

0.1 Iterative methods

As seen in the previous lecture, direct methods can theoretically compute exact solutions $\mathbf{x} \in \mathbb{R}^N$ to linear systems in the form of:

$$\mathbf{A}\mathbf{x} = \mathbf{b}$$

with matrix with real coefficients $\mathbf{A} \in M_N(\mathbb{R})$ and given data $\mathbf{b} \in \mathbb{R}^N$, in a determined finite number of steps.

As computing the inverse of the matrix is unrealistic, several methods were introduced based on factorizations of the type $\mathbf{A} = \mathbf{P}\mathbf{Q}$ where \mathbf{P} and \mathbf{Q} have a structure simplifying the resolution of the system: diagonal, banded, triangular.

Methods like LU, Cholevski take advantage of the existence of a decomposition involving triangular matrices while QR for example, involves the construction of an orthogonal basis. All methods prove to be quite expensive, hard to parallelize due to the sequential nature of the algorithm and prone to error propagation.

Iterative methods have been developed for:

- solving very large linear systems with direct methods is in practice not possible due to the complexity in term of computational operations and data,
- taking advantage of sparse system for which the structure of the matrix can result in dramatic speed-up (this is the case for numerical schemes for PDEs),
- using the fact that some systems like PDEs discretizations are already formulated in an iterative fashion.

In this section, we discuss briefly the computational properties of iterative methods for solving linear systems. Computing the exact solution is not a requirement anymore but instead the algorithm is supposed to converge asymptotically to the exact solution: the algorithm is stopped when the approximate solution is deemed *close enough* to the exact solution in a sense to be defined. A parameter used as stopping criterion triggers the completion of the algorithm.

The general idea of these methods is to introduce a splitting of the form:

$$\mathbf{A} = \mathbf{G} - \mathbf{H}$$

such the solution \mathbf{x} satisfies:

$$\mathbf{G}\mathbf{x} = \mathbf{b} + \mathbf{H}\mathbf{x}$$

Similarly to fixed-point methods we can define a sequence of approximate solutions (\mathbf{x}^k) satisfying relations of the form:

$$\mathbf{G}\hat{\mathbf{x}}^{k+1} = \mathbf{b} + \mathbf{H}\hat{\mathbf{x}}^k$$

with G invertible.

The matrix viewed as a linear mapping in \mathbb{R}^N , the counterpart of such approaches is given by the Brouwer Theorem in finite dimension, where a continuous mapping $f : \Omega \rightarrow \Omega$ with Ω compact of \mathbb{R}^N admits a fixed-point \mathbf{x}^* satisfying $f(\mathbf{x}^*) = \mathbf{x}^*$ and is contracting.

Methods introduced depend on the iteration defined by the splitting and call for several questions regarding the computational aspects:

1. How can the convergence be ensure?
2. How fast is the convergence?
3. How expensive is each iteration?
4. How does the algorithm behave with respect to numerical error?

The question of the convergence is addressed by proving an estimate on error vectors in terms of iteration error $\hat{\boldsymbol{\epsilon}}^k = \hat{\mathbf{x}}^{k+1} - \hat{\mathbf{x}}^k$ or global error: $\boldsymbol{\epsilon}^k = \hat{\mathbf{x}}^k - \mathbf{x}$. The convergence rate α means that $C > 0$, $|\boldsymbol{\epsilon}^{k+1}| \leq C|\boldsymbol{\epsilon}^k|^\alpha$.

For example, substituting $\hat{\mathbf{x}}^{k+1} = G^{-1}\mathbf{b} + G^{-1}H\hat{\mathbf{x}}^k$ in $\hat{\boldsymbol{\epsilon}}^k = \hat{\mathbf{x}}^{k+1} - \hat{\mathbf{x}}^k$ gives a relation between successive iteration errors:

$$\hat{\boldsymbol{\epsilon}}^k = G^{-1}H \hat{\boldsymbol{\epsilon}}^{k-1}$$

with $M = G^{-1}H$ the iteration matrix, and recursively $\hat{\boldsymbol{\epsilon}}^k = (G^{-1}H)^{k+1}\hat{\boldsymbol{\epsilon}}^0$. Convergence is then conditioned to the existence of a contraction factor $K < 1$ such that $\|\hat{\boldsymbol{\epsilon}}^k\|_\infty \leq K \|\hat{\boldsymbol{\epsilon}}^{k-1}\|_\infty$ ensuring decrease of the error.

