Utrecht, 26 oktober 2016

Basic Iterative Methods II

Gerard Sleijpen

Universiteit Utrecht
Department of Mathematics

http://www.staff.science.uu.nl/~sleij101/

Program Lecture 5

- Power Method & Richardson
- Filtering
- Shift-and-Invert & Preconditioning
- Polynomial Iteration
- Selecting Parameters

  1) single parameter a) static
  b) dynamic

  2) Multiple parameters a) static (Chebyshev)
  b) dynamic (GCR)

$\mathbf{A}$ is $n \times n$, $\mathbf{A}\mathbf{v}_j = \lambda_j \mathbf{v}_j$

$\mathbf{A}\mathbf{v} = \lambda \mathbf{v}$ **Shifted power**: $\widetilde{\mathbf{u}}_k = (\mathbf{A} - \sigma \mathbf{I})\mathbf{u}_k$

Scale $\mathbf{u}_{k+1} = \widetilde{\mathbf{u}}_k / \|\widetilde{\mathbf{u}}_k\|_2$

Theorem.
The $\mathbf{u}_k$ converge to (a multiple of) $\mathbf{v}_{j_0}$ if

$|\lambda_{j_0} - \sigma| > |\lambda_j - \sigma|$ all $j \neq j_0$:

and $\mathbf{u}_0$ has a component in the direction of $\mathbf{v}_{j_0}$

$\mathbf{v}_{j_0}$ is the **dominant** eigenvector of $\mathbf{A} - \sigma \mathbf{I}$,
and $\lambda_{j_0} - \sigma$ is the **dominant** eigenvalue.

Eventual error reduction is $\rho \equiv \max_{j \neq j_0} \frac{|\lambda_j - \sigma|}{|\lambda_{j_0} - \sigma|}$

$\mathbf{A}$ is $n \times n$, $\mathbf{A}\mathbf{v}_j = \lambda_j \mathbf{v}_j$

$\mathbf{A}\mathbf{v} = \lambda \mathbf{v}$ **Shifted power**: $\widetilde{\mathbf{u}}_k = (\mathbf{I} - \alpha \mathbf{A})\mathbf{u}_k$

Scale $\mathbf{u}_{k+1} = \widetilde{\mathbf{u}}_k / \mathbf{e}_1^* \widetilde{\mathbf{u}}_k$

Improvements are based on the fact that

$$f(\mathbf{A})\mathbf{v}_j = f(\lambda_j)\mathbf{v}_j.$$

Examples. $f(\mathbf{A}) = (\mathbf{I} - \alpha \mathbf{A})$

$f(\mathbf{A}) = \mathbf{I} + \gamma_1 \mathbf{A} + \ldots + \gamma_\ell \mathbf{A}^\ell = (\mathbf{I} - \alpha_1 \mathbf{A}) \ldots (\mathbf{I} - \alpha_\ell \mathbf{A})$

$f(\mathbf{A}) = (\mathbf{A} - \sigma \mathbf{I})^{-1}$

Combination. Cayley transform:

$$f(\mathbf{A}) = (\mathbf{A} - \mathbf{I})^{-1}(\mathbf{I} + \mathbf{A})$$

$\mathbf{A}\mathbf{v} = \lambda\mathbf{v}$ **Shifted power**: $\tilde{\mathbf{u}}_k = (\mathbf{I} - \alpha\mathbf{A})\mathbf{u}_k$

Scale $\mathbf{u}_{k+1} = \tilde{\mathbf{u}}_k / \mathbf{e}_1^*\tilde{\mathbf{u}}_k$

Improvements are based on the fact that

$$f(\mathbf{A})\mathbf{v}_j = f(\lambda_j)\mathbf{v}_j.$$

Examples. $f(\mathbf{A}) = (\mathbf{I} - \alpha\mathbf{A})$

$f(\mathbf{A}) = \mathbf{I} + \gamma_1\mathbf{A} + \ldots + \gamma_\ell\mathbf{A}^\ell = (\mathbf{I} - \alpha_1\mathbf{A})\ldots(\mathbf{I} - \alpha_\ell\mathbf{A})$

$f(\mathbf{A}) = (\mathbf{A} - \sigma\mathbf{I})^{-1}$

Combination. Cayley transform:

$$f(\mathbf{A}) = (\mathbf{A} - \sigma\mathbf{I})^{-1}(\mathbf{I} - \alpha\mathbf{A})$$

$\mathbf{A}\mathbf{v} = \lambda\mathbf{v}$ **Shifted power**: $\tilde{\mathbf{u}}_k = (\mathbf{I} - \alpha\mathbf{A})\mathbf{u}_k$

Scale $\mathbf{u}_{k+1} = \tilde{\mathbf{u}}_k / \mathbf{e}_1^*\tilde{\mathbf{u}}_k$

Improvements. Apply power method with

$f(\mathbf{A}) = (\mathbf{I} - \alpha_1\mathbf{A})\ldots(\mathbf{I} - \alpha_\ell\mathbf{A})$ or $f(\mathbf{A}) = (\mathbf{A} - \sigma\mathbf{I})^{-1}(\mathbf{I} - \alpha\mathbf{A})$

Equivalent interpretations.

