Numerical Linear Algebra for Computational Science and Information Engineering

Arnoldi and Lanczos Procedures

Nicolas Venkovic
nicolas.venkovic@tum.de

Chair of Computational Mathematics
School of Computation, Information and Technology
Technical University of Munich



Outline I

1. Arnoldi process
   Section 6.1 in Darve & Wootters (2021) 2
2. Arnoldi Rayleigh-Ritz for dominant eigenpairs 10
3. Shift-and-invert Arnoldi Rayleigh-Ritz for interior eigenpairs 14
4. Arnoldi harmonic Ritz for interior eigenpairs 16
5. Lanczos process
   Section 6.2 in Darve & Wootters (2021) 30

Krylov subspace methods for few eigenpairs

- So far we've seen:
  - **Power iterations**, **inverse iterations** and **Rayleigh quotient iterations** to compute a single eigenpair.
  - **QR iterations**, the **divide-and-conquer method** and the **method of bissection** to compute all the eigenpairs of a small-to-medium size and dense matrix.
  - **LOBPCG** to compute a few extremal generalized eigenpairs of a large, possibly sparse matrix pencil (A, B).
- **Krylov subspace methods** are another set of iterative methods to compute a few eigenpairs of a large matrix A.
  - We assume that the mapping $x \mapsto Ax$ can be operated efficiently, possibly because A is sparse or as a matrix-free subroutines..
  - We denote two methods in particular:
    - The **Arnoldi** process is meant for non-symmetric matrices, and
    - The **Lanczos** process, which was introduced later for symmetric matrices.

Arnoldi process

Section 6.1 in Darve & Wootters (2021)

Arnoldi process for Krylov subspace basis construction

- Given a vector $v \in \mathbb{R}^n$, the k-th Krylov subspace of $A \in \mathbb{R}^{n \times n}$ is

$$\mathcal{K}_k(A, v) := \operatorname{span}\{v, Av, \ldots, A^{k-1}v\}.$$

- Krylov subspaces are the cornerstone of the approximation to solutions of linear algebraic problems in relation to the matrix A.
- As we saw with the power method, the vectors $v, Av, A^2v, \ldots$ quickly converge to the leading eigenvector of A.
  As a result, the vectors $v, Av, \ldots, A^{k-1}v$ constitute an ill-conditioned basis of the Krylov subspace.
  This ill conditioning is undesirable for approximation purposes in finite arithmetic.
- Instead of the naive choice of monomial basis vectors, the Arnoldi process seeks to generate an orthonormal basis $q_1, \ldots, q_k$ for the Krylov subspace:

$$\operatorname{span}\{q_1, \ldots, q_k\} = \mathcal{K}_k(A, v).$$

  We use the notation $Q_k := [q_1, \ldots, q_k]$ so that, in exact arithmetic, we have

$$Q_k^T Q_k = I_k \quad \text{and} \quad \operatorname{range}(Q_k) = \mathcal{K}_k(A, v).$$

Arnoldi process for Krylov subspace basis construction, cont'd$_1$

- In its general form, the Arnoldi procedure is as follows:

Algorithm 1 Arnoldi : $(v, k) \mapsto Q_k := [q_1, \ldots, q_k]$

1: $q_1 := v/\|v\|_2$
2: **for** $j = 1, \ldots, k-1$ **do**
3: $\quad q_{j+1} := \Pi^{(j)} A q_j / \|\Pi^{(j)} A q_j\|_2$

where $\Pi^{(j)}$ is a projector onto $\mathrm{range}(Q_j)^\perp$.

- The basis Q_k is orthonormal by construction, i.e., $Q_k^T Q_k = I_k$.
- For $j = 1, \ldots, k$, the matrix $P_j := Q_j Q_j^H$ is an orthogonal projector onto

$$\mathrm{range}(Q_j) = \mathrm{span}\{q_1, \ldots, q_j\}$$

- Since P_j is an orthogonal projector, the matrix

$$\Pi^{(j)} := I_n - P_j = I_n - Q_j Q_j^T$$

is an orthogonal projector onto

$$\mathrm{null}(P_j) = \mathrm{null}(Q_j) = \mathrm{range}(Q_j)^\perp.$$

Arnoldi process for Krylov subspace basis construction, cont'd$_2$

- In its general form, the Arnoldi procedure is as follows:

Algorithm 1 Arnoldi : $(v, k) \mapsto Q_k := [q_1, \ldots, q_k]$

1: $q_1 := v/\|v\|_2$
2: **for** $j = 1, \ldots, k-1$ **do**
3: $\quad q_{j+1} := \Pi^{(j)} A q_j / \|\Pi^{(j)} A q_j\|_2$

where $\Pi^{(j)}$ is a projector onto $\mathrm{range}(Q_j)^\perp$.

- We can show that $\mathrm{range}(Q_k) = \mathcal{K}_k(A, v)$ by induction:
  The base case is trivially satisfied as $\mathrm{span}\{q_1\} = \mathcal{K}_1(A, v)$.
  Using the hypothesis

$$\mathrm{range}(Q_j) = \mathrm{span}\{q_1, \ldots, q_j\} = \mathcal{K}_j(A, v) \quad \text{and} \quad q_j = p_{j-1}(A)v$$

where p_{j-1} is a polynomial of degree $j-1$, we get

$$q_{j+1} \propto \Pi^{(j)} A q_j = A q_j - Q_j (Q_j^T A q_j) = A p_{j-1}(A) v - \underbrace{Q_j (Q_j^T A q_j)}_{\in \mathcal{K}_j(A,v)}$$

so that $q_{j+1} = p_j(A)v$ for some polynomial p_j of degree j and $\mathrm{range}(Q_{j+1}) = \mathcal{K}_{j+1}(A, v)$. □

Reduction to Hessenberg form

- The Arnoldi procedure reduces A into Hessenberg. That is,

$$\underline{H}_k := Q_{k+1}^T A Q_k$$

is upper Hessenberg. This is proven as follows:

- First by construction, we have

$$q_{j+1} = \Pi^{(j)} A q_j / \|\Pi^{(j)} A q_j\|_2 \implies \Pi^{(j)} A q_j = \|\Pi^{(j)} A q_j\|_2 q_{j+1}.$$

- Then, we have

$$\begin{aligned} A q_j &= \Pi^{(j)} A q_j + Q_j Q_j^T A q_j \\ &= \|\Pi^{(j)} A q_j\|_2 q_{j+1} + \sum_{\ell=1}^{j} q_\ell q_\ell^T A q_j \end{aligned}$$

which we left multiply by q_i^T to obtain

$$h_{ij} = q_i^T A q_j = \|\Pi^{(j)} A q_j\|_2 q_i^T q_{j+1} + \sum_{\ell=1}^{j} (q_\ell^T A q_j) q_i^T q_\ell$$

so that $h_{ij} = 0$ for $i > j+1$. $\square$

The Arnoldi relation

- A key property of the Arnoldi process is the relation $AQ_k = Q_{k+1}\underline{H}_k$. It is proven as follows:

- First, we have:

$$h_{j+1,j} = q_{j+1}^T A q_j = \|\Pi^{(j)} A q_j\|_2 q_{j+1}^T q_{j+1} + \sum_{\ell=1}^{j} q_\ell q_\ell^T q_{j+1}^T A q_j = \|\Pi^{(j)} A q_j\|_2$$

so that $Aq_j = \sum_{\ell=1}^{j+1} h_{\ell j} q_\ell = \sum_{\ell=1}^{k+1} h_{\ell j} q_\ell$ for $j = 1, \ldots, k$.