In terms of the matrix M , this translate for the spectral radius $\rho(M)$ as $\rho(M) < 1$ since in that case $\lim_{k \rightarrow \infty} M^k \hat{\boldsymbol{\epsilon}}^0 = 0_{\mathbb{R}^N}$. The smaller the spectral radius, the faster the convergence.

Each method is described briefly and qualitatively with just the necessary ingredients to discuss practical implementations.

0.2 Relaxation methods

Consider the relations for each row $i = 1, \dots, N$:

$$\mathbf{x}_i = \frac{1}{a_{ii}} \left(b_i - \sum_{i \neq j} a_{ij} \mathbf{x}_j \right) \quad (1)$$

Let us introduce two methods based on constructing sequences of approximate solutions $(\hat{\mathbf{x}}^k)$, $k \geq 1$ given an initial guess $\hat{\mathbf{x}}^0 \in \mathbb{R}^N$ and then associated relaxation methods.

0.2.1 Jacobi, methods of simultaneous displacements

$$\hat{\mathbf{x}}_i^{k+1} = \frac{1}{a_{ii}} \left(b_i - \sum_{i \neq j} a_{ij} \hat{\mathbf{x}}_j^k \right) \quad (2)$$

Convergence: the global error $\boldsymbol{\epsilon}^k$ is controlled by

$$\|\boldsymbol{\epsilon}^{k+1}\| \leq \sum_{i \neq j} \left| \frac{a_{ij}}{a_{ii}} \right| \|\boldsymbol{\epsilon}^k\| \leq K^k \|\boldsymbol{\epsilon}^1\|$$

It is then enough if the matrix is strictly diagonally dominant. Expressing the iteration error gives directly that $M = G^{-1}H$ such that $\rho(M) < 1$.

Algorithm: the splitting is

$$A = D - H$$

with $D = \text{diag}(A)$, thus

$$\hat{\mathbf{x}}^{k+1} = D^{-1}(\mathbf{b} + H\hat{\mathbf{x}}^k)$$

Implementation:

1. Parallelization component by component is possible since there is only dependency on $\hat{\mathbf{x}}^k$.
2. Memory requirement for storing both $\hat{\mathbf{x}}^{k+1}$ and $\hat{\mathbf{x}}^k$ at each iteration.

0.2.2 Gauss–Seidel, methods of successive displacements

In Jacobi iterations, notice that sequential ordered computation of terms

$$\hat{\mathbf{x}}_i^{k+1} = \frac{1}{a_{ii}} \left(b_i - \sum_{i \neq j} a_{ij} \hat{\mathbf{x}}_j^k \right) \quad (3)$$

involves components $\hat{\mathbf{x}}_j^k$ which are also computed for $\hat{\mathbf{x}}^{k+1}$ if $j < i$.

$$\hat{\mathbf{x}}_i^{k+1} = \frac{1}{a_{ii}} \left(b_i - \sum_{i < j} a_{ij} \hat{\mathbf{x}}_j^{k+1} - \sum_{i > j} a_{ij} \hat{\mathbf{x}}_j^k \right) \quad (4)$$

Algorithm: the splitting is

$$A = L - R_0$$

with $L = D + L_0$ lower-triangular matrix and R_0 strict upper-triangular matrix, thus

$$\hat{\mathbf{x}}^{k+1} = D^{-1}(\mathbf{b} - L_0\hat{\mathbf{x}}^{k+1} + R_0\hat{\mathbf{x}}^k)$$

or

$$L\hat{\mathbf{x}}^{k+1} = D^{-1}(\mathbf{b} + R_0\hat{\mathbf{x}}^k)$$

Recast under the usual form:

$$\hat{\mathbf{x}}^{k+1} = \mathbf{L}^{-1}(\mathbf{b} + \mathbf{R}_0\hat{\mathbf{x}}^k)$$

and the iteration matrix is $\bar{\mathbf{M}} = \mathbf{L}^{-1}\mathbf{R}_0$.

