

A Grassmann-Rayleigh Quotient Iteration for Computing Invariant Subspaces

P.-A. Absil[‡], R. Mahony[§], R. Sepulchre[‡], P. Van Dooren[¶]

[‡] Institut Montefiore, Université de Liège,
Bât. B28 Systèmes, Grande Traverse 10, 4000 Liège, Belgium.
absil@montefiore.ulg.ac.be, R.Sepulchre@ulg.ac.be

[§] Monash University, Clayton, Victoria, Australia.

[¶] Université Catholique de Louvain, Belgium.

Abstract

The classical Rayleigh Quotient Iteration (RQI) computes a 1-dimensional invariant subspace of a symmetric matrix A with cubic convergence. We propose a generalization of the RQI which computes a p -dimensional invariant subspace of A . The geometry of the algorithm on the Grassmann manifold $\text{Gr}(p, n)$ is developed to show cubic convergence and to draw connections with recently proposed Newton algorithms on Riemannian manifolds.

Keywords: Rayleigh quotient iteration, invariant subspace, Grassmann manifold.

AMS Mathematics Subject Classification 2000: 65F15.

1 Introduction

There exist many classical methods for computing a single eigenpair (eigenvector and eigenvalue) of a symmetric matrix A . Among them are the power method and the (shifted) inverse iteration, with the Rayleigh quotient iteration (RQI) as a particular case [Par74, Par98]. The RQI is of particular interest because of its cubic convergence and its potential use in the shifted QR algorithm [Wat82, Par98].

In some cases, especially for multiple or clustered eigenvalues, it is advisable to compute the whole invariant subspace spanned by the corresponding eigenvectors. Methods have been proposed to achieve this task ([Ste73], [DMW83], [Cha84], unified in [Dem87]) but they only display linear convergence (as it is the case of the classical subspace iteration of Rutishauser [Rut69, Rut70]), or quadratic convergence but at a high numerical cost of $O(n^3)$ per iteration.

In the present paper, we propose a generalization of the RQI dealing with p -dimensional subspaces of $\mathbb{R}^n$. The property of cubic convergence of the classical RQI

extends to the generalized algorithm. Moreover, the numerical cost of each iteration is shown to be as low as $O(np^2)$ after a preliminary reduction to condensed form.

The generalized RQI we propose for invariant subspace computation is defined as follows.

Algorithm 1.1 *Pick an orthonormal basis $X_{(0)}$ in $\mathbb{R}^{n \times p}$, i.e. $X_{(0)}^T X_{(0)} = I_p$. Then, for $k = 0, 1, 2, \dots$, repeat the following:*

1. *Compute the solution $Z \in \mathbb{R}^{n \times p}$ of the Sylvester equation*

$$T_{X_{(k)}} Z \equiv AZ - ZX_{(k)}^T AX_{(k)} = X_{(k)}. \quad (1)$$

2. *Compute $X_{(k+1)} := \text{qf}(Z)$, where $\text{qf}(Z)$ denotes the Q -factor of the QR decomposition of Z .*

The iterates in Algorithm 1.1 are orthogonal matrices $X_{(k)}$ but our interest is in fact directed towards a p -dimensional invariant *subspace* of A . We show in Section 3 that Algorithm 1.1 indeed defines an algorithm on the set of p -dimensional linear subspaces of $\mathbb{R}^n$. This set can be endowed with a manifold structure, called the *Grassmann manifold* $\text{Gr}(p, n)$. The geometrical structure of $\text{Gr}(p, n)$ is used to define a *distance* between subspaces (Section 2) which is instrumental in the proof of cubic local convergence in the sense of subspaces (Section 4).

This structure is also helpful in establishing connections between Algorithm 1.1 and other algorithms for refining invariant subspaces estimations (Section 6). These other algorithms can be classified into two categories: methods based on perturbation theory [Ste73], all leading up to solving a Riccati equation by iteration [Dem87]; and optimization-oriented methods taking benefit of the fact that the invariant subspaces of A are the stationary points of a well chosen generalization of the Rayleigh quotient on the Grassmann manifold [EAS98].