1. Diminish unwanted components. **Filtering.**
2. Amplify wanted components
3. Improve distribution eigenvalues. **Preconditioning.**

Preconditioning

Purpose. To improve the distribution of the eigenvalues in order to speed up convergence.

For eigenvalue computation:
make the wanted eigenvector (strongly) dominant.
Shift & Invert can be a feasible strategy

For linear systems: cluster the eigenvalues round 1.
Precondition with a matrix **M** for which

- $\Lambda(\mathbf{M}^{-1}\mathbf{A})$ clusters 'better' round 1 than $\Lambda(\mathbf{A})$
- the system $\mathbf{M}\mathbf{u} = \mathbf{r}$ can efficiently be solved for $\mathbf{u}$.

For eigenvalue computation:
A and $\mathbf{M}^{-1}\mathbf{A}$ generally do not have the same eigenvectors.

$\mathbf{A}\mathbf{v} = \lambda\mathbf{v}$ **Shifted power**: $\tilde{\mathbf{u}}_k = (\mathbf{I} - \alpha\mathbf{A})\mathbf{u}_k$

Scale $\mathbf{u}_{k+1} = \tilde{\mathbf{u}}_k / \mathbf{e}_1^*\tilde{\mathbf{u}}_k$

Improvements. Apply power method with

$f(\mathbf{A}) = (\mathbf{I} - \alpha_1\mathbf{A})\ldots(\mathbf{I} - \alpha_\ell\mathbf{A})$ or $f(\mathbf{A}) = (\mathbf{A} - \sigma\mathbf{I})^{-1}(\mathbf{I} - \alpha\mathbf{A})$

$\mathbf{A}\mathbf{x} = \mathbf{b}$ **Richardson:** $\mathbf{x}_{k+1} = \mathbf{x}_k + \alpha(\mathbf{b} - \mathbf{A}\mathbf{x}_k)$

Polynomial version: Select α_k per step.

Purpose: Diminish all components 'equally' well.

Richardson (with relax. par.)

```
Select x_0, α, tol, k_max
Compute r_0 = b − A x_0
for k = 0, 1, 2, ..., k_max do
    If ||r|| ≤ tol, break, end if
    u_k = r_k
    c_k = A u_k
    x_{k+1} = x_k + α u_k
    r_{k+1} = r_k − α c_k
end do
```

$\mathbf{u}_k$ **search direction** (for the approximate)

Note. Update $\mathbf{r}_k$ of the form $\mathbf{A}\mathbf{u}_k$ with $\mathbf{u}_k$ update $\mathbf{x}_k$.

Richardson (with relax. par.)

```
Select x, α, tol, k_max
Compute r = b − A x
for k = 0, 1, 2, ..., k_max do
    If ||r|| ≤ tol, break, end if
    u = r
    c = A u
    x ← x + α u
    r ← r − α c
end do
```

This is a 'memory friendly' version.

$\leftarrow$: new value replaces old one.

Polynomial iteration

```
Select x, α_1, ..., α_ℓ, tol, k_max
Compute r = b − A x
for k = 0, 1, 2, ..., k_max do
    If ||r|| ≤ tol, break, end if
    u = r
    c = A u
    j = k mod ℓ, α = α_{j+1}
    x ← x + α u
    r ← r − α c
end do
```

General remarks for linear systems.

- **The preconditioned system.**
For ease of discussion assume no preconditioning:
if preconditioner replace $\mathbf{A}$ by $\mathbf{M}^{-1}\mathbf{A}$ and $\mathbf{b}$ by $\mathbf{M}^{-1}\mathbf{b}$.

- **Consistent updates.**
We update $\mathbf{r}$ and $\mathbf{x}$ consistently:
update $\mathbf{r}$ by vectors $-\mathbf{c}$ of the form $\mathbf{c} = \mathbf{A}\mathbf{u}$ with $\mathbf{u}$ explicitly avaliable and update $\mathbf{x}$ by $\mathbf{u}$
$\mathbf{x}_{k+1} = \mathbf{x}_k + \alpha_k \mathbf{u}_k, \quad \mathbf{c}_k = \mathbf{A}\mathbf{u}_k, \quad \mathbf{r}_{k+1} = \mathbf{r}_k - \alpha_k \mathbf{c}_k$

- **The shifted system.**
Assume $\mathbf{x}_0 = \mathbf{0}$.
If $\mathbf{x}_0 \neq \mathbf{0}$, solve $\mathbf{A}\mathbf{x} = \mathbf{r}_0 \equiv \mathbf{b} - \mathbf{A}\mathbf{x}_0$.

