

Anderson Acceleration for Fixed-Point Iteration

Homer Walker
Mathematical Sciences Department
Worcester Polytechnic Institute

Joint in part with Peng Ni. Thanks to Juan Meza, Chao Yang, and Yousef Saad.
Research supported in part by NSF Grant DMS-0915183.

Fixed-Point Iteration

Given $g : \mathbb{R}^n \rightarrow \mathbb{R}^n, \dots$

Fixed-point problem: Find $x_* \in \mathbb{R}^n$ such that $x_* = g(x_*)$.

Fixed-Point Iteration: Given x_0 .

For $k = 0, 1, \dots$

Set $x_{k+1} = g(x_k)$.

Fixed-point iteration is very widely used in CS&E.

WHY?

Typically, ...

- Convergence is linear at best, often slow, often in doubt.
- “Globalization” is unavailable.
- The problem can be recast as $f(x_*) = 0$, where $f(x) \equiv g(x) - x$, for which there are many very effective algorithms and codes.

But there are often good reasons ...

- ▶ ease of implementation,
- ▶ second derivative information unnecessary,
- ▶ desirable structure preserved,
- ▶ low cost per iteration,
- ▶ parallel advantages,
- ▶ constraints satisfied.

Example: The Self-Consistent Field (SCF) Iteration

Context: Electronic structure computations.

Physics problem: Find the ground state energy and electron wave functions for a many-electron system.

Leads to the ...

Kohn–Sham (1965) equations: $H(X)X = X\Lambda$, $X^T X = I$.

SCF Iteration: Given X_0 .

For $k = 0, 1, \dots$

Evaluate $H(X_k)$ and solve $H(X_k)X_{k+1} = X_{k+1}\Lambda_{k+1}$.

Example: The Expectation-Maximization (EM) Algorithm

Formalized by Dempster, Laird, and Rubin (1977).

Context: Statistical estimation (maximum-likelihood) using *incomplete data*.

General idea: Determine the next approximate MLE to *maximize* the *expectation* of the complete-data log-likelihood function, given the observed incomplete data and the current approximate MLE.

Marvelous property: The likelihood function *increases* at each iteration.

Particular problem: Estimate the parameters in a *mixture density*

$$p(x|\Phi) = \sum_{i=1}^m \alpha_i p_i(x|\phi_i)$$

using an “unlabeled” sample on the mixture.

Example: The EM Algorithm (cont.)

Assume *normal densities*: For each i , $\phi_i = (\mu_i, \Sigma_i)$ and

$$p_i(x|\phi_i) = \frac{1}{(2\pi)^{n/2}(\det \Sigma_i)^{1/2}} e^{-(x-\mu_i)^T \Sigma_i^{-1}(x-\mu_i)/2}$$

EM iteration: For $i = 1, \dots, m$,

$$\alpha_i^+ = \frac{1}{N} \sum_{k=1}^N \frac{\alpha_i p_i(x_k|\phi_i)}{p(x_k|\Phi)},$$


$$\mu_i^+ = \left\{ \sum_{k=1}^N x_k \frac{\alpha_i p_i(x_k|\phi_i)}{p(x_k|\Phi)} \right\} / \left\{ \sum_{k=1}^N \frac{\alpha_i p_i(x_k|\phi_i)}{p(x_k|\Phi)} \right\},$$

$$\Sigma_i^+ = \left\{ \sum_{k=1}^N (x_k - \mu_i^+)(x_k - \mu_i^+)^T \frac{\alpha_i p_i(x_k|\phi_i)}{p(x_k|\Phi)} \right\} / \left\{ \sum_{k=1}^N \frac{\alpha_i p_i(x_k|\phi_i)}{p(x_k|\Phi)} \right\}.$$


EM Iterations Demo

A Univariate Normal Mixture.