- Then we have:

$$\begin{aligned} AQ_k &= \begin{bmatrix} Aq_1 & \ldots & Aq_k \end{bmatrix} \\ &= \begin{bmatrix} \sum_{\ell=1}^{k+1} h_{\ell 1} q_\ell & \ldots & \sum_{i=1}^{k+1} h_{\ell k} q_\ell \end{bmatrix} \\ &= \begin{bmatrix} Q_{k+1} h_{1:k+1,1} & \ldots & Q_{k+1} h_{1:k+1,k} \end{bmatrix} = Q_{k+1}\underline{H}_k \\ &= \begin{bmatrix} Q_k h_{1:k,1} & \ldots & Q_k h_{1:k,k} + h_{k+1,k} q_k \end{bmatrix} = Q_k H_k + h_{k+1,k} q_k e_k^T \end{aligned}$$

where $H_k := Q_k^T A Q_k$.

CGS-based Arnoldi procedure

- Deploying the procedure with the orthogonal projector $\Pi^{(j)} := I_n - Q_j Q_j^T$ leads to CGS-based Arnoldi:

Algorithm 2 CGS-based Arnoldi: $(v, k) \mapsto Q_k$

1: $q_1 := v/\|v\|_2$
2: **for** $j = 1, \ldots, k-1$ **do**
3: $\quad w := Aq_j$
4: $\quad h_{1:j,j} := Q_j^T w$ ▷ $h_{1:j,j} = Q_j^T A q_j$
5: $\quad w := w - Q_j h_{1:j,j}$ ▷ $w = \Pi^{(j)} A q_j$
6: $\quad h_{j+1,j} := \|w\|_2$ ▷ $h_{j+1,j} = \|\Pi^{(j)} A q_j\|_2$
7: $\quad q_{j+1} := w/h_{j+1,j}$

- The main computation effort of CGS-based Arnoldi lies in two BLAS-2 operations per iteration in addition to one matrix-vector product with A.
- We see that the Hessenberg matrix $\underline{H}_k$ is recovered as a by-product of the application of the projectors $\Pi^{(j)}$ for $j = 1, \ldots, k$.
- We saw previously that CGS leads to poorly conditioned bases.
  In practice, we seldom use the CGS-based Arnoldi procedure and rather resort to MGS- or CGS2-based implementations.

MGS-based Arnoldi procedure

- In the MGS-based implementation of the Arnoldi procedure, we resort to the mathematical equivalent (in exact arithmetic) orthogonal projector

$$\Pi^{(j)} := (I_n - q_j q_j^T) \dots (I_n - q_1 q_1^T).$$

Algorithm 3 MGS-based Arnoldi

1: $q_1 := v/\|v\|_2$
2: **for** $j = 1, \dots, k-1$ **do**
3: $\quad w := Aq_j$
4: $\quad$ **for** $i = 1, \dots, j$ **do**
5: $\quad\quad h_{ij} := q_i^T w$ $\quad \triangleright\ h_{ij} = q_i^T A q_j$
6: $\quad\quad w := w - h_{ij} q_i$ $\quad \triangleright\ w = \Pi^{(j)} A q_j$
7: $\quad h_{j+1,j} := \|w\|_2$ $\quad \triangleright\ h_{j+1} = \|\Pi^{(j)} A q_j\|_2$
8: $\quad q_{j+1} := w/h_{j+1,j}$

- The main computational effort is MGS-based Arnoldi is a loop of BLAS-1 operations in addition to one matrix-vector product with A per iteration. MGS produces better conditioned bases at smaller arithmetic intensity.
- The Hessenberg matrix $\underline{H}_k$ is still a by-product of the algorithm.

CGS2-based Arnoldi procedure

- Alternatively, a CGS2-based implementation of the Arnoldi procedure relies on the orthogonal projector

$$\Pi^{(j)} := (I_n - Q_j Q_j^T)(I_n - Q_j Q_j^T).$$

Algorithm 4 CGS2-based Arnoldi

$q_1 := v/\|v\|_2$
for $j = 1, \ldots, k-1$ **do**
  $w := Aq_j$
  $f := Q_j^T w$
  $w := w - Q_j f$
  $g := Q_j^T w$
  $w := w - Q_j g$ ▷ $w = \Pi^{(j)} Aq_j$
  $h_{1:j,j} := f + g$
  $h_{j+1,j} := \|w\|_2$ ▷ $h_{j+1} = \|\Pi^{(j)} Aq_j\|_2$
  $q_{j+1} := w/h_{j+1,j}$

- Although CGS2 requires twice the number of FLOPs as CGS and MGS, it often achieves better performance than MGS on distributed and hierarchical memories due to higher arithmetic intensity.

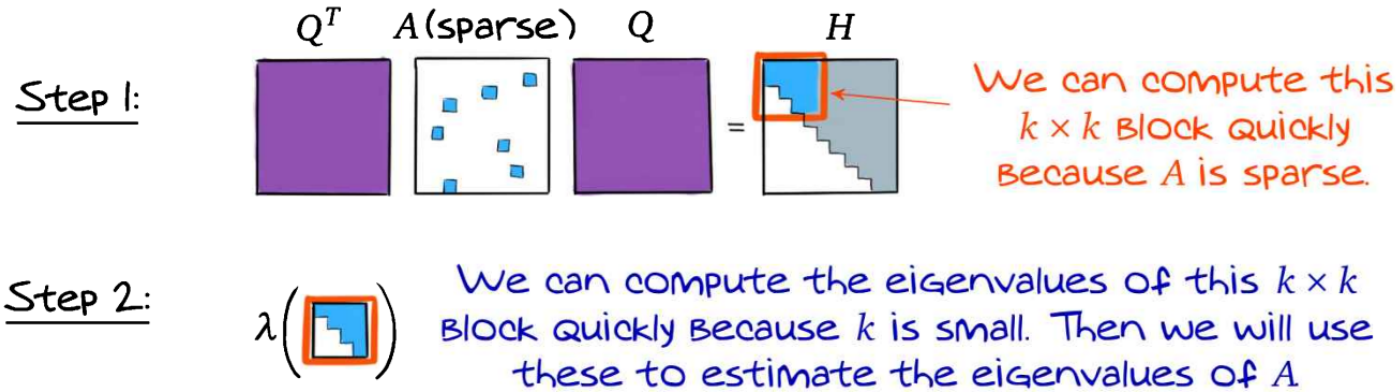
Arnoldi Rayleigh-Ritz for dominant eigenpairs

Reduction to Hessenberg form

- For $k = n$ with $Q := Q_n$, the basis Q forms a similarity transformation to Hessenberg form with

$$A = QHQ^T$$

  where $H = Q^T AQ$.
- By property of similarity transformations, the eigenvalues of H are the same as those of A. This can be exploited to find eigenvalues of A.
- Instead of considering the full Hessenberg matrix H, we approximate eigenpairs of A with eigenpairs of a leading block H_k of H with $k \ll n$.



Darve, E., & Wootters, M. (2021). Numerical linear algebra with Julia. Society for Industrial and Applied Mathematics.

- As it turns out, the eigenpairs of the leading block H_k of H are good approximations of some eigenpairs of A.

Arnoldi procedure with Rayleigh-Ritz vectors

- **Eigenvectors with eigenvalues whose norms are the largest** among the spectrum of A tend to be **well approximated by Rayleigh-Ritz** projections, as explained by Parlett (1998) and Saad (2011).
- Rayleigh-Ritz projections are **commonly defined with respect to Krylov subspaces** whose bases Q_k are obtained by a Arnoldi procedure:
  - Then, a **Rayleigh-Ritz vector** $y \in \text{range}(Q_k)$ approximates an eigenvector of A with the **Ritz value** λ such that $Ay - \lambda y \perp \text{range}(Q_k)$. That is, we **search for** $(\lambda, \hat{y}) \in \mathbb{C} \times \mathbb{C}^k \setminus \{0\}$ s.t. $z^H (Ay - \lambda y) = 0$ $\forall\, z \in \text{range}(Q_k)$ with $y = Q_k \hat{y}$. This **simplifies to**

$$Q_k^T (AQ_k\hat{y} - \lambda Q_k\hat{y}) = 0$$

$$H_k\hat{y} - \lambda\hat{y} = 0 \quad \Longrightarrow \quad \boxed{H_k\hat{y} = \lambda\hat{y}}$$

    where use is made of the Arnoldi relation and $Q_k^T Q_k = I_k$.
  - $nev < k$ **dominant eigenpairs** $\{(\lambda_\ell, \hat{y}_\ell)\}_{\ell=1}^{nev}$ **of** H_k are **used to approximate dominant eigenpairs of** A **with** $\{(\lambda_\ell, y_\ell)\}_{\ell=1}^{nev}$ where $y_\ell := Q_k \hat{y}_\ell$.

