF-CGLS solver always delivers much better approximation than the GMRES and the GMRES-TSVD* method (though the latter often has a lower optimal relative error).

4.2 The second set of experiments

We consider 2D image restoration moderately ill-posed problems, where the available images are affected by a spatially invariant blur and Gaussian white noise. In this setting, given a point-spread function (PSF) that describes how a single pixel is deformed, a blurring process is modeled as a 2D convolution of the PSF and an exact discrete finite image $X^{\text{ex}} \in \mathbb{R}^{n \times n}$. Here and in the following, a PSF is represented as a 2D image $P \in \mathbb{R}^{q \times q}$,

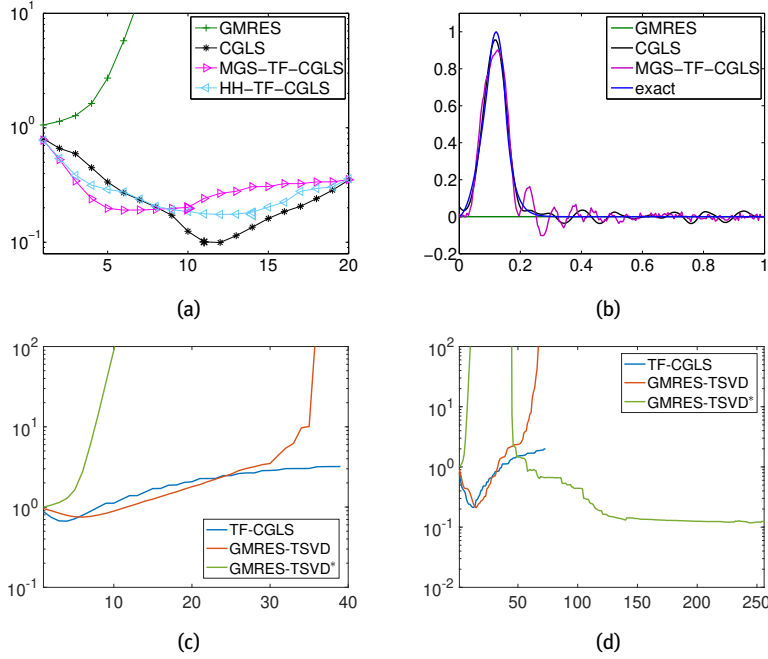


Fig. 4: Test problem heat, with $N = 256$, and $\hat{\varepsilon} = 10^{-2}$. (a) Relative error history for GMRES, CGLS, TF-CGLS with modified Gram-Schmidt implementation (MGS-TF-CGLS), and TF-CGLS with Householder implementation (HH-TF-CGLS). (b) Best approximations achieved by GMRES, CGS, and MGS-TF-CGLS. (c) Relative error history for GMRES-TSVD*, GMRES-TSVD, and MGS-TF-CGLS (with 40 Arnoldi iterations). (d) Relative error history for GMRES-TSVD* (with 256 Arnoldi iterations), GMRES-TSVD, and MGS-TF-CGLS (with 74 Arnoldi iterations).

Tab. 1: Average results over 20 runs of some moderate-scale test problems with nonsymmetric coefficient matrix from [13], with noise level $\hat{\varepsilon} = 10^{-2}$. The TF-CGLS relative error is the one attained when stopping criterion (3.14) is satisfied. The minimum attainable relative errors (opt) and the ones attained when the discrepancy principle is satisfied (DP) are reported for all the methods, together with the corresponding average number of iterations. The average number of iterations required by TF-CGLS to satisfy the stopping criteria (3.12) and (3.14) is also reported.