2 Grassmann manifold

We denote by $\text{Gr}(p, n)$ the set of all p -dimensional linear subspaces of $\mathbb{R}^n$ (“linear” will be omitted in the sequel). A p -dimensional subspace $\mathcal{Y}$ of $\mathbb{R}^n$ shall be represented by an $n \times p$ full column rank matrix Y whose columns span this space $\mathcal{Y}$. For ease of reference, we denote by $\text{ST}(p, n)$, termed the *noncompact Stiefel manifold*, the set of the $n \times p$ real matrices with full column rank. Throughout this text, we use the notation $\lfloor Y \rfloor$ for the column space of Y

$$\lfloor Y \rfloor := \text{colspace}(Y) = \{Y\alpha : \alpha \in \mathbb{R}^p\}. \quad (2)$$

The representation of $\mathcal{Y}$ by Y is not unique. Indeed, $\lfloor Y \rfloor = \lfloor Z \rfloor$ if and only if there exists a nonsingular p -by- p matrix M such that $Z = YM$.

We are about to derive a way of uniquely representing the elements of certain subsets of $\text{Gr}(p, n)$. Consider the *Stiefel manifold* $\text{St}(p, n)$ (e.g. [EAS98]), defined as the set of all n -by- p orthogonal real matrices,

$$\text{St}(p, n) = \{X \in \mathbb{R}^{n \times p} : X^T X = I_p\}, \quad (3)$$

not to be confused with the noncompact $\text{ST}(p, n)$ defined above. Consider a point W in $\text{St}(p, n)$ and pick $W_\perp$ in $\text{St}(n-p, n)$ such that $\lfloor W \rfloor$ and $\lfloor W_\perp \rfloor$ are orthogonal, that is $W_\perp^T W = 0$. Then $(W|W_\perp) \in O_n$ and its columns form an orthonormal basis of $\mathbb{R}^n$. Any Y in $\text{ST}(p, n)$ admits the decomposition

$$Y = WY_1 + W_\perp Y_2, \quad (4)$$

where $Y_1 = W^T Y$ and $Y_2 = W_\perp^T Y$. Note that Y_1 and Y_2 depend on the basis $(W|W_\perp)$. If Y_1 is invertible, then there is a unique $K \in \mathbb{R}^{(n-p) \times p}$ such that $\lfloor Y \rfloor = \lfloor W + W_\perp K \rfloor$. From (4), one deduces that $K = Y_2 Y_1^{-1} = (W_\perp^T Y)(W^T Y)^{-1}$. This K is called the *local expression* or *local coordinates* of $\lfloor Y \rfloor$ relative to the reference $(W|W_\perp)$. The mapping $\lfloor Y \rfloor \mapsto K$ defines a *coordinate chart*. Any other reference point $(W|W_\perp)$ in O_n defines another coordinate chart. These charts make $\text{Gr}(p, n)$ a *manifold* in the sense of the differential geometry (see e.g. [HM94]). This manifold is the so called *Grassmann manifold* of p -dimensional linear subspaces of $\mathbb{R}^n$.

The *distance* between two subspaces $\lfloor Y \rfloor$ and $\lfloor Z \rfloor$ is defined as

$$\text{dist}(\lfloor Y \rfloor, \lfloor Z \rfloor) = \|P_Y - P_Z\|_2, \quad (5)$$

where P_Y and P_Z are the orthogonal projectors on $\lfloor Y \rfloor$ and $\lfloor Z \rfloor$ respectively [Ste73]. The following theorem (proved in [AMSD00]) establishes that $\|K\|_2$ is a valuable measure of the distance (5) in view of assessing the rate of convergence of $\lfloor X + X_\perp K \rfloor$ to $\lfloor X \rfloor$.