- ▶ $p_i(x|\phi_i) = \frac{1}{\sqrt{2\pi\sigma_i^2}} e^{-(x-\mu_i)^2/(2\sigma_i^2)}$ for $i = 1, \dots, 5$.
- ▶ Sample of 10,000 observations.
 - $[\alpha_1, \dots, \alpha_5] = [.2, .3, .3, .1, .1]$
 - $[\mu_1, \dots, \mu_5] = [0, 1, 2, 3, 4]$,
 - $[\sigma_1^2, \dots, \sigma_5^2] = [.2, 2, .5, .1, .1]$.
- ▶ EM iterations on the means: $\mu_i^+ = \left\{ \sum_{k=1}^N x_k \frac{\alpha_i p_i(x_k|\phi_i)}{p(x_k|\Phi)} \right\} / \left\{ \sum_{k=1}^N \frac{\alpha_i p_i(x_k|\phi_i)}{p(x_k|\Phi)} \right\}$.



Sample histogram and graph of the mixture PDF.



Convergence of the EM iterates.

Anderson Acceleration

Derived from a method of [D. G. Anderson](#), *Iterative procedures for nonlinear integral equations*, J. Assoc. Comput. Machinery, 12 (1965), 547–560.

Suppose we have a fixed-point iteration $x_+ = g(x)$, $g : \mathbf{R}^n \rightarrow \mathbf{R}^n$.

Anderson Acceleration: Given x_0 and $m \geq 1$.

Set $x_1 = g(x_0)$.

Iterate: For $k = 1, 2, \dots$

Set $m_k = \min\{m, k\}$.

Set $F_k = (f_{k-m_k}, \dots, f_k)$, where $f_i = g(x_i) - x_i$.

Solve $\min_{\alpha \in \mathbf{R}^{m_k+1}} \|F_k \alpha\|_2$ s. t. $\sum_{i=0}^{m_k} \alpha_i = 1$.

Set $x_{k+1} = \sum_{i=0}^{m_k} \alpha_i g(x_{k-m_k+i})$.

Acceleration Methods — Perspectives

Mathematicians and numerical analysts have extensively studied ...

- ▶ *vector-extrapolation* methods, especially polynomial methods: (reduced-rank, minimal-polynomial, modified minimal-polynomial);
- ▶ *vector* and *topological ε -algorithms*.

See the book by Brezinski & Redivo-Zaglia (1991); surveys by Brezinski (2000), Jbilou–Sadok (2000), Smith–Ford–Sidi (1987), ...

Anderson acceleration is in a distinct category:

- ▶ Eirola–Nevanlinna (1989) and U. Yang (1995) methods and related methods based on quasi-Newton updating;
- ▶ “charge-mixing” methods for electronic-structure computations: Pulay (1980, 1982), Kerker (1981), Thomas–Fermi (Raczkowski et al. 2001), Broyden (Kresse–Furthmüller 1996), ...

See

- ▶ [H. Fang and Y. Saad \(2008\)](#), *Two classes of multiseant methods for nonlinear acceleration*, Numer. Linear Algebra Appl., 16, 197-221;
- ▶ [C. Yang, J. C. Meza, B. Lee, and L.-W. Wang \(2008\)](#), *KSSOLV — a MATLAB toolbox for solving the Kohn–Sham equations*, ACM Trans. Math. Softw., 36, 1-35;
- ▶ [G. Kresse and J. Furthmüller \(1996\)](#), *Efficient iterative schemes for ab initio total-energy calculations using a plane-wave basis set*, Phys. Rev. B, 54, 169-186.

Anderson Acceleration — Rationale

Suppose $g(x) = Ax + b$ for $A \in \mathbb{R}^{n \times n}$ and $b \in \mathbb{R}^n$.

Then $x_{k+1} = \sum_{i=0}^{m_k} \alpha_i g(x_{k-m_k+i}) = g(\sum_{i=0}^{m_k} \alpha_i x_{k-m_k+i})$. Thus

$$x_{k+1} = g(x_{min})$$

where $x_{min} = \sum_{i=0}^{m_k} \alpha_i x_{k-m_k+i}$ has minimal residual within the affine subspace containing $\{x_{k-m_k+i}\}_{i=0, \dots, m_k}$.