B. N. Parlett, The symmetric eigenvalue problem. Society for Industrial and Applied Mathematics (1998).

Y. Saad, Numerical methods for large eigenvalue problems: revised edition. Society for Industrial and Applied Mathematics (2011).

Arnoldi procedure with Rayleigh-Ritz vectors, cont'd$_1$

- A desirable property of the Rayleigh-Ritz approximation y_ℓ is that **the Ritz value λ_ℓ equates the corresponding Rayleigh quotient**:

$$y_\ell^H A y_\ell = (Q_k \hat{y}_\ell)^H A Q_k \hat{y}_\ell = \hat{y}_\ell^H Q_k^T A Q_k \hat{y}_\ell = \hat{y}_\ell^H H_k \hat{y}_\ell = \lambda_\ell \hat{y}_\ell^H \hat{y}_\ell = \lambda_\ell$$

  where $\hat{y}_\ell$ is assumed to have unit length.

- The **eigen-residual** $\tilde{r}_\ell := A y_\ell - \lambda_\ell y_\ell$ of the Rayleigh-Ritz vector y_ℓ is s.t.

$$\begin{aligned}
\tilde{r}_\ell &= A Q_k \hat{y}_\ell - \lambda_\ell Q_k \hat{y}_\ell \\
&= Q_k H_k \hat{y}_\ell + h_{k+1,k} q_{k+1} e_k^T \hat{y}_\ell - \lambda_\ell Q_k \hat{y}_\ell \\
&= \lambda_\ell Q_k \hat{y}_\ell + h_{k+1,k} q_{k+1} e_k^T \hat{y}_\ell - \lambda_\ell Q_k \hat{y}_\ell \\
&= h_{k+1,k} q_{k+1} e_k^T \hat{y}_\ell
\end{aligned}$$

$$\boxed{\tilde{r}_\ell = \beta_{k,\ell} q_{k+1}} \quad \text{where} \quad \boxed{\beta_{k,\ell} := h_{k+1,k} e_k^T \hat{y}_\ell}.$$

  - Essentially, **the eigen-residuals $\tilde{r}_1, \ldots, \tilde{r}_k$ of the Rayleigh-Ritz vectors $y_1, \ldots, y_k$** defined with respect to the **Krylov subspace $\mathcal{K}_k(A, q_1)$** are **all parallel, along the Arnoldi vector q_{k+1}**.

Arnoldi procedure with Rayleigh-Ritz vectors, cont'd$_2$

- From the fact that $\tilde{r}_\ell = \beta_{k,\ell} q_{k+1}$, the norm of the eigen-residual is such that $\|\tilde{r}_\ell\|_2^2 = |\beta_{k,\ell}|^2 q_{k+1}^T q_{k+1} = |\beta_{k,\ell}|^2$ where $|\beta_{k,\ell}| = |h_{k+1,k}|\,|e_k^T \hat{y}_\ell|$.
  - Consequently, **a stopping criterion of the form** $\|\tilde{r}_\ell\|_2 < \epsilon|\lambda_\ell|$ **can be checked efficiently** at every iteration **without** having to compute the **matrix-vector product** Ay_ℓ or even to assemble the vector $y_\ell := Q_k \hat{y}_\ell$.
- As explained earlier,
  - the **orthogonalization** which is at the root of the Arnoldi procedure has time complexity $\mathcal{O}(k^2 n)$,
  - the **reduced eigensolve** of H_m has time complexity $\mathcal{O}(k^3)$,
  - the **storage of the Arnoldi basis** in Q_k has space complexity $\mathcal{O}(kn)$ so that, **if convergence is not achieved** for some number k of iterations, it is **necessary to start the Arnoldi procedure over** with a new initial vector q_1.
- A **naive restart** of the Arnoldi procedure **can be highly detrimental to** the **convergence** of approximate eigenvectors. Some **care needs to be taken** so as **to reduce convergence slowdown**.

Shift-and-invert Arnoldi Rayleigh-Ritz for interior eigenpairs

Shift-and-invert spectral transformation

- Rayleigh-Ritz pairs (λ, y) converge first towards eigenpairs (θ, z) of A with the largest value of $|\theta|$.
  - In practice, we may want to approximate an eigenpair with eigenvalue θ close to some shift σ, i.e., with small value of $|\sigma - \theta|$. E.g., $\sigma = 0$.
  - Rayleigh-Ritz approximations (λ, y) of such eigenpairs (θ, z) in Krylov subspaces converge very slowly when $|\sigma|$ is small compared to the spectral radius of A.
- The shift-and-invert spectral transformation was introduced by Ericsson and Ruhe (1980) as a means to circumvent this issue:
  - Consider the eigenvalue problem given by

$$(A - \sigma I_n)^{-1} w = \vartheta w$$

where it is assumed that σ is not an eigenvalue of A. Then, we have

$$w = \vartheta(A - \sigma I_n)w$$
$$w = \vartheta A w - \vartheta\sigma w.$$

Ericsson, T., & Ruhe, A. (1980). The spectral transformation Lanczos method for the numerical solution of large sparse generalized symmetric eigenvalue problems. Mathematics of Computation, 35(152), 1251-1268.

Shift-and-invert spectral transformation, cont'd

Since σ is not an eigenvalue of the non-singular matrix A, the shift-and-invert operator $(A - \sigma I_n)^{-1}$ is not singular and $\vartheta \neq 0$ so that

$$Aw = \left(\sigma + \frac{1}{\vartheta}\right) w.$$

Essentially, $(\sigma + 1/\vartheta, w)$ is an eigenpair of A.

- Now, if an Arnoldi procedure is applied to $(A - \sigma I_n)^{-1}$, the corresponding Rayleigh-Ritz pairs will first converge to the eigenpairs (ϑ, w) of the shift-and-invert operator with largest $|\vartheta|$.
- However, when $|\vartheta|$ is maximized, the magnitude of $\sigma - (\sigma + 1/\vartheta)$ is minimized.Therefore, the Rayleigh-Ritz pairs of a shift-and-invert Arnoldi procedure will first converge to the eigenpairs of A with eigenvalues closest to σ.

- Shift-and-invert operators are implemented in ARPACK to compute interior eigenpairs.
- Shift-and-invert Arnoldi procedures rely on repetitive applications of the $(A - \sigma I_n)^{-1}$ operator.

Arnoldi harmonic Ritz for interior eigenpairs

Harmonic Ritz approximation of interior eigenpairs

- While shift-and-invert Arnoldi procedures allow fast convergence of Rayleigh-Ritz pairs towards interior eigenpairs,it comes at the cost of repeated applications of $(A - \sigma I_n)^{-1}$. However:
  1. Factorizing the shifted operator $A - \sigma I_n$ is not always possible.
  2. One may actually need to generate a basis for a Krylov subspace of A, and have little use for a basis of Krylov subspace of the shift-and-invert operator $(A - \sigma I_n)^{-1}$:
     - E.g., if interior eigenvectors of A are needed to restart GMRES when solving $Ax = b$.
- As a means to bypass the need to apply shift-and-invert operators, Morgan (1991) introduces a new projection method in which the shift-and-invert operator is applied implicitly:
  - Consider the case in which we are equipped with a basis for the search space range(P) stored in the columns of P.
  - Let $Q := (A - \sigma I_n)P$, and consider the Rayleigh-Ritz pairs of the shift-and-invert operator $(A - \sigma I_n)^{-1}$ with respect to range(Q).