Theorem 2.1 *Let $X \in \text{St}(p, n)$ and $(X|X_\perp) \in O_n$. Then the tangents of the canonical angles between the subspaces $\lfloor X \rfloor$ and $\lfloor X + X_\perp K \rfloor$ are the singular values of K , and*

$$\begin{aligned} \text{dist}(\lfloor X \rfloor, \lfloor X + X_\perp K \rfloor) &= \sin \text{atan} \|K\|_2 \\ &= \frac{\|K\|_2}{\sqrt{1 + \|K\|_2^2}} = \|K\|_2(1 + O(\|K\|_2^2)). \end{aligned} \quad (6)$$

3 Grassmann-Rayleigh Quotient Iteration

By analogy with the classical Rayleigh quotient iteration (see e.g. [Par74, Par98]), we consider the following *Grassmann-Rayleigh Quotient Iteration* (GRQI):

Algorithm 3.1 (GRQI) *Pick $\mathcal{Y}_{(0)}$ in Grassmann manifold. Then, for $k = 0, 1, 2, \dots$, repeat the following:*

1. *Pick $Y \in \text{ST}(p, n)$ such that $\lfloor Y \rfloor = \mathcal{Y}_{(k)}$.*
2. *Solve*

$$T_Y Z \equiv AZ - Z(Y^T Y)^{-1} Y^T A Y = Y \quad (7)$$

for $Z \in \mathbb{R}^{n \times p}$.

3. *Define $\mathcal{Y}_{(k+1)} := \lfloor Z \rfloor$.*

The following proposition states that Algorithm 3.1 indeed defines an algorithm on $\text{Gr}(p, n)$.

Proposition 3.2 *The sequence $\{\mathcal{Y}_{(k)}\}_{k \geq 0}$ generated by Algorithm 3.1 is independent of the choice of Y used to represent $\mathcal{Y}$. That is, if Y_a and Z_a (Y_b , Z_b respect.) verify (7) and if $\lfloor Y_a \rfloor = \lfloor Y_b \rfloor$, then $\lfloor Z_a \rfloor = \lfloor Z_b \rfloor$.*

Proof: If $AZ - Z(Y^T Y)^{-1} Y^T A Y = Y$, then a right-multiplication by M leads to $AZM - ZM((YM)^T(YM))^{-1}(YM)^T A Y M = YM$, which means that YM and ZM also verify (7). ■

Moreover, Algorithm 3.1 and our original Algorithm 1.1 are equivalent in the sense of the p -subspaces:

Proposition 3.3 *Algorithms 3.1 and 1.1 define the same subspace iteration, that is, if $\lfloor X_{(k)} \rfloor = \mathcal{Y}_{(k)}$, then $\lfloor X_{(k+1)} \rfloor = \mathcal{Y}_{(k+1)}$.*

Proof: Algorithm 1.1 is a realization of Algorithm 3.1 with the particular choice $Y_+ = \text{qf}(Z)$. Note that the qf operation does not alter the column space. ■

Clearly, for computational matters, it is preferable to represent elements of $\text{Gr}(p, n)$ by orthogonal bases in order to prevent loss of numerical rank, which would deteriorate the quality of the subspace representation. Thus, Algorithm 1.1 must be interpreted as a numerically reliable realization of Algorithm 3.1.

4 Convergence of the GRQI

4.1 Local convergence

The following theorem generalizes the result of cubic local convergence of the classical RQI [Par74, Par98].

Theorem 4.1 (Cubic convergence of the GRQI)

Suppose A is symmetric real. Let $\mathcal{V}$ be an invariant subspace of A , and $\mathcal{V}_\perp$ be its orthogonal complement. If $A|_{\mathcal{V}}$ and $A|_{\mathcal{V}_\perp}$ have no eigenvalue in common, then the iterates $\mathcal{Y}_{(k)}$ of the GRQI (Algorithm 3.1) converge cubically to $\mathcal{V}$ for all $\mathcal{Y}_{(0)}$ in a neighbourhood of $\mathcal{V}$.

The proof is detailed in [AMSD00]. Here is the outline.

We shall need the two following lemmas (proved in [AMSD00]).