Anderson Acceleration and GMRES

- ▶ *Anderson acceleration is “essentially equivalent” to GMRES on linear problems.*

Assume ...

- ▶ Anderson acceleration is not truncated, i.e., $m_k = k$ for each k .
- ▶ g is linear, i.e., $g(x) = Ax + b$ for $A \in \mathbf{R}^{n \times n}$ and $b \in \mathbf{R}^n$.
- ▶ $(I - A)$ is nonsingular.
- ▶ GMRES is applied to $(I - A)x = b$ with initial point x_0 .

Theorem. *Suppose these assumptions hold and that, for some $k > 0$, $r_{k-1}^{\text{GMRES}} \neq 0$ and also $\|r_{j-1}^{\text{GMRES}}\|_2 > \|r_j^{\text{GMRES}}\|_2$ for each j such that $0 < j < k$. Then*

$$\sum_{i=0}^k \alpha_i x_i^{\text{AA}} = x_k^{\text{GMRES}} \text{ and } x_{k+1}^{\text{AA}} = g(x_k^{\text{GMRES}}).$$

Anderson Acceleration and GMRES (cont.)

Consider ...

- ▶ $Ax = b$ for nonsingular $A \in \mathbb{R}^n$.
- ▶ An operator splitting $A = M - N$ for nonsingular $M \in \mathbb{R}^{n \times n}$.
- ▶ A stationary iteration $x_+ = g(x) \equiv M^{-1}Nx + M^{-1}b$.

Assume ...

- ▶ Anderson acceleration is not truncated, i.e., $m_k = k$ for each k .
- ▶ GMRES is applied to $M^{-1}Ax = M^{-1}b$ with initial point x_0 .

Corollary. Suppose these hold and that, for some $k > 0$, $r_{k-1}^{\text{GMRES}} \neq 0$ and also $\|r_{j-1}^{\text{GMRES}}\|_2 > \|r_j^{\text{GMRES}}\|_2$ for each j such that $0 < j < k$. Then

$$\sum_{i=0}^k \alpha_i x_i^{\text{AA}} = x_k^{\text{GMRES}} \text{ and } x_{k+1}^{\text{AA}} = g(x_k^{\text{GMRES}}) = M^{-1}N x_k^{\text{GMRES}} + M^{-1}b.$$

Anderson Acceleration and Multi-Secant Updating

In general ...

Eyert (1996), Fang–Saad (2008): *Successive Anderson acceleration iterates are related by*

$$x_{k+1} = x_k - B_k^{-1}f(x_k), \quad (\star)$$

where $f(x) \equiv g(x) - x$ and B_k is the second (“bad”) Broyden *multi-secant* update of $-I$ satisfying

$$B_k(x_{i+1} - x_i) = f(x_{i+1}) - f(x_i), \quad k - m_k \leq i \leq k - 1. \quad (\star\star)$$

- ▶ An *Anderson family* of methods of the form $(\star)$ is defined in Fang–Saad (2008). In this, Anderson acceleration is designated the *Type II* method.
- ▶ The Anderson family *Type I* method has the form $(\star)$, in which B_k is the first (“good”) Broyden multi-secant update of $-I$ satisfying $(\star\star)$.
- ▶ It is straightforward to cast the Type I method in the form of the Anderson acceleration algorithm.

Anderson Acceleration and Multi-Secant Updating (cont.)

- ▶ The Anderson family Type I method is “essentially equivalent” to the Arnoldi (FOM) method on linear problems.

Assume ...

- ▶ The Type I method is not truncated, i.e., $m_k = k$ for each k .
- ▶ g is linear, i.e., $g(x) = Ax + b$ for $A \in \mathbb{R}^{n \times n}$ and $b \in \mathbb{R}^n$.
- ▶ $(I - A)$ is nonsingular.
- ▶ The Arnoldi method is applied to $(I - A)x = b$ with initial point x_0 .