Convergence: the global error $\boldsymbol{\epsilon}^k$ is controlled by

$$\|\boldsymbol{\epsilon}^{k+1}\| \leq \frac{\sum_{i>j} \left| \frac{a_{ij}}{a_{ii}} \right|}{1 - \sum_{i<j} \left| \frac{a_{ij}}{a_{ii}} \right|} \|\boldsymbol{\epsilon}^k\| \leq \bar{K}^k \|\boldsymbol{\epsilon}^1\|$$

If the Jacobi contraction factor $K < 1$ then $\bar{K} < 1$. Expressing the iteration error gives directly that $\bar{\mathbf{M}} = \mathbf{L}^{-1}\mathbf{R}_0$ such that $\rho(\bar{\mathbf{M}}) < 1$.

Implementation:

1. Parallelization component by component is not possible easily since there is serialization for each row i due to the dependency on $\hat{\mathbf{x}}_j^{k+1}$, $j < i$.
2. Memory requirement is only for storing one vector of \mathbb{R}^N at each iteration.

0.2.3 Relaxation of Jacobi and Gauss-Seidel

Relaxation methods consists of adding a linear combination of the approximate solution at the previous iteration to minimize the spectral radius for convergence, using the relaxation parameter $\gamma \in (0, 1)$.

1. Jacobi Over-Relaxation (JOR):

$$\hat{\mathbf{x}}_i^{k+1} = (1 - \gamma) \hat{\mathbf{x}}_i^k + \gamma \frac{1}{a_{ii}} \left(b_i - \sum_{i \neq j} a_{ij} \hat{\mathbf{x}}_j^k \right) \quad (5)$$

which reads in matricial form

$$\hat{\mathbf{x}}^{k+1} = \mathbf{M}_\gamma \hat{\mathbf{x}}^k + \gamma \mathbf{D}^{-1} \mathbf{b}$$

with $\mathbf{M}_\gamma = (1 - \gamma)\mathbb{I} + \gamma \mathbf{D}^{-1}\mathbf{H}$

2. Successive Over-Relaxation (SOR):

$$\hat{\mathbf{x}}_i^{k+1} = (1 - \gamma) \hat{\mathbf{x}}_i^k + \gamma \frac{1}{a_{ii}} \left(b_i - \sum_{i<j} a_{ij} \hat{\mathbf{x}}_j^{k+1} - \sum_{i>j} a_{ij} \hat{\mathbf{x}}_j^k \right) \quad (6)$$

which reads in matricial form

$$\hat{\mathbf{x}}^{k+1} = \mathbf{M}_\gamma \hat{\mathbf{x}}^k + \gamma \mathbf{C} \mathbf{b}$$

with $\mathbf{M}_\gamma = (1 + \gamma \mathbf{D}^{-1} \mathbf{L}_0^{-1})^{-1} [(1 - \gamma)\mathbb{I} + \gamma \mathbf{D}^{-1} \mathbf{R}_0]$ and $\mathbf{C} = (1 + \gamma \mathbf{D}^{-1} \mathbf{L}_0^{-1})^{-1} \mathbf{D}^{-1}$

The relaxation parameter γ cannot be known *a priori* and is usually determined by heuristics.

0.2.4 Parallelization of Gauss–Seidel

Overcoming the serialization in Gauss–Seidel is possible if the matrix is sparse. Taking advantage of the fact that components does not all possess connectivities with each other: such dependencies can be built from the sparsity pattern then decoupled graphs identified:

1. Component Dependency-Graph: generate a graph to reorder entries such that dependencies are avoided.
2. Red–Black coloring: special case for two-dimensional problems.

0.3 Krylov-subspace methods

The idea of these methods is that the solution is decomposed on a sequence of orthogonal subspaces.

If A is symmetric definite positive it induces the corresponding scalar product:

$$\langle \mathbf{x}, \mathbf{y} \rangle = (\mathbf{Ax}, \mathbf{y}) = \mathbf{y}^T \mathbf{Ax}$$

with $(\mathbf{A}\cdot, \cdot)$ canonical scalar product in \mathbb{R}^N . The vectors $(\mathbf{e}_1, \dots, \mathbf{e}_N)$ are said A -conjugate if $\mathbf{e}_j^T \mathbf{A} \mathbf{e}_i = 0$ for $i \neq j$: they are orthogonal for the scalar-product induced by A .