Lemma 4.2 Let A, B be $p \times p$ symmetric matrices. Let C be a $p \times p$ matrix. Define $\text{gap}[A, B] = \min |\lambda(A) - \lambda(B)|$. Suppose $\text{gap}[A, B] > 0$. Then, for all E, F in $\mathbb{R}^{p \times p}$ with $\|E\|_F, \|F\|_F < \text{gap}[A, B]/2$, the equation

$$(A + E)Z - Z(B + F) = C$$

admits one and only one solution $Z \in \mathbb{R}^{p \times p}$, and

$$\|Z\|_2 \leq \frac{\|C\|_2}{\text{gap}[A, B] - \|E\|_F - \|F\|_F}.$$

Lemma 4.3 Let D be diagonal and let $\|\cdot\|$ denote any consistent matrix norm. Then for every $\epsilon > 0$ there exists a $\delta > 0$ with the following property: for every E such that $\|E\| < \delta$ and such that D and $D - E$ have no eigenvalues in common, the unique solution Z to

$$DZ - Z(D - E) = I$$

is invertible and its inverse satisfies

$$\|Z^{-1}\| \leq (1 + \epsilon)\|E\|.$$

A step of the GRQI (Algorithm 3.1) is defined by

$$\mathcal{Y} = [Y] \mapsto \mathcal{Y}_+ = [Z] \quad (8)$$

where Y and Z solve (7) repeated below

$$\mathbf{T}_Y Z \equiv AZ - Z(Y^T Y)^{-1} Y^T A Y = Y.$$

Take $V \in \text{St}(p, n)$ such that $[V] = \mathcal{V}$ and that $\Lambda_1 := V^T A V$ is diagonal. Also take $V_\perp \in \text{St}(n - p, n)$ such that $[V_\perp] = \mathcal{V}_\perp$ and that $\Lambda_2 := V_\perp^T A V_\perp$ is diagonal. The matrices Λ_1 and Λ_2 have no element in common since their elements are the eigenvalues of $A|_{\mathcal{V}}$ and $A|_{\mathcal{V}_\perp}$ respectively.

Let K be the unique $(n - p)$ -by- p matrix such that $\mathcal{Y} = [V + V_\perp K]$. So, K is the local representation of $\mathcal{Y}$ relative to $(V|V_\perp)$. By Theorem 2.1,

$$\text{dist}(\mathcal{V}, \mathcal{Y}) = \|K\|_2(1 + O(\|K\|_2^2)).$$

Decompose Z according to $Z = V Z_1 + V_\perp Z_2$, so that $K_+ := Z_2 Z_1^{-1}$ is the local representation of $\mathcal{Y}_+$ relative to $(V|V_\perp)$. Replacing $Y = V + V_\perp K$ and $Z = V Z_1 + V_\perp Z_2$ in (7), and premultiplying by V^T and $V_\perp^T$ respectively, leads to

$$\Lambda_1 Z_1 - Z_1(I_p + K^T K)^{-1}(\Lambda_1 + K^T \Lambda_2 K) = I_p \quad (9)$$

$$\Lambda_2 Z_2 - Z_2(I_p + K^T K)^{-1}(\Lambda_1 + K^T \Lambda_2 K) = K \quad (10)$$

$$K_+ = Z_2 Z_1^{-1}. \quad (11)$$

First note that

$$\begin{aligned} & (I + K^T K)^{-1}(\Lambda_1 + K^T \Lambda_2 K) \\ &= \Lambda_1 - (K^T K \Lambda_1 - K^T \Lambda_2 K + O(K^4)). \end{aligned}$$

Lemma 4.2 applied to (10) gives the inequality

$$\|Z_2\|_2 \leq \frac{\sqrt{p}}{\text{gap}[\Lambda_1, \Lambda_2] + O(\|K\|^2)} \|K\|_2 \quad (12)$$

for all K sufficiently small, and Lemma 4.3 applied to (9) leads to

$$\|Z_1^{-1}\| \leq (\|\Lambda_1\| + \|\Lambda_2\|)\|K\|^2 + O(K^4) \quad (13)$$

for all K sufficiently small. From the two inequalities above, one deduces

$$\|K_+\|_2 \leq \sqrt{p} \frac{\|A|_{\mathcal{V}}\|_2 + \|A|_{\mathcal{V}_\perp}\|_2}{\text{gap}[A|_{\mathcal{V}}, A|_{\mathcal{V}_\perp}]} \|K\|_2^3 + O(K^5), \quad (14)$$

that is, cubic convergence.