Theorem. Suppose these assumptions hold and that, for some $k > 0$, x_j^{Arnoldi} is defined for $0 < j \leq k$ and $r_{k-1}^{\text{Arnoldi}} \neq 0$. Then

$$\sum_{i=0}^k \alpha_i x_i^{\text{Type I}} = x_k^{\text{Arnoldi}} \quad \text{and} \quad x_{k+1}^{\text{Type I}} = g(x_k^{\text{Arnoldi}}).$$

Anderson Acceleration and Multi-Secant Updating (cont.)

Consider . . .

- ▶ $Ax = b$ for nonsingular $A \in \mathbb{R}^n$.
- ▶ An operator splitting $A = M - N$ for nonsingular $M \in \mathbb{R}^{n \times n}$.
- ▶ A stationary iteration $x_+ = g(x) \equiv M^{-1}Nx + M^{-1}b$.

Assume . . .

- ▶ The Type I method is not truncated, i.e., $m_k = k$ for each k .
- ▶ The Arnoldi method is applied to $M^{-1}Ax = M^{-1}b$ with initial point x_0 .

Corollary. Suppose these hold and that, for some $k > 0$, x_j^{Arnoldi} is defined for $0 < j \leq k$ and $r_{k-1}^{\text{Arnoldi}} \neq 0$. Then

$$\sum_{i=0}^k \alpha_i x_i^{\text{Type I}} = x_k^{\text{Arnoldi}} \text{ and } x_{k+1}^{\text{Type I}} = g(x_k^{\text{Arnoldi}}).$$

Implementation

The implementational issue is *solving the least-squares problem*:

$$\min_{\alpha \in \mathbb{R}^{m_k+1}} \|F_k \alpha\|_2 \quad \text{s.t.} \quad \sum_{i=0}^{m_k} \alpha_i = 1.$$

Note that ...

- F_{k+1} is obtained from F_k by adding a column on the right and, if $k \geq m$, also dropping a column on the left.
- In practice, additional columns may be dropped on the left to maintain acceptable conditioning.

These low-rank changes allow efficient QR solution of the least-squares problem by updating the Q and R factors at a cost of $O(m_k n)$ flops/iteration.


Good approach: As in Fang–Saad (2008), recast as an unconstrained problem

$$\min_{\gamma \in \mathbb{R}^{m_k}} \|f_k - \mathcal{F}_k \gamma\|_2,$$


where $\mathcal{F}_k = (\Delta f_{k-m_k}, \dots, \Delta f_{k-1})$ and $\Delta f_i = f_{i+1} - f_i$ for each i .

EM Iterations Demo (cont.)

- ▶ $p_i(x|\phi_i) = \frac{1}{\sqrt{2\pi\sigma_i^2}} e^{-(x-\mu_i)^2/(2\sigma_i^2)}$ for $i = 1, \dots, 5$.
- ▶ Sample of 10,000 observations.
 - $[\alpha_1, \dots, \alpha_5] = [.2, .3, .3, .1, .1]$
 - $[\mu_1, \dots, \mu_5] = [0, 1, 2, 3, 4]$,
 - $[\sigma_1^2, \dots, \sigma_5^2] = [.2, 2, .5, .1, .1]$.
- ▶ EM iterations on the means: $\mu_i^+ = \left\{ \sum_{k=1}^N x_k \frac{\alpha_i p_i(x_k|\phi_i)}{p(x_k|\Phi)} \right\} / \left\{ \sum_{k=1}^N \frac{\alpha_i p_i(x_k|\phi_i)}{p(x_k|\Phi)} \right\}$.



Sample histogram and graph of the mixture PDF.




Convergence of the EM iterates with (red) and without (blue) Anderson acceleration.

Example: EM Convergence and “Separation”

A Univariate Normal Mixture.