How to select the α_j and σ?

Static.
Select parameter(s) before starting the iteration.
Base selection on pre-knowledge of the spectrum.

Dynamic.
Let the computational process determine the parameter(s).
Computation based on information that becomes available during the iteration.

Static. Single parameter

Examples. $\mathbf{A}\mathbf{v}_0 = \lambda_0\mathbf{v}_0,\ \ \lambda_0 \in \Lambda(\mathbf{A})$ wanted eigenvalue.

- If $|\lambda_0-\mu| > |\lambda-\mu|$ for all other $\lambda \in \Lambda(\mathbf{A})$:

$$f(\mathbf{A}) = \mathbf{A}-\mu\mathbf{I}.$$

Shifted power method.

- If λ_0 closest to some target value τ is wanted:

$$f(\mathbf{A}) = (\mathbf{A}-\sigma\mathbf{I})^{-1} \text{ with } \sigma=\tau.$$

Inverse iteration or **Wielandt iteration**.

Static. Single parameter

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

- If all λ_j eigenvalues $\mathbf{A}$ in $[\lambda_-,\lambda_+] = [\mu-\rho,\mu+\rho] \subset (0,\infty)$:

$\mu = (\lambda_+ + \lambda_-)/2, \qquad \rho = (\lambda_+ - \lambda_-)/2.$

$f(\mathbf{A}) = \mathbf{I}-\alpha_{\text{opt}}\mathbf{A}$ with $\alpha_{\text{opt}} \equiv 1/\mu$,

$$\max|f(\lambda_j)| \leq \frac{\lambda_+-\lambda_-}{\lambda_++\lambda_-} = \frac{1-\frac{1}{\mathcal{C}}}{1+\frac{1}{\mathcal{C}}} \leq e^{-\frac{2}{\mathcal{C}}}, \quad \text{where} \quad \mathcal{C} \equiv \frac{\lambda_+}{\lambda_-}$$

Therefore, for Richardson with $\alpha = \alpha_{\text{opt}}$,

$$\|\mathbf{r}_{k+1}^{\text{Rich}}\| \lesssim \exp\left(-\frac{2}{\mathcal{C}}\right)\|\mathbf{r}_k^{\text{Rich}}\| \qquad k \text{ large.}$$

Dynamic. Single parameter

Examples. $\mathbf{A}\mathbf{v}_0 = \lambda_0\mathbf{v}_0,\ \ \lambda_0 \in \Lambda(\mathbf{A})$ wanted eigenvalue.

- $f(\mathbf{A}) = (\mathbf{A}-\sigma\mathbf{I})^{-1}$, with $\sigma = \sigma_k = \rho(\mathbf{u}_k) \equiv \frac{\mathbf{u}_k^*\mathbf{A}\mathbf{u}_k}{\mathbf{u}_k^*\mathbf{u}_k}$.

Rayleigh Quotient Iteration

The Rayleigh quotient $\rho(\mathbf{u}_k)$ is the 'best' available approximate eigenvalue at step k.

If RQI converges, it converges quadratically eventually.
For Hermitian $\mathbf{A}$, the asymptotic convergence is even cubic.

"If converges": **Example.** $\mathbf{A} = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$. $\mathbf{v}_0 = e_1$.

RQI:

+ Fast convergence (if convergence).
+ Can detect eigenvalues in the interior of the spectrum.
− No controle on what eigenvalue is going to be detected.
  Remedy: First a few steps of Wielandt iteration.
− The linear systems to be solved require
  a new LU-decompostion in each step.

Wielandt Iteration:

− Linear convergence.
+ Can detect eigenvalues in the interior of the spectrum.
+ Finds eigenvalue close to the shift.
+ The same LU-decomposition can used in each step.

Note. The fact that linear systems have to be solved may make the methods not feasible for huge n.

$f(\mathbf{A}) = (\mathbf{I}-\alpha_1\mathbf{A})\ldots(\mathbf{I}-\alpha_\ell\mathbf{A})$ or $f(\mathbf{A}) = (\mathbf{A}-\sigma\mathbf{I})^{-1}(\mathbf{I}-\alpha\mathbf{A})$

Dynamic. Single parameter

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

Select $f(\mathbf{A}) = \mathbf{I}-\alpha_k\,\mathbf{A}$ with α_k to minimize:

- **Minimal Residual**: $\|\mathbf{r}_{k+1}\|_2 = \|\mathbf{r}_k - \alpha_k\,\mathbf{c}_k\|_2$ minimal
- If $\mathbf{A}$ is positive definite
  Steepest descent: $\|\mathbf{x}-\mathbf{x}_{k+1}\|_A$ minimal

Convergence if $\mathrm{Re}(\lambda_j) > 0$ for all eigenvalues λ_j of $\mathbf{A}$.