Note that Lemma 4.3 cannot be used in the nongeneric case where Λ_1 and $(I + K^T K)^{-1}(\Lambda_1 + K^T \Lambda_2 K)$ have a common eigenvalue. This case corresponds to a singular $\mathbf{T}_Y$. An easy remedy when a singular $\mathbf{T}_Y$ shows up during the iteration is to slightly perturb the offending iterate so as to dismiss the singularity while staying in the neighbourhood where cubic convergence holds.

4.2 Global convergence

In [Par80, Par98], it is shown that the iterates $x_{(k)}$ of the classical RQI either converge to an eigenvector of A , or converge to the bisectors of a pair of eigenvectors of A , the latter situation being unstable under perturbations of $x_{(k)}$. The proof, due to Kahan, relies on the observation that Euclidean norm of the residual

$$r = (A - \rho_A(x)I)x, \quad (15)$$

where $\rho_A(x) = x^T A x$, is monotone decreasing under the (classical) RQI.

A straightforward generalization of (15) in the case where $X \in \text{St}(p, n)$ would be

$$\text{Res} = AX - XR_A(X) \quad (16)$$

where $R_A(X) = X^T AX$ is the *matrix Rayleigh quotient*.

Another way to measure the “eigenness” of a space $[X]$ is the principal angle between $[AX]$ and $[X]$. This measure is more fair because it is invariant under a multiplication of A by a constant.

Unfortunately, neither of the two distances proposed above shows a systematic decrease in numerical tests. This compromises the possibility of establishing a global convergence analysis along the lines of Kahan’s proof.

Nonetheless, in numerical simulations, the GRQI invariably converges to an invariant subspace.

5 Practical implementation of the algorithm

The practical relevance of the new method depends on two important issues: (i) keeping the computational cost of each iteration step as low as possible and (ii) detecting stagnation or convergence of the method.

5.1 Computational efficiency

The key equation of Algorithm 1.1 is the Sylvester equation

$$AZ - Z\hat{A} = X, \quad X^T X = I_p \quad (17)$$

where A and $\hat{A} = R_A(X) = X^T AX$ are both symmetric. It is easy to see that orthogonal similarity transformations on A and $\hat{A}$ do not alter the nature or solvability of the equation. Define indeed

$$A_u := U^T A U, \quad \hat{A}_v := V^T \hat{A} V, \quad X_{uv} := U^T X V \quad (18)$$

then $Z_{uv} := U^T Z V$ solves the transformed equation

$$A_u Z_{uv} - Z_{uv} \hat{A}_v = X_{uv} \quad (19)$$

if and only if Z solves (17). One shows moreover that both systems of equations have the same conditioning [DV85]. The matrix $\hat{A} = X^T AX$ changes at each step of the iteration but is small compared to A . Transforming the n -by- n matrix A to a more convenient form A_u thus needs to be done **only once**, whereas computing a special form of the p -by- p matrix $\hat{A}_v$ can be done at each step of the algorithm. For this reason, we assume A to be already in an “appropriate” form $A_u = A$ and dismiss the index u in the sequel. Efficient methods to solve (17) are typically based on an eigenvalue

decomposition of the smallest matrix $\hat{A}$ [BS72, GNL79]. Then $\hat{A}_v$ is diagonal:

$$\hat{A}_v = \text{diag}\{\rho_1, \dots, \rho_p\} = V^T \hat{A} V. \quad (20)$$

An implementation of Algorithm 1.1 based on such reductions then requires the following basic steps at each iteration:

$$\hat{A} := X^T A X \quad (21)$$

$$\text{find } V \text{ such that } V^T \hat{A} V = \text{diag}\{\rho_1, \dots, \rho_p\} \quad (22)$$

$$X_v := X V \quad (23)$$

$$\text{find } Z_v \text{ such that } A Z_v - Z_v \hat{A}_v = X_v \quad (24)$$

$$X_+ := \text{qf}(Z_v V^T). \quad (25)$$

In the literature, the eigenvalues ρ_i of $\hat{A} = R_A(X)$ are called the Ritz values of A with respect to $[X]$ and the corresponding columns of X_v are called the Ritz vectors.