| | | Rel.Error (DP) | it (DP) | Rel.Error (opt) | it (opt) | (3.12) | (3.14) |
|------------------|-------------|------------------------|---------|------------------------|----------|--------|--------|
| i_laplace(128,1) | GMRES | $5.9919 \cdot 10^{-1}$ | 4.3 | $4.8728 \cdot 10^{-1}$ | 4.2 | — | — |
| | CGLS | $4.1821 \cdot 10^{-2}$ | 3.8 | $3.2919 \cdot 10^{-2}$ | 4.5 | — | — |
| | TF-CGLS | $4.1778 \cdot 10^{-2}$ | 3.8 | $3.2827 \cdot 10^{-2}$ | 4.5 | 14.5 | 13.0 |
| | GMRES-TSVD* | $5.9919 \cdot 10^{-1}$ | 4.3 | $3.8461 \cdot 10^{-2}$ | 11.6 | — | — |
| i_laplace(128,3) | GMRES | $5.1047 \cdot 10^0$ | 5.9 | $8.2388 \cdot 10^{-1}$ | 3.0 | — | — |
| | CGLS | $5.3116 \cdot 10^{-1}$ | 7.1 | $4.7355 \cdot 10^{-1}$ | 8.3 | — | — |
| | TF-CGLS | $5.2570 \cdot 10^{-1}$ | 7.2 | $4.3790 \cdot 10^{-1}$ | 11.2 | 14.5 | 13.1 |
| | GMRES-TSVD* | $5.1047 \cdot 10^0$ | 5.9 | $3.0080 \cdot 10^{-1}$ | 10.8 | — | — |
| baart(256) | GMRES | $5.6437 \cdot 10^{-1}$ | 3.0 | $3.1134 \cdot 10^{-1}$ | 4.0 | — | — |
| | CGLS | $1.6634 \cdot 10^{-1}$ | 3.0 | $1.5284 \cdot 10^{-1}$ | 3.4 | — | — |
| | TF-CGLS | $1.6650 \cdot 10^{-1}$ | 3.0 | $1.5302 \cdot 10^{-1}$ | 3.7 | 8.7 | 19.8 |
| | GMRES-TSVD* | $5.6437 \cdot 10^{-1}$ | 3.0 | $4.1127 \cdot 10^{-2}$ | 5.0 | — | — |
| heat(256) | GMRES | $4.4102 \cdot 10^7$ | 39.3 | $1.0000 \cdot 10^0$ | 1.0 | — | — |
| | CGLS | $1.0535 \cdot 10^{-1}$ | 11.1 | $9.2036 \cdot 10^{-2}$ | 13.1 | — | — |
| | TF-CGLS | $4.3834 \cdot 10^0$ | 32.9 | $6.2625 \cdot 10^{-1}$ | 4.1 | 40.0 | 40.0 |
| | GMRES-TSVD* | $4.4700 \cdot 10^7$ | 39.5 | $9.0833 \cdot 10^{-1}$ | 7.0 | — | — |




Fig. 5: From left to right: exact image, where the cropped portion is highlighted; blow-up (600%) of the anisotropic Gaussian PSF; blurred and noisy available image, with $\hat{\varepsilon} = 2 \cdot 10^{-2}$.

with $q \ll n$, typically. One can immediately see that, if $P_{i,j} \neq 0$, $i, j = 1, \dots, q$, the deblurring problem is underdetermined since, when convolving P with X^{ex} , additional (and unavailable) information about the exact image outside X^{ex} should be incorporated. A popular approach to overcome this phenomenon is to impose boundary conditions within the blurring process, i.e., to prescribe the behavior of the exact image outside X^{ex} (see [1] and the references therein). A 2D image restoration problem can be rewritten as a linear system (1.1), where the 1D array b is obtained by stacking the columns of the 2D blurred and noisy image (so that $N = n^2$), and the square matrix A incorporates the convolution process together with the boundary conditions. Although popular choices such as zero or periodic boundary conditions are particularly simple to implement, they often give rise to unwanted artifacts during the restoration process. The use of reflective or anti-reflective boundary conditions usually gives better results, as a sort of continuity of the image outside X^{ex} is imposed. Antireflective boundary conditions (ARBC) were originally introduced in [24], and further analyzed in several papers (see [5] and the references therein).

When dealing with a nonsymmetric PSF and ARBC, matrix-vector products with A can be implemented by fast algorithms (see [24]), but the same is not true for matrix-vector products with A^T (as, to the best of our knowledge, there is no known algorithm that can efficiently exploit the structure of A^T). Therefore, in practice, A^T is often approximated by a matrix A' defined by first rotating of 180° the PSF P used to build A (so to obtain the PSF P') and then modeling the 2D convolution process with P' and ARBC. In other words, image deblurring problems with a nonsymmetric PSF and ARBC can be only handled by transpose-free solvers. As addressed in Section 1, the authors of [5] propose to solve the equivalent, and somewhat symmetrized, linear system $AA'y = b$ (with $x = A'y$) by GMRES: in the following, this method is referred to as RP-GMRES. Our experiments are created by considering two different grayscale test images of size 256×256 pixels, together with two different PSFs, and antireflective or reflective boundary conditions; the sharp images are artificially blurred, and noise of variable levels is added. Matrix-vector products are computed efficiently by using the routines in `Restore Tools` [18]. An extension to handle ARBC within `Restore Tools` is available at: <http://scienze-como.uninsubria.it/mdonatelli/Software/software.html>.