Morgan, R. B. (1991). Computing interior eigenvalues of large matrices. Linear Algebra and its Applications, 154, 289-309.

Harmonic Ritz approximation of interior eigenpairs, cont'd$_1$

That is, consider the pair $(\vartheta, Q\hat{y})$ such that

$$Q^H(A - \sigma I_n)^{-1} Q\hat{y} = \vartheta Q^H Q\hat{y},$$

which develops into the reduced generalized eigenvalue problem

$$P^T(A - \sigma I_n)^H P\hat{y} = \vartheta P^T(A - \sigma I_n)^H (A - \sigma I_n) P\hat{y}$$

which does not require any application of the shift-and-invert operator.

- Resulting from a Rayleigh-Ritz projection of the shift-and-invert operator $(A - \sigma I_n)^{-1}$, the pair $(\sigma + 1/\vartheta, Q\hat{y})$ should be a good approximation with respect to range(Q) of the eigenpair closest to σ.
- As good of an approximation $Q\hat{y}$ might be, $P\hat{y} = (A - \sigma I_n)^{-1} Q\hat{y}$ is the first power iterate of the shift-and-invert operator initiated with $Q\hat{y}$, so that $P\hat{y}$ should be an even slightly better approximation of the eigenvector with eigenvalue closest to σ.
- Stewart (2001) showed that solutions $(\lambda, Q\hat{y})$ for which $Q\hat{y}$ has unit norm are such that$\|Ay_i\| \leq |\lambda_i|$, so that it is guaranteed that $\|\tilde{r}_i\|_2$ is small if λ_i is near zero.

Morgan, R.B. (1991). Computing interior eigenvalues of large matrices. Linear Algebra and its Applications, 154, 289-309.
G. W. Stewart, Matrix Algorithms II: Eigensystems, SIAM, Philadelphia, (2001).

Harmonic Ritz approximation of interior eigenpairs, cont'd$_2$

- Consequently, Morgan (1991) proposes a Petrov-Galerkin projection and seeks for pairs $(\sigma + \lambda, y)$ to approximate eigenpairs of A near σ with respect to range(P), leading to the following procedure:

  Find λ and $y \in \text{range}(P)$ s.t. $(A - \sigma I_n)y - \lambda y \perp (A - \sigma I_n)\text{range}(P)$,

  which first converges to eigenpairs of A near σ, thus motivating the selection of reduced generalized eigenpairs $(\lambda, \hat{y})$ with smallest values of $|\lambda|$ such that

$$P^T(A - \sigma I_n)^H(A - \sigma I_n)P\hat{y} = \lambda P^T(A - \sigma I_n)^H P\hat{y}.$$

- The projection proposed by Morgan (1991) is first studied for symmetric matrices, then further analyzed and first referred to as harmonic Ritz by Paige et al. (1995) before being considered in the context of non-symmetric eigenvalue problems by Sleijpen and Van der Vorst (1996).

Morgan, R. B. (1991). Computing interior eigenvalues of large matrices. Linear Algebra and its Applications, 154, 289-309.

Paige, C. C., Parlett, B. N., & Van der Vorst, H. A. (1995). Approximate solutions and eigenvalue bounds from Krylov subspaces. Numerical linear algebra with applications, 2(2), 115-133.

Sleijpen, G. L., & Van der Vorst, H. A. (1996). A Jacobi–Davidson Iteration Method for Linear Eigenvalue Problems. Matrix, 17(2), 401-425.

Harmonic Ritz approximation of interior eigenpairs, cont'd$_3$

- To simplify what follows, let us define

$$G_1 := P^T(A-\sigma I_n)^H(A-\sigma I_n)P \text{ and } G_2 := P^T(A-\sigma I_n)P$$

  so that, assuming σ is not an eigenvalue of A, the reduced eigenpair $(\lambda, \hat{y})$ is such that $G_2^{-H}G_1\hat{y} = \lambda\hat{y}$.
- It is well established that the Rayleigh quotient ρ of y with respect to A is a better approximation of the eigenvalue of A near σ than $\sigma + \lambda$. The Rayleigh quotient can be efficiently computed as

$$\rho = \frac{y^HAy}{y^Hy} = \frac{\hat{y}^HP^TAP\hat{y}}{\hat{y}P^TP\hat{y}} = \sigma + \frac{\hat{y}^HP^T(A-\sigma I_n)P\hat{y}}{\hat{y}P^TP\hat{y}} = \sigma + \frac{\hat{y}^HG_2\hat{y}}{\hat{y}P^TP\hat{y}}$$

  so that, if $P^TP = I_k$ and $\hat{y}^H\hat{y} = 1$, then we have $\boxed{\rho = \sigma + \hat{y}^HG_2\hat{y}}$.
- It is also common to monitor convergence through stopping criteria defined with respect to the residual

$$\hat{r} := Ay - \rho y$$

Harmonic Ritz approximation of interior eigenpairs, cont'd$_4$

whose norm can also be efficiently computed as we have

$$\begin{aligned}
\hat{r}^H\hat{r} &= (Ay-\rho y)^H(Ay-\rho y)\\
&= ((A-\sigma I_n)y+(\sigma-\rho)y)^H((A-\sigma I_n)y+(\sigma-\rho)y)\\
&= y^H(A-\sigma I_n)^H(A-\sigma I_n)y+(\sigma-\rho)y^H(A-\sigma I_n)^H y\\
&\quad+\overline{(\sigma-\rho)}y^H(A-\sigma I_n)y+\overline{(\sigma-\rho)}(\sigma-\rho)y^H y\\
&= \hat{y}^H G_1\hat{y}+(\sigma-\rho)\hat{y}^H G_2^H\hat{y}+\overline{(\sigma-\rho)}\hat{y}^H G_2\hat{y}+\overline{(\sigma-\rho)}(\sigma-\rho)y^H y
\end{aligned}$$

where, once again, we assume $P^TP=I_k$ and $\hat{y}^H\hat{y}=1$ so that

$$\begin{aligned}
\hat{r}^H\hat{r} &= \hat{y}^H G_1\hat{y}+(\sigma-\rho)\hat{y}^H G_2^H\hat{y}+\overline{(\sigma-\rho)}\hat{y}^H G_2\hat{y}+\overline{(\sigma-\rho)}(\sigma-\rho)\\
&= \lambda\hat{y}^H G_2^H\hat{y}+(\sigma-\rho)\hat{y}^H G_2^H\hat{y}+\overline{(\sigma-\rho)}(\rho-\sigma)+\overline{(\sigma-\rho)}(\sigma-\rho)\\
&= (\sigma+\lambda-\rho)\hat{y}^H G_2^H\hat{y}\\
&= (\sigma+\lambda-\rho)\overline{\hat{y}^H G_2\hat{y}}
\end{aligned}$$

which leads to $\boxed{\hat{r}^H\hat{r}=(\sigma+\lambda-\rho)\overline{(\rho-\sigma)}}$.