Steps (22), (23) and (25) involve dense matrices and require $O(p^3)$, $O(np^2)$ and $O(np^2)$ floating point operations (flops), respectively. One would like the other two steps (21) and (24) to require at most $O(np^2)$ flops as well. This is possible by putting A in a special “condensed” form. An obvious one is to choose A tridiagonal. Step (21) then requires $O(np^2)$ flops and step (24) can be decoupled in the individual systems

$$(A - \rho_i I) z_i = x_i \quad (26)$$

for each column x_i of X_v and z_i of Z_v , respectively. Since these systems are tridiagonal, they each require $O(n)$ flops.

We should point out that the preliminary reduction of A to tridiagonal form requires $O(n^3)$ flops and becomes therefore the most time consuming step of the procedure. A more economical idea would be to reduce A to a banded matrix with bandwidth $2k+1$. Such a matrix can also be viewed as a block tridiagonal matrix with $k \times k$ blocks. The complexity of obtaining such a form is comparable to that of obtaining a “scalar” tridiagonal form but it can be computed more efficiently on parallel architectures [DDSvdV91]. Moreover, if $k^2 \simeq p$ then both steps (21) and (24) still require $O(np^2)$ flops.

In many applications (e.g. PDE’s) the matrix A has a special sparsity pattern that can also be exploited. One e.g. often encounters matrices A that have already a banded form and therefore do not need a preliminary reduction. Other forms of A allow cheap solutions for the systems (26) although A does not have a banded form. These include e.g. low rank modifications of banded matrices or matrices with special structure such as Toeplitz or Hankel matrices. Finally, we point out that in many applications A is sparse and should therefore not be transformed anymore to a condensed form. Instead, one can use iterative solvers to find a

solution z_i to (26) for each right hand side x_i . It turns out [EW96] that a high relative accuracy of (26) is only needed in the last few steps of the iteration and hence that a lot of flexibility can be built into the iterative procedure.

Our recommendation is thus to use this algorithm on a block tridiagonal form of A when A is dense and not to reduce A at all when A is sparse, but rather to use flexible iterative solvers for the equations (26).

5.2 Stopping criterion

For the stopping criterion of the algorithm, one has to detect whether the solution Z of (17) is such that $\lfloor Z \rfloor$ is nearly equal to $\lfloor X \rfloor$. A possibility is to rely on $\text{dist}(\lfloor X \rfloor, \lfloor Z \rfloor)$ defined in Section 2. From Theorem 2.1,

$$\text{dist}(\lfloor X \rfloor, \lfloor Z \rfloor) = \sin \arctan \|K_Z\|_2,$$

and

$$\begin{aligned} \|K\|_2 &= \|X_\perp K\|_2 = \|X_\perp X_\perp^T Z (X^T Z)^{-1}\|_2 \\ &= \|(I - XX^T)Z(X^T Z)^{-1}\|_2 \end{aligned}$$

the computation of which requires $O(np^2)$ flops. One can also use a criterion based on the “block-diagonality” of $(X|X_\perp)^T A (X|X_\perp)$ by checking the value of $\|X_\perp^T A X\|$, i.e. the norm of the residual Res (see Section 4.2). When A is suitably condensed, the computation requires $O(np^2)$ flops.

5.3 Comparison with independent RQI’s

The practical implementation (21)-(25) remains a realization of Algorithm 3.1 if we redefine (25) as

$$X_+ = \text{qf}(Z_v). \quad (27)$$

In fact, if we are not interested in the X ’s generated by Algorithm 1.1 but only by the $\lfloor X \rfloor$ ’s, we can content ourselves with an implementation of Algorithm 3.1, which can be as follows:

Algorithm 5.1 (Implementation of GRQI)