$f(\mathbf{A}) = (\mathbf{I}-\alpha_1\mathbf{A})\ldots(\mathbf{I}-\alpha_\ell\mathbf{A})$ or $f(\mathbf{A}) = (\mathbf{A}-\sigma\mathbf{I})^{-1}(\mathbf{I}-\alpha\mathbf{A})$

Static. Multiple parameter

Examples. $\mathbf{A}\mathbf{x}=\mathbf{b}$.
Suppose we have a set $\mathcal{E}\subset\mathbb{C}$ that contains all λ_i.

Select $f(\mathbf{A}) = (\mathbf{I}-\alpha_1\mathbf{A})\cdot\ldots\cdot(\mathbf{I}-\alpha_\ell\mathbf{A})$, i.e., α_j, such that

$$\nu \equiv \max\{|f(\zeta)| = |(1-\alpha_1\zeta)\cdot\ldots\cdot(1-\alpha_\ell\zeta)| \mid \zeta\in\mathcal{E}\}$$

is as small as possible.

Notation.
$\mathcal{P}_\ell$ is the set of all polynomials of degree at most ℓ.

$$\mathcal{P}_\ell^0 \equiv \{p\in\mathcal{P}_\ell \mid p(0)=1\}$$

Observation. $p\in\mathcal{P}_\ell$

$$p(0)=1 \quad\Leftrightarrow\quad p(\zeta) = (1-\alpha_1\zeta)\cdot\ldots\cdot(1-\alpha_\ell\zeta).$$

$f(\mathbf{A}) = (\mathbf{I}-\alpha_1\mathbf{A})\ldots(\mathbf{I}-\alpha_\ell\mathbf{A})$ or $f(\mathbf{A}) = (\mathbf{A}-\sigma\mathbf{I})^{-1}(\mathbf{I}-\alpha\mathbf{A})$

Static. Multiple parameter

Examples. $\mathbf{A}\mathbf{x}=\mathbf{b}$.
Suppose we have a set $\mathcal{E}\subset\mathbb{C}$ that contains all λ_i.

Select $f(\mathbf{A}) = (\mathbf{I}-\alpha_1\mathbf{A})\cdot\ldots\cdot(\mathbf{I}-\alpha_\ell\mathbf{A})$, i.e., α_j, such that

$$\nu \equiv \max\{|f(\zeta)| = |(1-\alpha_1\zeta)\cdot\ldots\cdot(1-\alpha_\ell\zeta)| \mid \zeta\in\mathcal{E}\}$$

is as small as possible.

This is a problem from approximation theory:

Find a polynomial in $\mathcal{P}_\ell^0$ that is as small as possible on $\mathcal{E}$.

Solutions for $\mathcal{E} = [\lambda_-,\lambda_+]\subset(0,\infty)$ (**Chebyshev pols**)

Approximate solutions for ellipses (Cheb.), polygons (Faber pols).

$f(\mathbf{A}) = (\mathbf{I} - \alpha_1 \mathbf{A}) \ldots (\mathbf{I} - \alpha_\ell \mathbf{A})$ or $f(\mathbf{A}) = (\mathbf{A} - \sigma \mathbf{I})^{-1}(\mathbf{I} - \alpha \mathbf{A})$

Static. Multiple parameter

Examples. $\mathbf{A}\mathbf{v} = \lambda \mathbf{v}$.
Suppose we have a set $\mathcal{E} \subset \mathbb{C}$ that contains all λ_i, except for the wanted eigenvalue $\lambda_0 \in \Lambda(\mathbf{A})$.

Select $f(\mathbf{A}) = (\mathbf{I} - \alpha_1 \mathbf{A}) \cdot \ldots \cdot (\mathbf{I} - \alpha_\ell \mathbf{A})$ such that with

$$\nu \equiv \max_{\zeta \in \mathcal{E}} |(1 - \alpha_1 \zeta) \cdot \ldots \cdot (1 - \alpha_\ell \zeta)|$$

$\nu / |f(\lambda_0)|$ is as small as possible.

Chebyshev polynomials

$$T_\ell(x) \equiv \frac{1}{2}(\zeta^\ell + \zeta^{-\ell}), \quad \text{where} \quad x \equiv \frac{1}{2}(\zeta + \zeta^{-1}) \quad (\zeta \in \mathbb{C}).$$

Exercise. For all $x \in \mathbb{C}$ we have

$$\begin{cases} T_0(x) = 1, \quad T_1(x) = x, \\ T_{k+1}(x) = 2\,x\,T_k(x) - T_{k-1}(x) \ \text{ for } \ k = 1, 2, \ldots. \end{cases}$$

Assume $[\lambda_-, \lambda_+] = [\mu - \rho, \mu + \rho] \subset (0, \infty)$.