Harmonic Ritz approximation of interior eigenpairs, cont'd$_5$

- The norm of the harmonic residual $\tilde{r} := Ay - (\sigma + \lambda)y$ can also be used to monitor convergence. Still assuming $P^T P = I_k$ and $\hat{y}^H \hat{y} = 1$, we then have

$$\begin{aligned}
\tilde{r}^H \tilde{r} &= (Ay - (\sigma + \lambda)y)^H (Ay - (\sigma + \lambda)y) \\
&= ((A - \sigma I_n)y - \lambda y)^H ((A - \sigma I_n)y - \lambda y) \\
&= \hat{y}^H P^T (A - \sigma I_n)^H (A - \sigma I_n) P\hat{y} - \lambda \hat{y}^H P^T (A - \sigma I_n)^H P\hat{y} \\
&\quad - \overline{\lambda} \hat{y}^H P^T (A - \sigma I_n) P\hat{y} + \lambda\overline{\lambda} \\
&= \hat{y}^H G_1 \hat{y} - \lambda \hat{y}^H G_2^H \hat{y} - \overline{\lambda} \hat{y}^H G_2 \hat{y} + \lambda\overline{\lambda} \\
&= \lambda \hat{y}^H G_2^H \hat{y} - \lambda \hat{y}^H G_2^H \hat{y} - \overline{\lambda} \hat{y}^H G_2 \hat{y} + \lambda\overline{\lambda} \\
&= \overline{\lambda}(\lambda - \hat{y}^H G_2 \hat{y})
\end{aligned}$$

where $\hat{y}^H G_2 \hat{y} = \rho - \sigma$ so that

$$\boxed{\tilde{r}^H \tilde{r} = (\sigma + \lambda - \rho)\overline{\lambda}}.$$

Harmonic Ritz approximation of interior eigenpairs, cont'd$_6$

- Alternatively, the harmonic Ritz pairs may be formed from non-shifted procedures. That is, let

$$G_1 := (AP)^H AP \quad \text{and} \quad G_2 := P^T AP,$$

then, the reduced eigenvalue problem of the shifted approximation

$$y \in \text{range}(P) \quad \text{such that} \quad (A - \sigma I_n)y - \lambda y \perp (A - \sigma I_n)\text{range}(P)$$

is obtained as follows:

$$\begin{aligned}
((A - \sigma I_n)P)^H AP\hat{y} &= (\sigma + \lambda)((A - \sigma I_n)P)^H P\hat{y} \\
[(AP)^H AP - \overline{\sigma} P^T AP]\hat{y} &= (\sigma + \lambda)(P^T A^H P - \overline{\sigma} I_k)\hat{y} \\
(G_1 - \overline{\sigma} G_2)\hat{y} &= (\sigma + \lambda)(G_2^H - \overline{\sigma} I_k)\hat{y} \\
(G_1 - \overline{\sigma} G_2)\hat{y} &= (\sigma + \lambda)(G_2 - \sigma I_k)^H \hat{y} \\
(G_2 - \sigma I_k)^{-H}(G_1 - \overline{\sigma} G_2)\hat{y} &= (\sigma + \lambda)\hat{y}
\end{aligned}$$

in which case we have $\rho = \hat{y}^H G_2 \hat{y}$.

Arnoldi procedure with harmonic Ritz vectors

- More can be said for the case in which the search space is Krylov and generated by an Arnoldi procedure, see Morgan and Zheng (1998).
- Consider the shifted procedure $\texttt{Arnoldi}(A - \sigma I_n,\, q_1,\, k) \mapsto (Q_{k+1}, \underline{H}_k)$ which returns an orthonormal basis $Q_k := [q_1 \,\ldots\, q_k]$ of $\mathcal{K}_k(A - \sigma I_n, q_1)$ such that $(A - \sigma I_n)Q_k = Q_{k+1}\underline{H}_k$ where $Q_{k+1} := [Q_k\, q_{k+1}]$ as well as $\underline{H}_k = Q_{k+1}^T(A - \sigma I_n)Q_k$ and $H_k = Q_k^T(A - \sigma I_n)Q_k$.
- Then, harmonic Ritz vectors $y \in \text{range}(Q_k)$ are such that

$$(A - \sigma I_n)y - \lambda y \perp (A - \sigma I_n)\text{range}(Q_k)$$

yields the following reduced generalized eigenvalue problem in which we search for non-trivial pairs $(\lambda, \hat{y}) \in \mathbb{C} \times \mathbb{C}^k$ such that $y = Q_k\hat{y}$ and

$$\begin{aligned} Q_k^T(A - \sigma I_n)^H(A - \sigma I_n)Q_k\hat{y} &= \lambda Q_k^T(A - \sigma I_n)^H Q_k\hat{y} \\ \underline{H}_k^H Q_{k+1}^T Q_{k+1}\underline{H}_k\hat{y} &= \lambda H_k^H\hat{y} \\ \underline{H}_k^H\underline{H}_k\hat{y} &= \lambda H_k^H\hat{y} \end{aligned}$$

Morgan, R. B., & Zeng, M. (1998). Harmonic projection methods for large non-symmetric eigenvalue problems. Numerical linear algebra with applications, 5(1), 33-55.

Arnoldi procedure with harmonic Ritz vectors, cont'd$_1$

- Reformulating the Arnoldi relation into

$$(A - \sigma I_n)Q_k = Q_k H_k + h_{k+1,k} q_{k+1} e_k^T$$

allows to rewrite the reduced eigenvalue problem of the harmonic Ritz projection as follows:

$$(H_k^H H_k + |h_{k+1,k}|^2 e_k e_k^T)\hat{y} = \lambda H_k^H \hat{y}$$

$$\boxed{(H_k + |h_{k+1,k}|^2 f e_k^T)\hat{y} = \lambda \hat{y}}$$

where $f := H_k^{-H} e_k \in \mathbb{C}^k$ and $e_k := I_k[:, k]$.

- Then, the harmonic eigen-residual $\tilde{r} := Ay - (\sigma + \lambda)y$ of a given harmonic Ritz approximate eigenpair $(\sigma + \lambda, y)$ with $y := Q_k \hat{y}$ is such that

$$\begin{aligned}
\tilde{r} &= AQ_m\hat{y} - (\sigma + \lambda)Q_m\hat{y} \\
&= (A - \sigma I_n)Q_k\hat{y} - \lambda Q_k\hat{y} = (Q_k H_k + h_{k+1,k} q_{k+1} e_k^T)\hat{y} - \lambda Q_k\hat{y} \\
&= Q_k(\lambda I_n - |h_{k+1,k}|^2 f e_k^T)\hat{y} + h_{k+1,k} q_{k+1} e_k^T \hat{y} - \lambda Q_k\hat{y} \\
&= h_{k+1,k}(e_k^T\hat{y})q_{k+1} - |h_{k+1,k}|^2 (e_k^T\hat{y})Q_k f
\end{aligned}$$

Arnoldi procedure with harmonic Ritz vectors, cont'd$_2$

which can be written

$$\begin{aligned}
\tilde{r} &= h_{k+1,k}(e_k^T\hat{y})Q_{k+1}\begin{bmatrix}0_{k\times 1}\\ 1\end{bmatrix} - |h_{k+1,k}|^2(e_k^T\hat{y})Q_{k+1}\begin{bmatrix}f\\ 0\end{bmatrix}\\
&= h_{k+1,k}(e_k^T\hat{y})Q_{k+1}\begin{bmatrix}0_{k\times 1}\\ 1\end{bmatrix} + h_{k+1,k}(e_k^T\hat{y})Q_{k+1}\begin{bmatrix}-\overline{h_{k+1,k}}f\\ 0\end{bmatrix}\\
&= h_{k+1,k}(e_k^T\hat{y})Q_{k+1}\begin{bmatrix}-\overline{h_{k+1,k}}f\\ 1\end{bmatrix}\\
&= \beta_k Q_{k+1}s
\end{aligned}$$

so that

$$\boxed{\tilde{r} = \beta_k Q_{k+1}s} \text{ where } \boxed{\beta_k := h_{k+1,k}e_k^T\hat{y}} \text{ and } \boxed{s := \begin{bmatrix}-\overline{h_{k+1,k}}f\\ 1\end{bmatrix}}.$$

- The norm of $\tilde{r}$ is then given by $\boxed{\|\tilde{r}\|_2 = |\beta_k|(|h_{k+1,k}|^2 f^H f + 1)^{1/2}}$.