- ▶ $p_i(x|\phi_i) = \frac{1}{\sqrt{2\pi\sigma_i^2}} e^{-(x-\mu_i)^2/(2\sigma_i^2)}$ for $i = 1, \dots, 3$.
- ▶ EM iterations on the means: $\mu_i^+ = \left\{ \sum_{k=1}^N x_k \frac{\alpha_i p_i(x_k|\phi_i)}{p(x_k|\Phi)} \right\} / \left\{ \sum_{k=1}^N \frac{\alpha_i p_i(x_k|\phi_i)}{p(x_k|\Phi)} \right\}$.
- ▶ Sample of 100,000 observations.
 - $[\alpha_1, \alpha_2, \alpha_3] = [.3, .3, .4]$, $[\sigma_1^2, \sigma_2^2, \sigma_3^2] = [1, 1, 1]$.
 - $[\mu_1, \mu_2, \mu_3] = [0, 2, 4]$, $[0, 1, 2]$, $[0, .5, 1]$.

Convergence of the EM iterates with (red) and without (blue) Anderson acceleration. With acceleration, the convergence does not slow significantly as the separation worsens.



Example: ALS in Nonnegative Matrix Factorization

Problem: Given $A \in \mathbb{R}^{m \times n}$, find $W \in \mathbb{R}^{m \times k}$ and $H \in \mathbb{R}^{k \times n}$ to solve

$$\min_{W \geq 0, H \geq 0} \|A - WH\|_F.$$

Alternating Nonnegative Least-Squares : Given W_0 .


For $k = 0, 1, \dots$

Evaluate $H_k = \arg \min_{H \geq 0} \|A - W_k H\|_F$.

Evaluate $W_{k+1} = \arg \min_{W \geq 0} \|A - WH_k\|_F$.

Apply to the *term-document matrix* from Example 1.1 of Eldén (2007) with $k = 3$:

$$A = \begin{pmatrix} 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 1 \\ 1 & 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 1 & 0 & 1 & 1 & 0 \\ 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 & 1 \\ 0 & 1 & 1 & 0 & 0 \end{pmatrix}$$






Convergence of the ALS iterates with (red) and without (blue) Anderson acceleration.

Example: Domain Decomposition, Linear Problem

A Convection-Diffusion Problem.

$$\begin{aligned}\Delta u + cu + du_x + eu_y &= f \quad \text{in } \mathcal{D} = [0, 1] \times [0, 1], \\ u &= 0 \quad \text{on } \partial\mathcal{D}.\end{aligned}$$

- ▶ Centered differences, 128×128 grid.
- ▶ 4×4 array of subdomains, 32×32 grid per subdomain.
- ▶ Restricted additive Schwarz (Cai-Sarkis, 1999), 3 grid lines overlap, no coarse grid.
- ▶ $m = 50$ in Anderson acceleration.



Convergence of the RAS iterates with (red) and without (blue) Anderson acceleration. The green curve shows convergence of RAS-preconditioned GMRES.


Example: Domain Decomposition, Easy Nonlinear Problem

The Bratu Problem.

$$\begin{aligned}\Delta u + \lambda e^u &= 0 \quad \text{in } \mathcal{D} = [0, 1] \times [0, 1], \\ u &= 0 \quad \text{on } \partial\mathcal{D}.\end{aligned}$$

- ▶ Centered differences, 128×128 grid.
- ▶ 4×4 array of subdomains, 32×32 grid per subdomain.
- ▶ Restricted additive Schwarz (Cai-Sarkis, 1999), 3 grid lines overlap, no coarse grid.
- ▶ $m = 50$ in Anderson acceleration.
- ▶ Newton-backtracking subdomain solves.

Convergence of the nonlinear RAS iterates with (red) and without (blue) Anderson acceleration. The green curve shows convergence of full-domain Newton's method with backtracking.




