Orthogonalization with a non-standard inner product with the application to approximate inverse preconditioning

Miro Rozložník *

Abstract

In this contribution we review the most important schemes used for orthogonalization of column vectors (stored in the matrix B) with respect to the non-standard inner product (induced by some symmetric positive definite matrix A) 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 for the classical Gram-Schmidt algorithm (CGS), modified Gram-Schmidt algorithm (MGS) algorithm and for yet another variant of sequential orthogonalization, which is motivated originally by the AINV preconditioner and which uses oblique projections. 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 expensive implementation based on eigenvalue decomposition (EIG) 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 $\kappa(A^{1/2}B)\kappa(B)$ (it essentially means the quadratic dependence on the condition number of the matrix $A^{1/2}B$). Since the orthogonality in the modified Gram-Schmidt algorithm depends only linearly on $\kappa(A^{1/2}B)$, 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 lead 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 [1].

References

*Joint work with Jiří Kopal, Miroslav Tůma and Alicja Smoktunowicz, Institute of Computer Science, Czech Academy of Sceinces, Prague, Czech Republic, email: miro@cs.cas.cz.

[1] 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. (see also http://www.cs.cas.cz/miro/krst10.pdf)