- Pick $Y_{(0)} \in \text{ST}(p, n)$. Then, for $k = 0, 1, 2, \dots$:*
1. *Orthonormalize the columns of $Y_{(k)}$ to get $X \in \text{St}(p, n)$.*
 2. *Compute $\hat{A} := X^T A X$.*
 3. *Find V such that $V^T \hat{A} V = \text{diag}\{\rho_1, \dots, \rho_p\}$.*
 4. *$X_v := XV$.*
 5. *For $i = 1, \dots, p$, find Z such that $AZ - Z \text{diag}\{\rho_1, \dots, \rho_p\} = X_v$. If this system is singular, dismiss the singularity by slightly perturbing the ρ_i ’s.*
 6. *$Y_{(k+1)} := Z$.*

Algorithm 5.1 shows that the GRQI step can be interpreted as p classical RQI steps in parallel preceded by an orthonormalization step. Because this orthonormalization is expensive ($O(np^2)$), one could prefer to

skip this step and to define $x_{i+} = z_i$ after (26), which would result in p independent Rayleigh quotient iterations. This simplification of the GRQI has two disadvantages. First, the local convergence property is lost. Indeed, there exist starting points arbitrarily close to certain invariant subspaces for which parallel evolution of p RQIs results in convergence to a different invariant subspace [BS89]. Second, parallel evolution of p RQIs without orthogonalization is very prone to column merging in the course of the iterations.

One could imagine to implement only a few steps of Algorithm 3.1 and then switch to p independent RQIs. It is difficult however to evaluate the right time to switch from one strategy to the other, because there is no bound on Ritz vectors [Par98]. If Algorithm 3.1 is abandoned too early, then the independent RQIs are likely to merge or to behave oddly as explained in the last paragraph. Moreover, the convergence of GRQI is so fast in the neighbourhood of an invariant subspace that one hardly has the time to make a decision.

6 Comparison with other methods

In comparison with other methods for computing invariant subspaces [Ste73, DMW83, Cha84, EAS98], the GRQI achieves a higher rate of convergence at a lower numerical cost after reduction to tridiagonal form (Table 1).

Method	Order of conv.	Num. cost
GRQI	cubic	$O(np^2)$
Newton on $\text{Gr}(p, n)$	cubic	$O(n^3)$
“Demmell iteration”	quadratic	$O(n^3)$
“Demmell Newton”	linear	$O(np^2)$

Table 1: Comparison of methods for invariant subspace computation. The last column gives the numerical cost by iteration for A block tridiagonal (see Section 5).

“Newton on $\text{Gr}(p, n)$ ” refers to the Newton algorithm [Smi94] on the manifold $\text{Gr}(p, n)$ for finding a stationary point of a generalized version of the classical Rayleigh quotient $\rho_A(\lfloor Y \rfloor) = \text{tr}(Y^T Y)^{-1} Y^T A Y$.

In [Dem87], Demmell unifies algorithms from [Cha84, DMW83, Ste73] by showing that they all attempt to solve the same Riccati equation

$$A_{22}K - KA_{11} = -A_{21} + KA_{12}K, \quad (28)$$

which solves the invariant subspace problem

$$AZ - ZB = 0$$

subject to the constraint

$$Z = W + W_\perp K, \quad K \in \mathbb{R}^{(n-p) \times p}.$$

There are two basic approaches to solving the Riccati equation (28): the iteration

$$A_{22}K_{i+1} - K_{i+1}A_{11} = -A_{21} + K_iA_{12}K_i, \quad (29)$$

which we refer to as “Demmel iteration” in the table above, and Newton’s method which reads

$$\begin{aligned} (A_{22} - K_iA_{12})K_{i+1} - K_{i+1}(A_{11} + A_{12}K_i) \\ = -A_{21} - K_iA_{12}K_i, \end{aligned} \quad (30)$$

referred to as “Demmel Newton” in our table.

7 Conclusion

We have shown how the Rayleigh quotient iteration can be extended to the Grassmann manifold $\text{Gr}(p, n)$ in view of computing a p -dimensional invariant subspace of a symmetric matrix A . Taking advantage of the geometry of $\text{Gr}(p, n)$ and of some properties of Sylvester equations, we showed cubic convergence of the algorithm. Comparison with other methods for computing invariant subspaces [Ste73, DMW83, Cha84, EAS98] leads to the conclusion that the Grassmann-Rayleigh Quotient Iteration achieves a higher rate of convergence at a lower numerical cost.

