# Regularized Nonlinear Acceleration. 

Alexandre d'Aspremont, CNRS \& D.I. ENS.

with Damien Scieur \& Francis Bach.
Support from ERC SIPA and ITN MacSeNet.

## Introduction

## Generic convex optimization problem

$$
\min _{x \in \mathbb{R}^{n}} f(x)
$$

## Introduction

Algorithms produce a sequence of iterates.


We only keep the last (or best) one. . .

## Introduction

Aitken's $\Delta^{2}$ [Aitken, 1927]. Given a sequence $\left\{s_{k}\right\}_{k=1, \ldots} \in \mathbb{R}^{\mathbb{N}}$ with limit $s_{*}$, and suppose

$$
s_{k+1}-s_{*}=a\left(s_{k}-s_{*}\right), \quad \text { for } k=1, \ldots
$$

We can compute $a$ using

$$
s_{k+1}-s_{k}=a\left(s_{k}-s_{k-1}\right) \quad \Rightarrow \quad a=\frac{s_{k+1}-s_{k}}{s_{k}-s_{k-1}}
$$

and get the limit $s^{*}$ by solving

$$
s_{k+1}-s^{*}=\frac{s_{k+1}-s_{k}}{s_{k}-s_{k-1}}\left(s_{k}-s^{*}\right)
$$

which yields

$$
s^{*}=\frac{s_{k-1} s_{k+1}-s_{k}^{2}}{s_{k+1}-2 s_{k}+s_{k-1}}
$$

This is Aitken's $\Delta^{2}$ and allows us to compute $s_{*}$ from $\left\{s_{k+1}, s_{k}, s_{k-1}\right\}$.

## Introduction

Aitken's $\Delta^{2}$ [Aitken, 1927], again. Given a sequence $\left\{s_{k}\right\}_{k=1, \ldots} \in \mathbb{R}^{\mathbb{N}}$ with limit $s_{*}$, and suppose that for $k=1, \ldots$,

$$
a_{0}\left(s_{k}-s_{*}\right)+a_{1}\left(s_{k+1}-s_{*}\right)=0 \quad \text { and } a_{0}+a_{1}=1 \quad \text { (normalization) }
$$

We have

$$
\begin{aligned}
\underbrace{\left(a_{0}+a_{1}\right)}_{=1} & s_{*}
\end{aligned}=a_{0} s_{k-1}+a_{1} s_{k}, ~ \begin{aligned}
& =a_{0}\left(s_{k}-s_{k-1}\right)+a_{1}\left(s_{k+1}-s_{k}\right)
\end{aligned}
$$

We get $s^{*}$ using

$$
\left[\begin{array}{ccc}
0 & s_{k+1}-s_{k} & s_{k}-s_{k-1} \\
-1 & s_{k} & s_{k-1} \\
0 & 1 & 1
\end{array}\right] \cdot\left[\begin{array}{c}
s^{*} \\
a_{1} \\
a_{0}
\end{array}\right]=\left[\begin{array}{l}
0 \\
0 \\
1
\end{array}\right] \Leftrightarrow \quad s^{*}=\frac{\left|\begin{array}{cc}
s_{k+1}-s_{k} & s_{k}-s_{k-1} \\
s_{k} & s_{k-1}
\end{array}\right|}{\left|\begin{array}{cc}
s_{k+1}-s_{k} & s_{k}-s_{k-1} \\
1 & 1
\end{array}\right|}
$$

Same formula as before, but generalizes to higher dimensions.

## Introduction

## Convergence acceleration. Consider

$$
s_{k}=\sum_{i=0}^{k} \frac{(-1)^{i}}{(2 i+1)} \quad \xrightarrow{k \rightarrow \infty} \quad \frac{\pi}{4}=0.785398 \ldots
$$

we have

| $k$ | $\frac{(-1)^{k}}{(2 k+1)}$ | $\sum_{i=0}^{k} \frac{(-1)^{i}}{(2 i+1)}$ | $\Delta^{2}$ |
| :--- | :--- | :--- | :--- |
| 0 | 1 | 1.0000 | - |
| 1 | -0.33333 | 0.66667 | - |
| 2 | 0.2 | 0.86667 | $\mathbf{0 . 7 9 1 6 7}$ |
| 3 | -0.14286 | $\mathbf{0 . 7 2 3 8 1}$ | $\mathbf{0 . 7 8 3 3 3}$ |
| 4 | 0.11111 | 0.83492 | $\mathbf{0 . 7 8 6 3 1}$ |
| 5 | -0.090909 | $\mathbf{0 . 7 4 4 0 1}$ | $\mathbf{0 . 7 8 4 9 2}$ |
| 6 | 0.076923 | 0.82093 | $\mathbf{0 . 7 8 5 6 8}$ |
| 7 | -0.066667 | $\mathbf{0 . 7 5 4 2 7}$ | $\mathbf{0 . 7 8 5 2 2}$ |
| 8 | 0.058824 | 0.81309 | $\mathbf{0 . 7 8 5 5 2}$ |
| 9 | -0.052632 | $\mathbf{0 . 7 6 0 4 6}$ | $\mathbf{0 . 7 8 5 3 1}$ |

## Introduction

## Convergence acceleration.

- Similar results apply to sequences satisfying

$$
\sum_{i=0}^{k} a_{i}\left(s_{n+i}-s_{*}\right)=0
$$

using Aitken's ideas recursively.
■ This produces Wynn's $\varepsilon$-algorithm [Wynn, 1956].

- See [Brezinski, 1977] for a survey on acceleration, extrapolation.
- Directly related to the Levinson-Durbin algo on AR processes.