Arnoldi procedure with harmonic Ritz vectors, cont'd$_3$

- When precise eigenvalues are wanted, it is preferred to use the Rayleigh quotient rather than $\sigma + \lambda$. Assuming $\hat{y}$ has unit norm, so does $y := Q_k \hat{y}$, and the Rayleigh quotient of y is given by

$$\rho = \sigma + \hat{y}^H H_k \hat{y} = \sigma + \lambda - |h_{k+1,k}|^2 (\hat{y}^H f)(e_k^T \hat{y}).$$

- Moreover, the norm of the eigen-residual $\hat{r} := Ay - \rho y$ is still such that

$$\|\hat{r}\|_2^2 = (\sigma + \lambda - \rho)\overline{(\rho - \sigma)}$$

and that of the harmonic residual $\tilde{r} := Ay - (\sigma + \lambda)y$ is still such that

$$\|\tilde{r}\|_2^2 = (\sigma + \lambda - \rho)\overline{\lambda}.$$

Arnoldi procedure with harmonic Ritz vectors, cont'd$_4$

- As mentioned before, we are interested by the case in which harmonic Ritz approximations are considered in the context of the non-shifted procedure $\texttt{Arnoldi}(A, q_1, k) \mapsto (Q_{k+1}, \underline{H}_k)$ which returns an orthonormal basis $Q_k := [q_1 \dots q_k]$ of $\mathcal{K}_k(A, q_1)$ such that $AQ_k = Q_{k+1}\underline{H}_k$ where $Q_{k+1} := [Q_k\, q_{k+1}]$, $\underline{H}_k = Q_{k+1}^T A Q_k$ and $H_k = Q_k^T A Q_k$.
  - Then, the harmonic Ritz vector $y \in \text{range}(Q_k)$ is still such that

$$(A - \sigma I_n)y - \lambda y \perp (A - \sigma I_n)\text{range}(Q_k)$$

but now yields the following reduced generalized eigenvalue problem:

$$((A - \sigma I_n)Q_k)^H (A - \sigma I_n)Q_k \hat{y} = \lambda((A - \sigma I_n)Q_k)^H Q_k \hat{y}$$
$$((A - \sigma I_n)Q_k)^H AQ_k \hat{y} - \sigma((A - \sigma I_n)Q_k)^H Q_k \hat{y} = \lambda((A - \sigma I_n)Q_k)^H Q_k \hat{y}$$

so that

$$((A - \sigma I_n)Q_k)^H AQ_k \hat{y} = (\sigma + \lambda)((A - \sigma I_n)Q_k)^H Q_k \hat{y}$$

Arnoldi procedure with harmonic Ritz vectors, cont'd$_5$

which develops as follows:

$$\begin{aligned}
((AQ_k)^T AQ_k - \overline{\sigma} Q_k^T AQ_k)\hat{y} &= (\sigma + \lambda)(Q_k^T A^T Q_k - \overline{\sigma} I_k)\hat{y} \\
(H_k^T H_k + |h_{k+1,k}|^2 e_k e_k^T - \overline{\sigma} H_k)\hat{y} &= (\sigma + \lambda)(H_k^T - \overline{\sigma} I_k)\hat{y} \\
((H_k^T - \overline{\sigma} I_k) H_k + |h_{k+1,k}|^2 e_k e_k^T)\hat{y} &= (\sigma + \lambda)(H_k - \sigma I_k)^H \hat{y} \\
((H_k - \sigma I_k)^H H_k + |h_{k+1,k}|^2 e_k e_k^T)\hat{y} &= (\sigma + \lambda)(H_k - \sigma I_k)^H \hat{y}
\end{aligned}$$

to finally yield

$$\boxed{(H_k + |h_{k+1,k}|^2 f e_k^T)\hat{y} = (\sigma + \lambda)\hat{y}} \text{ where } \boxed{f := (H_k - \sigma I_k)^{-H} e_k},$$

so that the expression for f differs from the shifted procedure. Still, the harmonic Ritz pairs should converge first to the eigenpairs of A closest to σ so that, now, we should not retain the least dominant reduced eigenpairs, but rather those with eigenvalues closest to σ.

Arnoldi procedure with harmonic Ritz vectors, cont'd$_6$

- Now, still assuming $\hat{y}^H \hat{y} = 1$, the Rayleigh quotient is given by

$$\rho = y^H A y = \hat{y}^H Q_k^T A Q_k \hat{y} = \hat{y}^H H_k \hat{y} = \sigma + \lambda - |h_{k+1,k}|^2 (\hat{y}^H f)(e_k^T \hat{y}).$$

- Irrespective of the basis generated, as long as it's orthonormal, we already saw the residual given by $\hat{r} := Ay - \rho y$ is such that

$$\hat{r}^H \hat{r} = (\sigma + \lambda - \rho)\overline{(\rho - \sigma)}.$$

- And the harmonic eigen-residual $\tilde{r} := Ay - (\sigma + \lambda)y$ is such that

$$\begin{aligned}
\tilde{r} &= AQ_k\hat{y} - (\sigma + \lambda)Q_k\hat{y} \\
&= (Q_k H_k + h_{k+1,k} q_{k+1} e_k^T)\hat{y} - (\sigma + \lambda)Q_k\hat{y} \\
&= Q_k H_k \hat{y} + h_{k+1,k} q_{k+1} e_k^T \hat{y} - (\sigma + \lambda)Q_k\hat{y} \\
&= (\sigma + \lambda)Q_k\hat{y} - |h_{k+1,k}|^2 Q_k f e_k^T \hat{y} + h_{k+1,m} q_{k+1} e_k^T \hat{y} - (\sigma + \lambda)Q_k\hat{y} \\
&= -|h_{k+1,k}|^2 (e_k^T \hat{y}) Q_k f + h_{k+1,k} q_{k+1} e_m^T \hat{y}
\end{aligned}$$

  which, similarly as before, can be recast into

$$\tilde{r} = \beta_k Q_{k+1} s \quad \text{where} \quad \beta_k := h_{k+1,k} e_k^T \hat{y} \quad \text{and} \quad s = \begin{bmatrix} -\overline{h_{k+1,k}} f \\ 1 \end{bmatrix}.$$

Lanczos process

Section 6.2 in Darve & Wootters (2021)

Lanczos process for symmetric matrices

- The Lanczos process is a specialized form of the Arnoldi process for **symmetric matrices**.
- When A is symmetric (i.e., $A = A^T$), the Hessenberg matrix $H_k = Q_k^T A Q_k$ is symmetric too.Consequently, it is **tridiagonal**:

$$T_k = \begin{bmatrix} \alpha_1 & \beta_1 & 0 & \cdots & 0 \\ \beta_1 & \alpha_2 & \beta_2 & \cdots & 0 \\ 0 & \beta_2 & \alpha_3 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & \alpha_k \end{bmatrix}$$

  where $\alpha_i = q_i^T A q_i$ are the diagonal elements and $\beta_i = q_i^T A q_{i+1} = q_{i+1}^T A q_i$ are the off-diagonal elements.
- This tridiagonal structure means that in the Arnoldi recurrence relation ($AQ_k = Q_{k+1}\underline{T}_k$), most terms vanish:

$$Aq_j = \beta_{j-1} q_{j-1} + \alpha_j q_j + \beta_j q_{j+1}.$$

- This three-term recurrence relation is the foundation of the Lanczos process.

Derivation of the Lanczos process

- From the three-term recurrence relation, we can derive the Lanczos algorithm:

$$\begin{aligned}\beta_{j-1}q_{j-1} + \alpha_j q_j + \beta_j q_{j+1} &= Aq_j \\ \beta_j q_{j+1} &= Aq_j - \alpha_j q_j - \beta_{j-1}q_{j-1} \\ q_{j+1} &= \frac{1}{\beta_j}(Aq_j - \alpha_j q_j - \beta_{j-1}q_{j-1}).\end{aligned}$$

  The coefficients are determined as:

$$\alpha_j = q_j^T A q_j \quad \text{and} \quad \beta_j = \|Aq_j - \alpha_j q_j - \beta_{j-1}q_{j-1}\|_2$$

- This leads to a much simpler algorithm compared to the full Arnoldi process:
  - we only need to maintain three vectors in memory at any time: q_{j-1}, q_j, and q_{j+1}.
  - the work done remains constant as the iteration count increases.