Theorem. With $x \equiv (\mu - \lambda)/\rho$ and $p_{\text{Cheb}}(\lambda) \equiv \dfrac{T_\ell(x)}{T_\ell(\mu/\rho)}$,

we have that $p_{\text{Cheb}} \in \mathcal{P}_\ell^0$ and for any $q \in \mathcal{P}_\ell^0$,

$$\max |p_{\text{Cheb}}(\lambda)| \leq \max |q(\lambda)|,$$

where the maxima are taken over all $\lambda \in [\lambda_-, \lambda_+]$.

Chebyshev polynomials

$$T_\ell(x) \equiv \frac{1}{2}(\zeta^\ell + \zeta^{-\ell}), \quad \text{where} \quad x \equiv \frac{1}{2}(\zeta + \zeta^{-1}) \quad (\zeta \in \mathbb{C}).$$

Exercise. For all $x \in \mathbb{C}$ we have

$$\begin{cases} T_0(x) = 1, \quad T_1(x) = x, \\ T_{k+1}(x) = 2\,x\,T_k(x) - T_{k-1}(x) \ \text{ for } \ k = 1, 2, \ldots. \end{cases}$$

Assume $[\lambda_-, \lambda_+] = [\mu - \rho, \mu + \rho] \subset (0, \infty)$.

Theorem. With $x \equiv (\mu - \lambda)/\rho$ and $p_{\text{Cheb}}(\lambda) \equiv \dfrac{T_\ell(x)}{T_\ell(\mu/\rho)}$,

we have that $p_{\text{Cheb}} \in \mathcal{P}_\ell^0$ and

$$\max |p_{\text{Cheb}}(\lambda)| = \frac{1}{|T_\ell(\mu/\rho)|} \leq 2 \exp\left(-\frac{2\ell}{\sqrt{\mathcal{C}}}\right),$$

where the max. is taken over all $\lambda \in [\lambda_-, \lambda_+]$ and $\mathcal{C} \equiv \frac{\lambda_+}{\lambda_-}$.

Chebyshev versus Richardson

Error reduction for spectrum in $[\lambda_-, \lambda_+] \subset (0, \infty)$.
Put $\mathcal{C} \equiv \frac{\lambda_+}{\lambda_-}$.

- Degree ℓ Chebychev.

$$\|\mathbf{r}_{k+\ell}^{\text{Cheb}(\ell)}\|_2 \lesssim 2 \exp\left(-\frac{2\ell}{\sqrt{\mathcal{C}}}\right) \|\mathbf{r}_k^{\text{Cheb}(\ell)}\|_2 \qquad k \text{ large}$$

- Richardson with optimal α.

$$\|\mathbf{r}_{k+\ell}^{\text{Rich}}\|_2 \lesssim \exp\left(-\frac{2\ell}{\mathcal{C}}\right) \|\mathbf{r}_k^{\text{Rich}}\|_2 \qquad k \text{ large}$$

Note. Chebyshev iteration is designed for spectra in intervals, but works well also for (narrow) ellipses around an interval.

Chebyshev

With $\mu \equiv \frac{\lambda_+ + \lambda_-}{2}$ and $\rho \equiv \frac{\lambda_+ - \lambda_-}{2}$ we have that

$$\mathbf{r}_k = \frac{\widetilde{\mathbf{r}}_k}{\gamma_k} \quad \text{with} \quad \widetilde{\mathbf{r}}_k \equiv T_k(\tfrac{1}{\rho}(\mu\mathbf{I} - \mathbf{A}))\mathbf{r}_0, \quad \gamma_k \equiv T_k(\tfrac{\mu}{\rho})$$

$T_{k+1}(x) = 2xT_k(x) - T_{k-1}(x)$ implies that

$$\gamma_{k+1} = 2\tfrac{\mu}{\rho}\gamma_k - \gamma_{k-1} \quad \text{and} \quad \widetilde{\mathbf{r}}_{k+1} = \tfrac{2\mu}{\rho}\widetilde{\mathbf{r}}_k - \tfrac{2}{\rho}\mathbf{A}\widetilde{\mathbf{r}}_k - \widetilde{\mathbf{r}}_{k-1}.$$

Hence,

$$\mathbf{r}_{k+1} = \frac{2\mu\gamma_k}{\rho\gamma_{k+1}}\mathbf{r}_k - \frac{2\gamma_k}{\rho\,\gamma_{k+1}}\mathbf{A}\mathbf{r}_k - \frac{\gamma_{k-1}}{\gamma_{k+1}}\mathbf{r}_{k-1}$$

$$\mathbf{x}_{k+1} = \frac{2\mu\gamma_k}{\nu\gamma_{k+1}}\mathbf{x}_k + \frac{2\gamma_k}{\rho\gamma_{k+1}}\mathbf{r}_k - \frac{\gamma_{k-1}}{\gamma_{k+1}}\mathbf{x}_{k-1}$$

Note that the update of the residual also uses an additional 'older' residual.