References

- [AMSD00] P.-A. Absil, R. E. Mahony, R. Sepulchre, and P. Van Dooren, *A Grassmann-Rayleigh Quotient Iteration for computing invariant subspaces*, Technical report, Montefiore Institute, University of Liège, Bât. B28 Systèmes, Grande Traverse 10, 4000 Liège, Belgium, 2000.
- [BS72] R. H. Bartels and G. W. Stewart, *Algorithm 432: Solution of the matrix equation $AX + XB = C$* , Commun. ACM **15** (1972), 820–826.
- [BS89] S. Batterson and J. Smillie, *The dynamics of Rayleigh Quotient Iteration*, SIAM J. Numer. Anal. **26** (1989), no. 3, 624–636.
- [Cha84] F. Chatelin, *Simultaneous Newton’s iteration for the eigenproblem*, Computing, Suppl. **5** (1984), 67–74.
- [DDSvdV91] J. Dongarra, I. Duff, D. Sorensen, and H. van der Vorst, *Solving linear systems on vector and shared memory computers*, SIAM, 1991.
- [Dem87] J. W. Demmel, *Three methods for refining estimates of invariant subspaces*, Computing **38** (1987), 43–57.
- [DMW83] J. J. Dongarra, C. B. Moler, and J. H. Wilkinson, *Improving the accuracy of computed eigenvalues and eigenvectors*, SIAM J. Numer. Anal. **20** (1983), no. 1, 23–45.
- [DV85] P. Van Dooren and M. Verhaegen, *On the use of unitary state-space transformations*, Linear algebra and its Role in Linear Systems Theory (B. N. Datta, ed.), AMS Contemporary Mathematics Series, vol. 47, American Mathematical Society, 1985, pp. 447–463.
- [EAS98] A. Edelman, T. A. Arias, and S. T. Smith, *The geometry of algorithms with orthogonality constraints*, SIAM J. Matrix Anal. Appl. **20** (1998), no. 2, 303–353.
- [EW96] S. Eisenstat and H. Walker, *Choosing the forcing terms in an inexact newton method*, SIAM J. Scientific Comp. **17** (1996), 16–32.
- [GNL79] G. H. Golub, S. Nash, and C. F. Van Loan, *A Hessenberg-Schur method for the problem $AX + XB = C$* , IEEE Trans. Automatic Control **24** (1979), 909–913.
- [HM94] U. Helmke and J. B. Moore, *Optimization and dynamical systems*, Springer, 1994.
- [Par74] B. N. Parlett, *The Rayleigh Quotient Iteration and some generalizations for nonnormal matrices*, Mathematics of Computation **28** (1974), no. 127, 679–693.
- [Par80] B. N. Parlett, *The symmetric eigenvalue problem*, Prentice-Hall, Inc., Englewood Cliffs, N.J. 07632, 1980.
- [Par98] B. N. Parlett, *The symmetric eigenvalue problem*, Society for Industrial and Applied Mathematics, Philadelphia, 1998, slightly revised edition of Par80.
- [Rut69] H. R. Rutishauser, *Computational aspects of F. L. Bauer’s simultaneous iteration method*, Numerische Mathematik **13** (1969), 4–13.
- [Rut70] H. R. Rutishauser, *Simultaneous iteration method for symmetric matrices*, Numerische Mathematik **16** (1970), 205–223.
- [Smi94] S. T. Smith, *Optimization techniques on Riemannian manifolds*, Hamiltonian and gradient flows, algorithms and control (Anthony Bloch, ed.), Fields Institute Communications, vol. 3, American Mathematical Society, 1994, pp. 113–136.
- [Ste73] G. W. Stewart, *Error and perturbation bounds for subspaces associated with certain eigenvalue problems*, SIAM Review **25** (1973), no. 4, 727–764.
- [Wat82] D. S. Watkins, *Understanding the QR algorithm*, SIAM Review **24** (1982), no. 4, 427–440.