Lanczos algorithm

- The Lanczos algorithm can be formulated as follows:

Algorithm 5 Lanczos

1: Choose a starting vector q_1 with $\|q_1\|_2 = 1$
2: Set $\beta_0 = 0$ and $q_0 = 0$
3: **for** $j = 1, 2, \ldots, k$ **do**
4: $\quad v = Aq_j$
5: $\quad \alpha_j = q_j^T v$
6: $\quad v = v - \alpha_j q_j - \beta_{j-1} q_{j-1}$
7: $\quad \beta_j = \|v\|_2$
8: $\quad q_{j+1} = v/\beta_j$

- At each iteration, the procedure require a fixed number of BLAS-1 operations along with a matrix-vector product.
- After k steps, we have:
  - An orthonormal basis $Q_k = [q_1, q_2, \ldots, q_k]$ for the subspace $\mathcal{K}_k(A, q_1)$,
  - A tridiagonal matrix $T_k = Q_k^T A Q_k$ with diagonal elements α_i and off-diagonal elements β_i.

The Lanczos relation

- Similar to the Arnoldi relation, we have the Lanczos relation:

$$AQ_k = Q_k T_k + \beta_k q_{k+1} e_k^T$$

where T_k is the tridiagonal matrix.

- We can also write:

$$AQ_k = Q_{k+1} \underline{T}_k$$

where $\underline{T}_k$ is the $(k+1) \times k$ tridiagonal matrix:

$$\underline{T}_k = \begin{bmatrix} \alpha_1 & \beta_1 & 0 & \cdots & 0 \\ \beta_1 & \alpha_2 & \beta_2 & \cdots & 0 \\ 0 & \beta_2 & \alpha_3 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & \alpha_k \\ 0 & 0 & 0 & \cdots & \beta_k \end{bmatrix}.$$

Lanczos Rayleigh-Ritz for dominant eigenpairs

- **Eigenvectors with eigenvalues whose norms are the largest** among the spectrum of A are **well approximated by Rayleigh-Ritz** projections.
- A **Rayleigh-Ritz vector** $y \in \text{range}(Q_k)$ approximates an eigenvector of A with the **Ritz value** λ such that $Ay - \lambda y \perp \text{range}(Q_k)$. That is, we **search for** $(\lambda, \hat{y}) \in \mathbb{R} \times \mathbb{R}^k \setminus \{0\}$ s.t. $z^T (Ay - \lambda y) = 0 \; \forall \; z \in \text{range}(Q_k)$ with $y = Q_k \hat{y}$. This **simplifies to**

$$Q_k^T (AQ_k\hat{y} - \lambda Q_k \hat{y}) = 0$$

$$T_k\hat{y} - \lambda\hat{y} = 0 \quad \Longrightarrow \quad \boxed{T_k\hat{y} = \lambda\hat{y}}$$

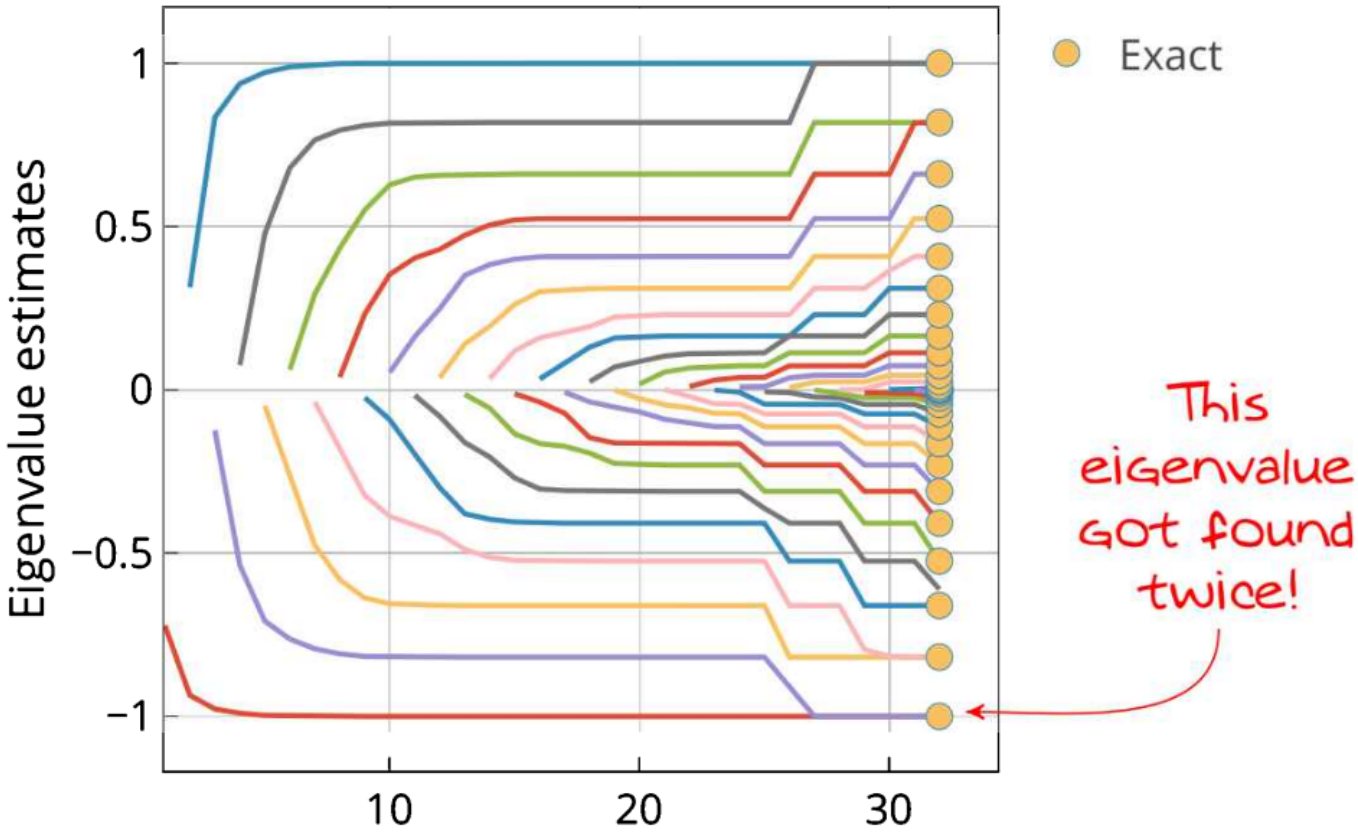
  where use is made of the Lanczos relation and $Q_k^T Q_k = I_k$.
- The eigen-residual $\tilde{r} := Ay - \lambda y$ can be computed as:

$$\tilde{r} = AQ_k\hat{y} - \lambda Q_k\hat{y} = \beta_k (e_k^T \hat{y}) q_{k+1}.$$

  This means $\|\tilde{r}\|_2 = \beta_k |e_k^T \hat{y}|$, thus providing a simple way to assess convergence without explicitly computing Ay.