0.3.1 Principle of descent methods: Steepest Gradient

Minimisation of the residual:

$$\mathbf{x}^* = \operatorname{argmin}_{\mathbf{x}} J(\mathbf{x}) = \frac{1}{2} \langle \mathbf{x}, \mathbf{x} \rangle - \langle \mathbf{b}, \mathbf{x} \rangle$$

Construct a sequence of solutions to approximate minimization problems, given $\hat{\mathbf{x}}^k$:

$$J(\hat{\mathbf{x}}^{k+1}) \leq J(\hat{\mathbf{x}}^k)$$

where $\hat{\mathbf{x}}^{k+1} = \hat{\mathbf{x}}^k + \alpha_{k+1} \mathbf{e}^{k+1}$, with α_{k+1} a descent factor and \mathbf{e}_{k+1} a direction.

For the Steepest Gradient:

1. take the direction given by $-\nabla J(\hat{\mathbf{x}}^k) = \mathbf{b} - \mathbf{A}\hat{\mathbf{x}}^k$ which is the residual $\mathbf{r}_k = \mathbf{b} - \mathbf{A}\hat{\mathbf{x}}^k$, thus $\hat{\mathbf{x}}^{k+1} = \hat{\mathbf{x}}^k + \alpha_{k+1} \mathbf{r}_k$.
2. choose the descent factor α^{k+1} minimizing the functional $J(\hat{\mathbf{x}}^k + \alpha_{k+1} \mathbf{r}_k)$:

$$\alpha_{k+1} = \frac{\mathbf{r}_k^T \mathbf{b}}{\mathbf{r}_k^T \mathbf{A} \mathbf{r}_k}$$

The speed of convergence is bounded by $\mathcal{O}(1 - \mathcal{C}(\mathbf{A})^{-1})$ with $\mathcal{C}(\mathbf{A})$ the conditioning of A . The gradient direction may not be optimal, Conjugate Gradient methods improve the choice of (\mathbf{e}_k) .

0.3.2 Conjugate Gradient

The Conjugate Gradient (CG) is a Krylov-subspace algorithm for symmetric positive definite matrices.

Given $\hat{\mathbf{x}}^0$, $(\hat{\mathbf{x}}^k)$ is a sequence of solutions to approximate k -dimensional minimisation problems.

For the Conjugate Gradient:

1. take the direction \mathbf{e}_{k+1} such that $(\mathbf{e}_1, \dots, \mathbf{e}_k, \mathbf{e}_{k+1})$ is A -conjugate, thus $\hat{\mathbf{x}}^{k+1} = \hat{\mathbf{x}}^k + \alpha_{k+1}\mathbf{e}_{k+1}$.
2. choose the descent factor α^{k+1} minimizing the functional $J(\hat{\mathbf{x}}^k + \alpha_{k+1}\mathbf{r}_k)$, which is defined by

$$\alpha_j = \frac{\mathbf{e}_j^T \mathbf{b}}{\mathbf{e}_j^T \mathbf{A} \mathbf{e}_j}$$

and with $\mathbf{e}_j^T \mathbf{b} \neq 0$ (unless the exact solution is reached).

The construction of $(\mathbf{e}_1, \dots, \mathbf{e}_{k+1})$ is done by orthogonalization of residuals by Gram-Schmidt:

$$\mathbf{e}_{k+1} = \mathbf{r}_k - \frac{\mathbf{e}_k^T \mathbf{A} \mathbf{r}_{k-1}}{\mathbf{e}_k^T \mathbf{A} \mathbf{e}_k}$$

so that $\mathbf{r}_{k+1} = \mathbf{b} - \mathbf{A}\hat{\mathbf{x}}^{k+1} = \mathbf{r}^k - \alpha_{k+1}\mathbf{A}\mathbf{e}_{k+1}$

After N steps, the A -conjugate basis of \mathbb{R}^N is done and the exact solution is reached:

$$\mathbf{x} = \sum_{j=1}^N \alpha_j \hat{\mathbf{x}}^j$$

For any k , the speed of convergence is bounded by

$$\mathcal{O} \left(\frac{1 - \sqrt{\mathcal{C}(\mathbf{A})}}{1 + \sqrt{\mathcal{C}(\mathbf{A})}} \right)^{2k}$$

in the norm induced by A , with $\mathcal{C}(A)$ the conditioning of A .

The Conjugate Gradient can therefore be seen as a direct methods but in practice:

- the iterative computation of the A -conjugate basis suffers from the same issue of numerical error propagation as the QR factorization leading to a loss of orthogonality,
- the convergence is slow, which makes it unrealistic to compute the exact solution for large systems,

so it is used as an iterative method.