Degree ℓ Chebyshev versus Chebyshev

Error reduction for spectrum in $[\lambda_-, \lambda_+] \subset (0, \infty)$.
Put $\mathcal{C} \equiv \frac{\lambda_+}{\lambda_-}$.

- Degree ℓ Chebychev.

$$\|\mathbf{r}_{j\ell}^{\text{Cheb}(\ell)}\|_2 \leq \mathcal{C}_E\, 2^j \exp\left(-\frac{2j\ell}{\sqrt{\mathcal{C}}}\right)\|\mathbf{r}_0\|_2 \qquad k \text{ large}$$

- Chebyshev

$$\|\mathbf{r}_{j\ell}^{\text{Cheb}}\|_2 \leq \mathcal{C}_E\, 2 \exp\left(-\frac{2j\ell}{\sqrt{\mathcal{C}}}\right)\|\mathbf{r}_0\|_2 \qquad k \text{ large}$$

Here, $\mathcal{C}_E$ some constant like $\mathcal{C}_E = \|\mathbf{V}\|_2\|\mathbf{V}^{-1}\|\|_2$, the conditioning of the basis of eigenvectors.

Chebyshev

```
Select x_0, tol, kmax, μ, ρ
Compute r_0 = b − A x_0
Set ν_0 = μ, r_1 = r_0 − (1/μ) A r_0, x_1 = x_0 + (1/μ) r_0
for k = 1,...,kmax do
    If ||r|| ≤ tol, break, end if
    ν_k = 2μ − ρ^2/ν_{k−1}
    α_k = 2μ/ν_k, β_k = 2/ν_k, γ_k = ρ^2/(ν_{k−1} ν_k),
    r_{k+1} = α_k r_k − β_k A r_k − γ_k r_{k−1}
    x_{k+1} = α_k x_k + β_k r_k − γ_k x_{k−1}
end for
```

With $\mu, \rho \in \mathbb{R}, \rho > 0$ such that

$$\Lambda(\mathbf{A}) \subset [\mu - \rho, \mu + \rho] \subset (0, \infty).$$

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

Summary.

- $\mathbf{r}_k$ is of the form $p_k(\mathbf{A})\mathbf{r}_0$ with $p_k \in \mathcal{P}_k^0$.

Examples. $p_k(x) = (1 - \alpha x)^k$ Richardson,
$p_{m\ell}(x) = (\prod_{j=1}^{\ell}(1 - \alpha_j x))^m$ Polynomial,
$p_k(x) = T_k(\frac{\mu - x}{\rho})/T_k(\frac{\mu}{\rho})$ Chebyshev, ...

- Since $p_k(0) = 1$ we have that $p_k(x) = 1 - xq_{k-1}(x)$ for some polynomial q_{k-1} of degree $k - 1$ and

$$\mathbf{r}_k = \mathbf{r}_0 - \mathbf{A}q_{k-1}(\mathbf{A})\mathbf{r}_0, \qquad \mathbf{x}_k = \mathbf{x}_0 + q_{k-1}(\mathbf{A})\mathbf{r}_0.$$

- Consistent update of $\mathbf{r}_k$ and $\mathbf{x}_k$,

$\mathbf{x}_{k+1} = \mathbf{x}_k + \alpha_k\mathbf{u}_k$, $\mathbf{c}_k = \mathbf{A}\mathbf{u}_k$, $\mathbf{r}_{k+1} = \mathbf{r}_k - \alpha_k\mathbf{c}_k$

i.e., no need to gather explicit information on q_{k-1}.

$\mathbf{Ax} = \mathbf{b}$

Summary.
Let $\mathcal{K}_k(\mathbf{A}, \mathbf{r}_0)$ be the **Krylov subspace of order** k generated by $\mathbf{A}$ and $\mathbf{r}_0$:

$$\begin{aligned} \mathcal{K}_k(\mathbf{A}, \mathbf{r}_0) &\equiv \text{span}(\mathbf{r}_0, \mathbf{A}\mathbf{r}_0, \ldots, \mathbf{A}^{k-1}\mathbf{r}_0) \\ &= \{q(\mathbf{A})\mathbf{r}_0 \mid q \in \mathcal{P}_{k-1}\}. \end{aligned}$$

Then

$$\mathbf{r}_k \in \mathbf{r}_0 + \mathbf{A}\mathcal{K}_k(\mathbf{A}, \mathbf{r}_0) \subset \mathcal{K}_{k+1}(\mathbf{A}, \mathbf{r}_0),$$
$$\mathbf{x}_k \in \mathcal{K}_k(\mathbf{A}, \mathbf{r}_0).$$

Dynamic. Multiple parameter

Find the residual in the Krylov subspace $\mathcal{K}_{k+1}(\mathbf{A}, \mathbf{r}_0)$ with 'smallest' norm. Use also 'older' residuals in the update process.