In the first experiment the blurred image is cropped in order to reduce the effect of the chosen boundary conditions. The maximum number m_{max} of Arnoldi iterations for Algorithm 1 is set to 50, and only the stopping criterion (3.14) is considered with $\tau' = 10^{-15}$.

1. **Anisotropic Gaussian blur.** For this experiment, the elements $P_{i,j}$ of the PSF P are analytically given by the following expression

$$p_{i,j} = \exp\left(-\frac{1}{2(s_1^2 s_2^2 - \rho^4)} (s_2^2(i-k)^2 - 2\rho^2(i-k)(j-\ell) + s_1^2(j-\ell)^2)\right)$$

where $i, j = 1, \dots, d$, and $[k, \ell]$ is the center of the PSF. The values $s_1 = 4$, $s_2 = 1.3$, $\rho = 2$, and $d = 21$ are considered, and the noise level is $\hat{\varepsilon} = 2 \cdot 10^{-2}$. ARBC are imposed. The test data are displayed in Fig. 5. Figure 6 shows the best restorations achieved by each method; relative errors and the corresponding number of iterations are displayed in the caption. The GMRES restored image still appears pretty noisy




Fig. 6: The lower row displays blow-ups (200%) of the restored images in the upper row. From left to right: standard GMRES method (0.1483, $m = 4$); TF-CGLS method (0.1344, $m = 14$, $k = 12$); right-preconditioned GMRES (0.1354, $m = 19$).



Fig. 7: From left to right: exact image; blow-up (200%) of the anisotropic Gaussian PSF; blurred and noisy available image, with $\hat{\varepsilon} = 5 \cdot 10^{-2}$.

and blurred and, though the relative error for the RP-GMRES restoration is slightly larger than the TF-CGLS one, the two images are visually very similar and the latter greatly improves the standard GMRES one. It should also be emphasized that the cost of each RP-GMRES iteration is dominated by two matrix-vector products (one with A , and one with A'). Therefore, the cost of computing the GMRES, TF-CGLS, and RP-GMRES restorations is dominated by 4, 14, and 38 matrix-vector products, respectively. Therefore, for this experiment, TF-CGLS can deliver a solution whose quality is almost identical to the RP-GMRES one, with great computational savings.

2. **Atmospheric blur.** The test data for this experiment are displayed in Fig. 7. The PSF, of size 256×256 pixels and available within [18], models a realistic atmospheric blur. Reflective boundary conditions are imposed, so that multiplications with A^T can be easily computed. The noise level is $\hat{\varepsilon} = 5 \cdot 10^{-2}$. Figure 8 shows the best restorations achieved by the GMRES, the TF-CGLS, and the CGLS methods; relative errors and the corresponding number of iterations are displayed in the caption. Also for this test problem, the TF-CGLS method delivers much better solutions than the GMRES methods; also, TF-CGLS proves to be much more efficient than CGLS, as its computational cost is dominated by 18 matrix-vector products with A (versus the 80 matrix-vector products required by CGLS).




Fig. 8: The lower row displays blow-ups (200%) of the restored images in the upper row. From left to right: standard GMRES method ($4.0018 \cdot 10^{-1}$, $m = 4$); TF-CGLS method ($2.7855 \cdot 10^{-1}$, $m = 18$, $k = 5$); CGLS method ($2.7619 \cdot 10^{-1}$, $m = 40$).

5 Conclusions

This paper presented a new class of transpose-free CG-like methods, which appear to be competitive with their standard counterparts. These methods are particularly meaningful when the transpose of the coefficient matrix is not easily available, and they represent a very valid alternative to the standard GMRES method, as they can successfully handle situations where the latter performs badly (e.g., when the SVD components of the original matrix A are heavily mixed in the GMRES approximation subspace).

When compared with CGLS, the new transpose-free CG-like methods have similar performances, and they can be employed also to solve well-posed problems, provided that some insight into the SVD behavior of the projected problems is available.

Funding: This work was partially supported by GNCS-INdAM and FRA-University of Trieste. P. Novati is a member of the INdAM research group GNCS.