Lanczos Rayleigh-Ritz in finite precision

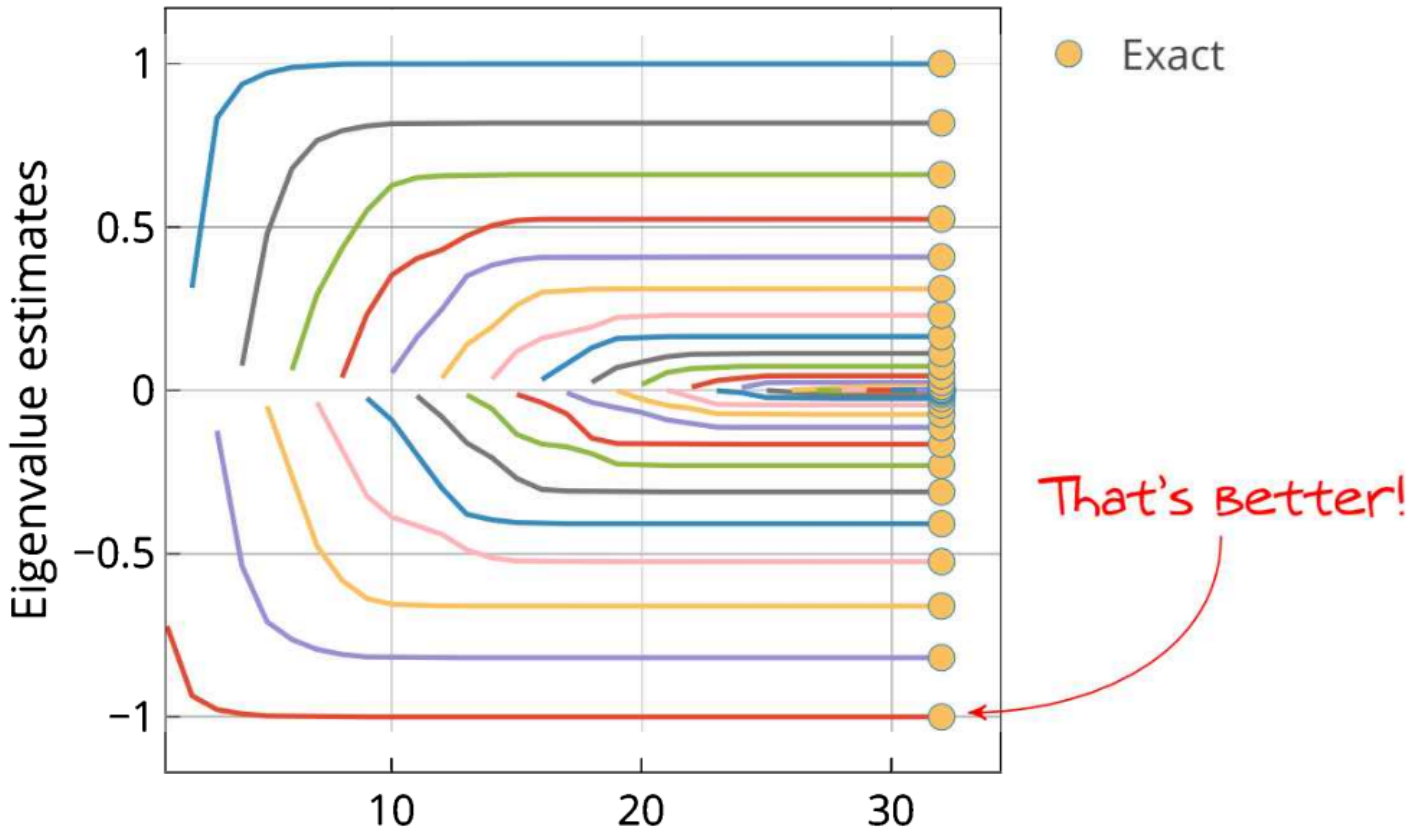
- **Loss of orthogonality:** In finite precision arithmetic, the Lanczos vectors quickly lose orthogonality, which can lead to:
  - Multiple copies of the same eigenvalue appearing (ghost eigenvalues)
  - Inaccurate eigenvalue approximations



Darve, E., & Wootters, M. (2021). Numerical linear algebra with Julia. Society for Industrial and Applied Mathematics.

Reorthogonalization strategies

- Different strategies exist to circumvent the issue of loss of orthogonality in finite precision.
  - **Full reorthogonalization**: Explicitly orthogonalize each new vector against all previous vectors.



Darve, E., & Wootters, M. (2021). Numerical linear algebra with Julia. Society for Industrial and Applied Mathematics.

Reorthogonalization strategies, cont'd

- The most straightforward fix to loss of orthogonality is full reorthogonalization:

Algorithm 6 Lanczos with full reorthogonalization

1: Choose a starting vector q_1 with $\|q_1\|_2 = 1$
2: Set $\beta_0 = 0$ and $q_0 = 0$
3: **for** $j = 1, 2, \ldots, k$ **do**
4: $\quad v = Aq_j$
5: $\quad \alpha_j = q_j^T v$
6: $\quad v = v - \alpha_j q_j - \beta_{j-1} q_{j-1}$
7: $\quad$ **for** $i = 1, 2, \ldots, j$ **do**
8: $\quad\quad v = v - (q_i^T v) q_i$ ▷ Reorthogonalization step
9: $\quad \beta_j = \|v\|_2$
10: $\quad q_{j+1} = v/\beta_j$

Full reorthogonalization turns Lanczos back into Arnoldi. Alternatives are

- **Selective reorthogonalization:** Only reorthogonalize when necessary, based on loss of orthogonality measures,
- **Partial reorthogonalization:** Reorthogonalize against a subset of previous vectors.

Lanczos harmonic Ritz for interior eigenpairs

- **Eigenvectors with interior eigenvalues** are **better approximated by harmonic Ritz** projections.
- A **harmonic Ritz vector** $y \in \text{range}(Q_k)$ approximates an eigenvector of A with the **harmonic Ritz value** $\sigma + \lambda$ in the vicinity of σ such that

$$(A - \sigma I_n)y - \lambda y \perp (A - \sigma I_n)\text{range}(Q_k).$$

Let the columns of Q_k form an orthonormal basis of $\mathcal{K}_k(A, q_1)$. Then, the **search for** for the pair $(\sigma + \lambda, \hat{y})$ is such that

$$\begin{aligned}
Q_k^T(A - \sigma I_n)^T(A - \sigma I_n)Q_k\hat{y} &= \lambda Q_k^T(A - \sigma I_n)^T Q_k\hat{y} \\
((AQ_k)^T AQ_k - \sigma Q_k^T AQ_k)\hat{y} &= (\sigma + \lambda)(Q_k^T AQ_k - \sigma I_k)\hat{y} \\
(T_k T_k + \beta_k^2 e_k e_k^T - \sigma T_k)\hat{y} &= (\sigma + \lambda)(T_k - \sigma I_k)\hat{y} \\
((T_k - \sigma I_k)T_k + \beta_k^2 e_k e_k^T)\hat{y} &= (\sigma + \lambda)(T_k - \sigma I_k)\hat{y} \\
(T_k + \beta_k^2 f e_k^T)\hat{y} &= (\sigma + \lambda)\hat{y}
\end{aligned}$$

where $f := (T_k - \sigma I_k)^{-1} e_k$.

Lanczos harmonic Ritz for interior eigenpairs

- Note that fe_k^T is not symmetric so that the **reduced eigenpairs of** the **harmonic Ritz procedure** are **generally complex**.

  However, they do **converge towards** the **real eigenpairs**.
- Similarly as with Arnoldi, the Rayleigh quotient is given by

$$\rho = \sigma + \lambda - \beta_k^2 (\hat{y}^H f)(e_k^T \hat{y})$$

whereas the residual $\hat{r} := Ay - \rho y$ is such that

$$\hat{r}^H \hat{r} = (\sigma + \lambda - \rho)\overline{(\rho - \sigma)}.$$

Summary of Krylov subspace methods

- We studied two main Krylov subspace methods for eigenvalue problems:
  - **Arnoldi process:** For general matrices, produces a Hessenberg matrix H_k, requires orthogonalization against all previously formed vectors.
  - **Lanczos process:** For symmetric matrices, produces a tridiagonal matrix T_k, relies on a three terms recurence formula.
- Both methods:
  - Construct an orthonormal basis for the Krylov subspace $\mathcal{K}_k(A, v)$.
  - Can be used with either Rayleigh Ritz or harmonic Ritz projections.
- Key advantages of Krylov subspace methods:
  - Only require matrix-vector products, ideal for large sparse matrices.
  - Can find several eigenvalues simultaneously.
- Modern implementations use:
  - Restarting techniques to limit memory requirements and increasing computational cost (for Arnoldi).
  - Reorthogonalization strategies for numerical stability of Lanczos.