Example: Domain Decomposition, Hard Nonlinear Problem


Transonic Flow in a Duct.

$$\begin{aligned} [A(x)\rho(u)u_x]_x &= 0, & 0 < x < 2, \\ u(0) &= 0, & u(2) = u_R, \end{aligned} \quad \text{where} \quad \begin{aligned} \rho(u) &= \left(1 + \frac{\gamma-1}{2}(1-u^2)\right)^{1/(\gamma-1)} \\ A(x) &= 0.4 + 0.6(x-1)^2 \end{aligned}$$

- ▶ Finite-difference discretization and parameter choices as in Cai–Keyes–Young (2001): $u_R = 1.15$, $\gamma = 1.4$.
- ▶ 512 interior gridpoints, 8 subdomains, 64 gridpoints per subdomains.
- ▶ Restricted additive Schwarz (Cai–Sarkis, 1999), 8 grid lines overlap, no coarse grid.
- ▶ $m = 20$ in Anderson acceleration.
- ▶ “matrix-free” Newton-GMRES-backtracking subdomain solves.



Nonlinear RAS iterations with (red) and without (blue) Anderson acceleration; “matrix-free” Newton-GMRES-backtracking iterations (green).



Mach distribution and shock location for the duct flow problem.


Example: Picard Iteration, Easy Problem

The Bratu Problem.

$$\begin{aligned}\Delta u + \lambda e^u &= 0 \quad \text{in } \mathcal{D} = [0, 1] \times [0, 1], \\ u &= 0 \quad \text{on } \partial\mathcal{D}.\end{aligned}$$

- ▶ Centered differences, 256×256 grid.
- ▶ Picard iteration: $u^+ = -\Delta^{-1}(\lambda e^u)$.
- ▶ $m = 10$ in Anderson acceleration.

Convergence of the Picard iterates with (red) and without (blue) Anderson acceleration. The green curve shows convergence of Newton's method with backtracking.



Example: Picard Iteration, Hard Problem





A Variably Saturated Flow Problem.

Ref.: I. Yavneh and G. Dardyk, *A multilevel nonlinear method*, SISC, 28 (2006), 24-46.

$$\begin{aligned} [g(u)u_x]_x &= 0, \quad 0 < x < 1, \\ u(0) &= -2, \quad u(1) = 1, \end{aligned} \quad \text{where} \quad \begin{aligned} g(u) &= \begin{cases} 1, & u \geq 0 \\ k_s \psi(|u|), & u < 0 \end{cases} \\ \psi(\theta) &= [1 + \alpha \theta^p]^{-q/2} \left(1 - \frac{(\alpha \theta)^{p-1}}{[1 + (\alpha \theta)^p]^q} \right)^2 \end{aligned}$$

- ▶ Discretization and parameter choices as in Yavneh-Dardyk (2006), §4.1:
 - $n = 256$;
 - $k_s = 1$;
 - $\alpha = 0.5, 1.0$;
 - $p = 2.5, 1.8$ and $q = 1 - p$.
- ▶ “For a given domain and boundary conditions, this problem generally becomes more difficult for smaller p and larger α . For $p < 2$, g' fails to satisfy a Lipschitz condition at $u = 0$, and $g(u)$ actually tends to a step function for $p \rightarrow 1+$. For α that is not small compared to one, the solution exhibits a thin boundary layer near $x = 0$ (in our example), which narrows quickly as α grows.”

Example: Picard Iteration, Variably Saturated Flow (cont.)



Picard iterations with (red) and without (blue) Anderson acceleration; finite-difference Newton-backtracking iterations (green).

In Conclusion ...

- Anderson acceleration is successfully used in at least one important application and shows promise in the examples given here.
- Theoretical results give *some* insights into its behavior.
 - ▶ Without truncation, it is essentially equivalent to GMRES on linear problems (but *not* a recommended substitute).
 - ▶ It is related to multi-secant (“bad” Broyden) updating methods.
- Efficiency and (relatively) good conditioning can be obtained through updated *QR* factorizations.
- Limited experiments suggest fast local convergence is possible — *if* enough past residuals are retained.

But ...

- There is no satisfactory local convergence theory.
 - There are no “globalizations”.
 - There are possible numerical weaknesses.
-
- ▶ **Much algorithmic research remains to be done.**
 - ▶ **Many possible applications remain to be explored.**