Generalized Conjugate Residuals

```
Select x_0, k_max, tol
Compute r_0 = b - A x_0
for k = 0, 1, ..., k_max do
   break if ||r_k||_2 <= tol
   u_k = r_k, c_k = A u_k
   for j = 0, ..., k-1 do
      β_j = c_j^* c_k / σ_j
      u_k ← u_k - β_j u_j
      c_k ← c_k - β_j c_j
   end for
   σ_k = c_k^* c_k,   α_k = c_k^* r_k / σ_k
   x_{k+1} = x_k + α_k u_k
   r_{k+1} = r_k - α_k c_k
end for
```

Generalized Conjugate Residuals

GCR is an optimal Krylov subspace solver:

Theorem. Assume $\mathbf{x}_0 = \mathbf{0}$: $\mathbf{r}_0 = \mathbf{b}$.
The GCR approximate solution $\mathbf{x}_k$ at step k is the vector in $\mathcal{K}_k(\mathbf{A}, \mathbf{r}_0)$ with smallest residual norm:

$$\|\mathbf{r}_k\|_2 = \|\mathbf{r}_0 - \mathbf{A}\mathbf{x}_k\|_2 \le \|\mathbf{r}_0 - \mathbf{A}\tilde{\mathbf{x}}\|_2 \quad (\tilde{\mathbf{x}} \in \mathcal{K}_k(\mathbf{A}, \mathbf{r}_0)).$$

In particular, $\|\mathbf{r}_k^{\text{GCR}}\|_2 \le \|\mathbf{r}_k^{\text{Cheb}}\|_2$.

Hence, if $\Lambda(\mathbf{A}) \subset [\lambda_-, \lambda_+] \subset (0, \infty)$, then, with $\mathcal{C} \equiv \frac{\lambda_+}{\lambda_-}$,

$$\|\mathbf{r}_k^{\text{GCR}}\|_2 \le \mathcal{C}_E\, 2 \exp\left(-\frac{2k}{\sqrt{\mathcal{C}}}\right) \|\mathbf{r}_0\|_2.$$

Here, $\mathcal{C}_E$ some constant like $\mathcal{C}_E = \|\mathbf{V}\|_2 \|\mathbf{V}^{-1}\|_2$,

the conditioning of the basis of eigenvectors.

Chebyshev versus GCR

Chebyshev.
+ No inner products
+ Short recurrences (three term recurrences)
− Not the smallest residuals with appr. sol. from $\mathcal{K}_k(\mathbf{A}, \mathbf{r}_0)$.
− Sensitive to the estimate of the hull of the spectrum.
− Only effective if spectrum in a narrow ellipse in a half plane as $\mathbb{C}^+$.

GCR.
+ Smallest residual with appr. sol. from $\mathcal{K}_k(\mathbf{A}, \mathbf{r}_0)$.
+ Flexible (any information can be used for $\mathbf{u}_k$)
+ Stable
− **Growing recurrences with increasing step number k: increasing computational costs, increasing storage demands.**

Flexible GCR

In the preceding transparancies, GCR has been constructed as an **optimal Krylov subspace solver**.

However, GCR can be turned into a **supspace solver!**:

If

$\mathbf{u}_k = \mathbf{r}_k$

is replaced by

`Solve approximately` $\mathbf{A}\mathbf{u}_k = \mathbf{r}_k$ `for` $\mathbf{u}_k$

then we search for an approximate solution in the search subspace span$(\mathbf{u}_0, \ldots, \mathbf{u}_{k-1})$ and GCR finds the one with smallest residual.

Exercise. Exact solve of $\mathbf{A}\mathbf{u}_k = \mathbf{r}_k$ leads to $\mathbf{x}_{k+1} = \mathbf{x}$.

Flexible GCR

`Solve approximately` $\mathbf{A}\mathbf{u}_k = \mathbf{r}_k$ `for` $\mathbf{u}_k$

Examples.

- $\mathbf{u}_k = \mathbf{r}_k$: standard GCR — searches $\mathcal{K}_k(\mathbf{A}, \mathbf{r}_0)$
- `Solve` $\mathbf{M}\mathbf{u}_k = \mathbf{r}_k$ `for` $\mathbf{u}_k$: preconditioned GCR
  searches the Krylov subspace $\mathbf{M}^{-1}\mathcal{K}_k(\mathbf{A}\mathbf{M}^{-1}, \mathbf{r}_0)$.
- `Use` ℓ `steps of GCR to solve` $\mathbf{A}\mathbf{u}_k = \mathbf{r}_k$: nested GCR
  solution in $\mathcal{K}_{\ell k}(\mathbf{A}, \mathbf{r}_0)$
- `Use GCR to solve` $\mathbf{A}\mathbf{u}_k = \mathbf{r}_k$ `to rel. res. acc.` 0.1
- At step $k = 0, 1, \ldots, \ell$ use information on the solution (as $\mathbf{u}_k$ representing singularities, etc.)
- At step $k = 0, \ldots, \ell$ use a '$\mathbf{u}_j$' from GCR run for $\mathbf{A}\mathbf{x} = \widetilde{\mathbf{b}}$.

