

ORTHOGONALIZATION with a NON-STANDARD INNER PRODUCT and APPROXIMATE INVERSE PRECONDITIONING

Jiří Kopal, Miro Rozložník, Miroslav Tůma and Alicja Smoktunowicz

Abstract

One of the most important and frequently used preconditioning techniques for solving symmetric positive definite systems $Ax = b$ is based on computing the approximate inverse factorization in the form $A^{-1} = ZZ^T$, where Z is upper triangular [1]. It is also a well-known fact that the factor Z can be computed column by column by means of the of an A -orthogonalization process applied to the unit basis vectors e_1, \dots, e_n . As noted in [3] such A -orthogonalization also produces the Cholesky factor of the matrix $A = U^T U$ where $U^{-1} = Z$. This fact has been exploited to construct efficient sparse approximate approximate inverse preconditioners [1, 2, 3]. In a more general setting, given the symmetric positive definite matrix A and the nonsingular matrix $Z^{(0)}$, we look for the factors Z and U so that $Z^{(0)} = ZU$ with $Z^T A Z = I$ and the upper triangular U is a Cholesky factor of the matrix $(Z^{(0)})^T A Z^{(0)} = U^T U$.

There are several ways how to compute the matrices Z and U . If we have the spectral decomposition $A = V\Lambda V^T$, the factor U can be obtained from the standard QR decomposition $\Lambda^{1/2} V^T Z^{(0)} = QU$. The factor Z can be then recovered as $Z = V\Lambda^{-1/2} Q$. Probably the most straightforward and frequently used approach is the Gram-Schmidt orthogonalization, which consecutively A -orthogonalizes the columns of $Z^{(0)}$ against previously computed vectors using the orthogonalization coefficients that form the factor U . In the classical Gram-Schmidt algorithm (CGS), the A -orthogonal vectors are computed via matrix-vector updates which are relatively easy to parallelize. The rearrangement of this scheme has led to the modified Gram-Schmidt algorithm (MGS) with better numerical properties. Introducing sequential orthogonalization however destroys desirable parallel properties of the algorithm. We will discuss also yet another variant of sequential orthogonalization, which is motivated originally by the AINV preconditioner and which uses oblique projections [3]. We will refer to this scheme as the AINV orthogonalization. The main motivation for the development of approximate inverse techniques came from parallel processing and so the early papers on inverse factorization did not study numerical properties of algorithms. However, concerns on robustness and accuracy became very quickly an important aspect and resulted into a significant progress in recent preconditioning techniques. While the initial schemes like the basic AINV algorithm were based on oblique projections or on the CGS orthogonalization [3], the development lead to their stabilization both in terms of the orthogonalization scheme (MGS in the SAINV algorithm [2]) and in terms of appropriate computation of diagonal entries in U (one-sided versus stabilized versions of AINV [2, 3]).

From a numerical point of view, all these techniques may produce vectors which are far from orthogonal. The orthogonality between computed vectors is however crucial for the quality of the preconditioner constructed in the approximate inverse factorization. Given some approximation \bar{Z} to Z such that $A^{-1} \approx \bar{Z}\bar{Z}^T$, we are especially interested in the loss of orthogonality between the columns of \bar{Z} measured by the 2-norm of the matrix $\bar{Z}^T A \bar{Z} - I$. It is a well-known fact that eigenvalues of $\bar{Z}^T A \bar{Z}$ determine the convergence rate of the preconditioned conjugate gradient method applied to $\bar{Z}^T A \bar{Z} y = \bar{Z}^T b$ where $x = \bar{Z} y$. Therefore the orthogonal basis problem is of a primary interest for this application. While for the case of the standard inner product there exists a complete rounding error analysis for all main orthogonalization schemes [4, 7], the numerical properties of the schemes with a non-standard inner product are much less understood.

In this contribution we review the most important schemes used for orthogonalization with respect to the non-standard inner product and give the worst-case bounds for corresponding quantities computed in finite precision arithmetic. We formulate our results on the loss of orthogonality and on the factorization error (measured by $\|\bar{Z}^T A \bar{Z} - I\|$ and $\|Z^{(0)} - \bar{Z}\bar{U}\|$) in terms of quantities proportional to the roundoff unit u , in terms of the condition number $\kappa(A)$ which represents an upper bound for the relative error in computing the A -inner product as well as the condition number of the matrix $A^{1/2}Z^{(0)}$ which plays an important role in the factorization $(Z^{(0)})^T A Z^{(0)} \approx \bar{U}^T \bar{U}$.

Although all orthogonalization schemes are mathematically equivalent, their numerical behavior can be significantly different. It follows from our analysis that while the factorization error is quite comparable for all these schemes, the orthogonality between computed vectors can be significantly lost and it depends on the condition number $\kappa(A)$. This is the case also for the eigenvalue-based (EIG) implementation and Gram-Schmidt with reorthogonalization (CGS2). The classical Gram-Schmidt algorithm and AINV orthogonalization behave very similarly and generate vectors with the orthogonality that besides $\kappa(A)$ depends also on the factor $\kappa(A^{1/2}Z^{(0)})\kappa(Z^{(0)})$ (it essentially means the quadratic dependence on the condition number of the matrix $A^{1/2}Z^{(0)}$). Since the orthogonality in the modified Gram-Schmidt algorithm depends only linearly on $\kappa(A^{1/2}Z^{(0)})$, MGS appears to be a good compromise between expensive EIG or CGS2 and less accurate CGS or AINV. Indeed in the context of approximate inverse preconditioning the stabilization of AINV has led to the SAINV algorithm which uses exactly the MGS orthogonalization. We treat separately the particular case of a diagonal A which is extremely useful in the context of weighted least squares problems. One can show then that local errors arising in the computation of a non-standard inner product do not play an important role here and that the numerical behavior of these schemes is almost identical to the behavior of the orthogonalization schemes with the standard inner product. For all these results we refer to [6]. We believe that these results are an initial step towards understanding the behavior of practical strategies in approximate inverse preconditioning and will stimulate further research of schemes that use some inexact orthogonalization with appropriate dropping criterion and lead to some sparse approximation of Z and U . For a survey of such preconditioning techniques we refer to [1].

References

- [1] M. Benzi. Preconditioning techniques for large linear systems: a survey. *J. Comput. Phys.*, 182(2):418–477, 2002.
- [2] M. Benzi, J. K. Cullum, and M. Tůma, Robust approximate inverse preconditioning for the conjugate gradient method, *SIAM J. Sci. Comput.*, 22(4), 1318–1332, 2000.
- [3] M. Benzi, C. D. Meyer, and M. Tůma, A sparse approximate inverse preconditioner for the conjugate gradient method, *SIAM J. Sci. Comput.*, 17(5), 1135–1149, 1996.
- [4] L. Giraud, J. Langou, M. Rozložník, and J. van den Eshof. Rounding error analysis of the classical Gram-Schmidt orthogonalization process, *Num. Math.*, 101: 87-100, 2005.
- [5] L. Giraud, J. Langou, and M. Rozložník. On the loss of orthogonality in the Gram-Schmidt orthogonalization process. *Comput. Math. Appl.*, 50(7):1069–1075, 2005.
- [6] J. Kopal, M. Rozložník, A. Smoktunowicz and M. Tůma. Rounding error analysis of orthogonalization with a non-standard inner product. 2010, submitted to *Num. Math.*
- [7] A. Smoktunowicz, J.L. Barlow, and J. Langou. A note on the error analysis of classical Gram-Schmidt. *Numer. Math.*, 105(2):299–313, 2006.