References

- [1] S. Berisha and J. G. Nagy, Iterative image restoration. In: *Academic Press Library in Signal Processing* (Eds. R. Chellappa and S. Theodoridis), Vol. 4, Elsevier, 2014, pp. 193–243.
- [2] D. Calvetti, B. Lewis, and L. Reichel, GMRES-type methods for inconsistent systems. *Linear Algebra Appl.* **316** (2000), 157–169.
- [3] D. Calvetti, B. Lewis, and L. Reichel, On the regularizing properties of the GMRES method. *Numer. Math.* **91** (2002), 605–625.
- [4] D. Calvetti, S. Morigi, L. Reichel, and F. Sgallari, Tikhonov regularization and the L-curve for large discrete ill-posed problems. *J. Comput. Appl. Math.* **123** (2000), 423–446.
- [5] M. Donatelli, D. Martin, and L. Reichel, Arnoldi methods for image deblurring with anti-reflective boundary conditions. *Appl. Math. Comput.* **253** (2015), 135–150.
- [6] Y. Dong, P. C. Hansen, M. E. Hochstenbach, and N. A. B. Riis, Fixing nonconvergence of algebraic iterative reconstruction with an unmatched backprojection. *SIAM J. Sci. Comput.* **41** (2019), No. 3, A1822–A1839.

- [7] S. Gazzola and P. Novati, Inheritance of the discrete Picard condition in Krylov subspace methods. *BIT* **56** (2016), No. 3, 893–918.
- [8] S. Gazzola, P. Novati, and M. R. Russo, On Krylov projection methods and Tikhonov regularization. *Electron. Trans. Numer. Anal.* **44** (2015), 83–123.
- [9] S. Gazzola, S. Noschese, P. Novati, and L. Reichel, Arnoldi decomposition, GMRES, and preconditioning for linear discrete ill-posed problems. *Appl. Numer. Math.* **142** (2019), 102–121.
- [10] G. H. Golub and C. F. Van Loan, *Matrix Computations*, 3rd ed., Johns Hopkins University Press, Baltimore, MD, 1996.
- [11] M. Hanke, *Conjugate Gradient Type Methods for Ill-Posed Problems*. Longman, Essex, UK, 1995.
- [12] M. Hanke, On Lanczos based methods for the regularization of discrete ill-posed problems. *BIT* **41** (2001), 1008–1018.
- [13] P. C. Hansen, Regularization tools: a Matlab package for analysis and solution of discrete ill-posed problems. *Numer. Algorithms* **6** (1994), 1–35.
- [14] P. C. Hansen, *Rank-deficient and discrete ill-posed problems*. SIAM, Philadelphia, PA, 1998.
- [15] P. C. Hansen and T. K. Jensen, Noise propagation in regularizing iterations for image deblurring. *Electron. Trans. Numer. Anal.* **31** (2008), 204–220.
- [16] T. K. Jensen and P. C. Hansen, Iterative regularization with minimum-residual methods. *BIT* **47** (2007), 103–120.
- [17] I. Moret, A note on the superlinear convergence of GMRES. *SIAM J. Numer. Anal.* **34** (1997), 513–516.
- [18] J. G. Nagy, K. M. Palmer, and L. Perrone, Iterative methods for image deblurring: a Matlab object oriented approach. *Numer. Algorithms* **36** (2004), 73–93.
- [19] P. Novati, Some properties of the Arnoldi based methods for linear ill-posed problems. *SIAM J. Numer. Anal.* **55** (2017), 1437–1455.
- [20] P. Novati and M. R. Russo, A GCV based Arnoldi–Tikhonov regularization method. *BIT* **54** (2014), No. 2, 501–521.
- [21] D. P. O’Leary and J. A. Simmons, A bidiagonalization–regularization procedure for large scale discretizations of ill-posed problems. *SIAM J. Sci. Stat. Comp.* **2** (1981), No. 4, 474–489.
- [22] J. R. Ringrose, *Compact Non-Self-Adjoint Operators*. Van Nostrand Reinhold Company, London, 1971.
- [23] Y. Saad, *Iterative Methods for Sparse Linear Systems*, 2nd ed., SIAM, Philadelphia, PA, 2003.
- [24] S. Serra-Capizzano, A note on anti-reflective boundary conditions and fast deblurring models. *SIAM J. Sci. Comput.* **25** (2003), 1307–1325.