⋮

GCR and Krylov subspace solvers

GCR is a subspace solver

Pros
- Flexible (any information can be exploited)

Cons
- Higher computational costs per step

Krylov subspace solvers

Pros
- Krylov subspace structure can be exploited to save computational costs per step [to be implemented ∗)].
- Polynomial approximation theory provides insight in convergence behaviour

Cons
- Sensitive to rounding errors [if ∗)].
- Not flexible (only fixed preconditioners are allowed).

$\mathbf{A}^* = \mathbf{A}$ Conjugate Residuals

```
Select x_0, k_max, tol
Compute r_0 = b − A x_0
for k = 0, 1, ..., k_max do
    break if ||r_k||_2 ≤ tol
    u_k = r_k,   c_k = A u_k
    β_{k−1} = c*_{k−1} c_k / σ_{k−1}
    u_k ← u_k − β_{k−1} u_{k−1}
    c_k ← c_k − β_{k−1} c_{k−1}
    σ_k = c*_k c_k,   α_k = c*_k r_k / σ_k
    x_{k+1} = x_k + α_k u_k
    r_{k+1} = r_k − α_k c_k
end for
```

$\mathbf{A}^* = \mathbf{A}$ Conjugate Residuals

3 DOTs:

$\beta_{k-1} = \frac{\mathbf{c}_{k-1}^* \mathbf{A} \mathbf{r}_k}{\sigma_{k-1}}, \quad \sigma_k = \mathbf{c}_k^* \mathbf{c}_k, \quad \rho_k \equiv \mathbf{c}_k^* \mathbf{r}_k, \quad \alpha_k = \frac{\rho_k}{\sigma_k}$

Save 1 DOT:

$$\beta_{k-1} = \frac{\mathbf{c}_{k-1}^* \mathbf{A} \mathbf{r}_k}{\mathbf{c}_{k-1}^* \mathbf{c}_{k-1}} = \frac{[\mathbf{r}_k - \mathbf{r}_{k-1}]^* \mathbf{A} \mathbf{r}_k}{[\mathbf{r}_k - \mathbf{r}_{k-1}]^* \mathbf{c}_{k-1}} = -\frac{\mathbf{r}_k^* \mathbf{A} \mathbf{r}_k}{\mathbf{r}_{k-1}^* \mathbf{c}_{k-1}} = -\frac{\rho_k}{\rho_{k-1}}$$

Here we used that $\alpha_{k-1} \mathbf{c}_{k-1} = \mathbf{r}_k - \mathbf{r}_{k-1}$

$\mathbf{r}_k \perp \mathbf{c}_{k-1}, \quad \mathbf{r}_k \perp \mathbf{A} \mathbf{r}_{k-1}$

$\mathbf{c}_k = \mathbf{A} \mathbf{r}_k - \beta_{k-1} \mathbf{c}_{k-1}$

$\mathbf{c}_k \perp \mathbf{c}_{k-1}$

Exercise. $\sigma_k = \mathbf{c}_k^* \mathbf{A} \mathbf{r}_k, \quad \rho_k = \mathbf{r}_k^* \mathbf{A} \mathbf{r}_k \in \mathbb{R}.$

$\mathbf{A}^* = \mathbf{A} > 0$ Conjugate Gradient

Suppose $\mathbf{A}$ is **positive definite**, i.e., $\mathbf{A}^* = \mathbf{A} > 0$.

Property. $(\mathbf{x}, \mathbf{y}) \equiv \mathbf{y}^* \mathbf{A}^{-1} \mathbf{x}$ is an inner product: the **$\mathbf{A}^{-1}$ inner product**.

Replace standard inner product by the $\mathbf{A}^{-1}$ inner product.

$\mathbf{r}^* \mathbf{c}_1 \rightsquigarrow \mathbf{r}^* \mathbf{A}^{-1} \mathbf{c}_1 = \mathbf{r}^* \mathbf{r} = \|\mathbf{r}\|_2^2$ **Norm r comes for free!**

$\mathbf{c}^* \mathbf{c} \rightsquigarrow \mathbf{c}^* \mathbf{A}^{-1} \mathbf{c} = \mathbf{c}^* \mathbf{u}$ **No $\mathbf{A}^{-1}$ needed!**

$\mathbf{r}_k \perp \mathbf{A} \mathbf{r}_j \rightsquigarrow \mathbf{r}_k \perp_{\mathbf{A}^{-1}} \mathbf{A} \mathbf{r}_j \Leftrightarrow \mathbf{r}_k \perp \mathbf{r}_j$: **orthogonal residuals**.

Additional saving of

1 DOT (norm **r** for free) and 1 AXPY $\rightsquigarrow$ CG