Example algorithm on first steps:

1. Given $\hat{\mathbf{x}}^0 = 0$, set $\mathbf{r}_0 = \mathbf{b} - \mathbf{A}\hat{\mathbf{x}}^0$ and $\mathbf{e}_1 = \mathbf{r}_0$,

2. Take $\hat{\mathbf{x}}_1 = \alpha_1 \mathbf{e}_1$, then $\alpha_1 \mathbf{e}_1^T \mathbf{A} \mathbf{e}_1 = \mathbf{e}_1^T \mathbf{b}$, thus

$$\alpha_1 = \frac{\mathbf{r}_0^T \mathbf{b}}{\mathbf{r}_0^T \mathbf{A} \mathbf{r}_0}$$

3. Compute the residual:

$$\mathbf{r}_1 = \mathbf{b} - \mathbf{A}\hat{\mathbf{x}}^1$$

4. Compute the direction:

$$\mathbf{e}_2 = \mathbf{r}_1 - \frac{\mathbf{e}_1^T \mathbf{A} \mathbf{r}_1}{\mathbf{e}_1^T \mathbf{A} \mathbf{e}_1} \mathbf{e}_1$$

5. Compute the factor:

$$\alpha_2 = \frac{\mathbf{e}_2^T \mathbf{b}}{\mathbf{e}_2^T \mathbf{A} \mathbf{e}_2}$$

6. Update the solution:

$$\hat{\mathbf{x}}^2 = \hat{\mathbf{x}}^1 + \alpha_2 \mathbf{e}_2$$

7. ...

The algorithm iteration reads:

1. Compute the residual:

$$\mathbf{r}_k = \mathbf{b} - \mathbf{A}\hat{\mathbf{x}}^k$$

2. Compute the direction:

$$\mathbf{e}_{k+1} = \mathbf{r}_k - \frac{\mathbf{e}_k^T \mathbf{A} \mathbf{r}_k}{\mathbf{e}_k^T \mathbf{A} \mathbf{e}_k} \mathbf{e}_k$$

3. Compute the factor:

$$\alpha_{k+1} = \frac{\mathbf{e}_{k+1}^T \mathbf{b}}{\mathbf{e}_{k+1}^T \mathbf{A} \mathbf{e}_{k+1}}$$

4. Update the solution:

$$\hat{\mathbf{x}}^{k+1} = \hat{\mathbf{x}}^k + \alpha_{k+1} \mathbf{e}_{k+1}$$

which requires two matrix-vector multiplications per loop, $\mathbf{A}\hat{\mathbf{x}}^k$ then $\mathbf{A}\mathbf{e}_{k+1}$. Using $\mathbf{r}_{k+1} = \mathbf{r}_k - \alpha_{k+1} \mathbf{A}\mathbf{e}_{k+1}$ saves one matrix-vector multiplication.

While the residual norm $\varrho_k = \|\mathbf{r}_k\|_2$ is big:

1. Compute the projection:

$$\beta_k = -\frac{\varrho_k}{\varrho_{k-1}}$$

2. Compute the direction:

$$\mathbf{e}_{k+1} = \mathbf{r}_k + \beta_k \mathbf{e}_k$$

3. Compute the factor:

$$\mathbf{w} = \mathbf{A}\mathbf{e}_{k+1}; \alpha_{k+1} = \frac{\varrho_k}{\mathbf{e}_{k+1}^T \mathbf{w}}$$

4. Update the solution:

$$\hat{\mathbf{x}}^{k+1} = \hat{\mathbf{x}}^k + \alpha_{k+1} \mathbf{e}_{k+1}$$

5. Update the residual:

$$\mathbf{r}^{k+1} = \mathbf{r}^k - \alpha_{k+1} \mathbf{w}$$

0.3.3 Preconditioners

While seeing the Conjugate Gradient as a pure iterative method relieves from concerns regarding orthogonality loss, the convergence is still slow as soon as the condition number of the matrix is bad.

Preconditioning the system consists in finding a non-singular symmetric matrix \mathbf{C} such that $\tilde{\mathbf{A}} = \mathbf{C}^{-1}\mathbf{A}\mathbf{C}^{-1}$ and the conjugate gradient is applied to

$$\tilde{\mathbf{A}}\tilde{\mathbf{x}} = \tilde{\mathbf{b}}$$

with $\tilde{\mathbf{x}} = \mathbf{C}^{-1}\mathbf{x}$ and $\tilde{\mathbf{b}} = \mathbf{C}^{-1}\mathbf{b}$.

With:

- $\mathbf{M} = \mathbf{C}^2$
- $\mathbf{e}_k = \mathbf{C}^{-1}\tilde{\mathbf{e}}_k$
- $\hat{\mathbf{x}}_k = \mathbf{C}^{-1}\tilde{\hat{\mathbf{x}}}_k$
- $\mathbf{z}_k = \mathbf{C}^{-1}\tilde{\mathbf{r}}_k$
- $\mathbf{r}_k = \mathbf{C}\tilde{\mathbf{r}}_k = \mathbf{b} - \mathbf{A}\hat{\mathbf{x}}_k$

and \mathbf{M} is a symmetric positive definite matrix called the preconditioner.

While the residual norm $\varrho_k = \|\mathbf{r}_k\|_2$ is big:

1. Solve:

$$\mathbf{M}\mathbf{z}_k = \mathbf{r}_k$$

2. Compute the projection:

$$\beta_k = -\frac{\mathbf{z}_k^T \mathbf{r}_k}{\mathbf{z}_{k-1}^T \mathbf{r}_{k-1}}$$

3. Compute the direction:

$$\mathbf{e}_{k+1} = \mathbf{z}_k + \beta_k \mathbf{e}_k$$

4. Compute the factor:

$$\alpha_{k+1} = \frac{\mathbf{z}_k^T \mathbf{r}_k}{\mathbf{e}_{k+1}^T \mathbf{A} \mathbf{e}_{k+1}}$$

5. Update the solution:

$$\hat{\mathbf{x}}^{k+1} = \hat{\mathbf{x}}^k + \alpha_{k+1} \mathbf{e}_{k+1}$$

6. Update the residual:

$$\mathbf{r}^{k+1} = \mathbf{r}^k - \alpha_{k+1} \mathbf{w}$$

The linear system $\mathbf{M}\mathbf{z}_k = \mathbf{r}_k$ should be easy to solve and can lead to fast convergence, typically $\mathcal{O}(\sqrt{N})$. Since

$$\mathbf{M}\mathbf{z}_k = \mathbf{b} - \mathbf{A}\hat{\mathbf{x}}_k$$

Then an iterative relation appears:

$$\hat{\mathbf{x}}_{k+1} = \mathbf{M}^{-1}(\mathbf{b} - \mathbf{A}\hat{\mathbf{x}}_k)$$

therefore iterative methods like Jacobi, Gauss-Seidel and relaxation methods can be used.

0.4 Power method

This method is used for finding the dominant eigenvalues of a matrix $\mathbf{A} \in M_N(\mathbb{R})$ of N eigenvectors (\mathbf{v}_i) with associated eigenvalues (λ_i) ordered in decreasing module. The eigenvalues are either real or conjugate complex pairs.

Given a random vector \mathbf{x}^0 , construct a sequence of vectors $(\hat{\mathbf{x}}^k)$ such that

$$\hat{\mathbf{x}}^{k+1} = \mathbf{A}\hat{\mathbf{x}}^k$$

then $\forall k \geq 0$

$$\hat{\mathbf{x}}^k = \sum_{i=0}^{N-1} \lambda_i^k \xi_i \mathbf{v}_i$$

for some coefficients (ξ_i) .

Assume that λ_0 is a dominant real eigenvalue and $\xi_0 \neq 0$, then

$$\hat{\mathbf{x}}^k = \lambda_0^k (\xi_0 \mathbf{v}_0 + \mathbf{r}_k)$$

with the residual \mathbf{r}_k defined as

$$\mathbf{r}_k = \lambda_0^{-k} \sum_{i=1}^{N-1} \lambda_i^k \xi_i \mathbf{v}_i$$

and $\lim_{k \rightarrow \infty} \mathbf{r}_k = O_{\mathbb{R}^N}$. To the limit $\hat{\mathbf{x}}_{k+1} \approx \lambda_0 \hat{\mathbf{x}}^k \approx \lambda_0 \xi_0 \mathbf{v}_0$ almost parallel to the first eigenvector.

- This method is fast to compute the spectral radius for the Jacobi method and relaxation parameters.
- The convergence is geometric and the speed depends on the ratio $|\lambda_1/\lambda_0|$.
- If the matrix is symmetric, the convergence speed can be doubled.
- If λ_0 is very large or very small then taking high powers lead to numerical issues, the algorithm requires a normalization.