■ Vector case: focus on Minimal Polynomial Extrapolation [Sidi et al., 1986].

Overall: a simple postprocessing step.

## Outline

- Introduction
- Minimal Polynomial Extrapolation
- Regularized MPE
- Numerical results


## Minimal Polynomial Extrapolation

Quadratic example. Minimize

$$
f(x)=\frac{1}{2}\|B x-b\|_{2}^{2}
$$

using the basic gradient algorithm, with

$$
x_{k+1}:=x_{k}-\frac{1}{L}\left(B^{T} B x_{k}-b\right) .
$$

we get

$$
x_{k+1}-x^{*}:=\underbrace{\left(\mathbf{I}-\frac{1}{L} B^{T} B\right)}_{A}\left(x_{k}-x^{*}\right)
$$

since $B^{T} B x^{*}=b$.

This means $x_{k+1}-x^{*}$ follows a vector autoregressive process.

## Minimal Polynomial Extrapolation

We have

$$
\sum_{i=0}^{k} c_{i}\left(x_{i}-x^{*}\right)=\sum_{i=1}^{k} c_{i} A^{i}\left(x_{0}-x^{*}\right)
$$

and setting $\mathbf{1}^{T} c=1$, yields

$$
\left(\sum_{i=0}^{k} c_{i} x_{i}\right)-x^{*}=p(A)\left(x_{0}-x^{*}\right), \quad \text { where } p(v)=\sum_{i=1}^{k} c_{i} v^{i}
$$

- Setting $c$ such that $p(A)\left(x_{0}-x^{*}\right)=0$, we would have

$$
\mathrm{x}^{*}=\sum_{\mathrm{i}=0}^{\mathrm{k}} \mathrm{c}_{\mathrm{i}} \mathrm{x}_{\mathrm{i}}
$$

- Get the limit by averaging iterates (using weights depending on $x_{k}$ ).
- We typically do not observe $A$ (or $x^{*}$ ).
- How do we extract $c$ from the iterates $x_{k}$ ?


## Minimal Polynomial Extrapolation

We have

$$
\begin{aligned}
x_{k}-x_{k-1} & =\left(x_{k}-x^{*}\right)-\left(x_{k-1}-x^{*}\right) \\
& =(A-\mathbf{I}) A^{k-1}\left(x_{0}-x^{*}\right)
\end{aligned}
$$

hence if $p(A)=0$, we must have

$$
\sum_{i=1}^{k} c_{i}\left(x_{i}-x_{i-1}\right)=(A-\mathbf{I}) p(A)\left(x_{0}-x^{*}\right)=0
$$

so if $(A-\mathbf{I})$ is nonsingular, the coefficient vector $c$ solves the linear system

$$
\left\{\begin{array}{l}
\sum_{i=1}^{k} c_{i}\left(x_{i}-x_{i-1}\right)=0 \\
\sum_{i=1}^{k} c_{i}=1
\end{array}\right.
$$

and $p(\cdot)$ is the minimal polynomial of $A$ w.r.t. $\left(x_{0}-x^{*}\right)$.

## Approximate Minimal Polynomial Extrapolation

## Approximate MPE.

- For $k$ smaller than the degree of the minimal polynomial, we find $c$ that minimizes the residual

$$
\left\|(A-\mathbf{I}) p(A)\left(x_{0}-x^{*}\right)\right\|_{2}=\left\|\sum_{i=1}^{k} c_{i}\left(x_{i}-x_{i-1}\right)\right\|_{2}
$$

- Setting $U \in \mathbb{R}^{n \times k+1}$, with $U_{i}=x_{i+1}-x_{i}$, this means solving

$$
\begin{equation*}
c^{*} \triangleq \underset{1}{\operatorname{argmin}}\|U c\|_{2} \tag{AMPE}
\end{equation*}
$$

in the variable $c \in \mathbb{R}^{k+1}$.

- Also known as Eddy-Mešina method [Mešina, 1977, Eddy, 1979] or Reduced Rank Extrapolation with arbitrary $k$ (see [Smith et al., 1987, §10]).


## Uniform Bound

Chebyshev polynomials. Crude bound on $\left\|U c^{*}\right\|_{2}$ using Chebyshev polynomials, to bound error as a function of $k$, with

$$
\begin{aligned}
\left\|\sum_{i=0}^{k} c_{i}^{*} x_{i}-x^{*}\right\|_{2} & =\left\|(I-A)^{-1} \sum_{i=0}^{k} c_{i}^{*} U_{i}\right\|_{2} \\
& \leq\left\|(I-A)^{-1}\right\|_{2}\left\|p(A)\left(x_{1}-x_{0}\right)\right\|_{2}
\end{aligned}
$$

We have

$$
\begin{aligned}
\left\|p(A)\left(x_{1}-x_{0}\right)\right\|_{2} & \leq\|p(A)\|_{2}\left\|\left(x_{1}-x_{0}\right)\right\|_{2} \\
& =\max _{i=1, \ldots, n}\left|p\left(\lambda_{i}\right)\right|\left\|\left(x_{1}-x_{0}\right)\right\|_{2}
\end{aligned}
$$

where $0 \leq \lambda_{i} \leq \sigma$ are the eigenvalues of $A$. It suffices to find $p(\cdot) \in \mathbb{R}_{k}[x]$ solving

$$
\inf _{\left\{p \in \mathbb{R}_{k}[x]: p(1)=1\right\}} \sup _{v \in[0, \sigma]}|p(v)|
$$

Explicit solution using modified Chebyshev polynomials.

## Uniform Bound using Chebyshev Polynomials



Chebyshev polynomials $T_{3}(x, \sigma)$ and $T_{5}(x, \sigma)$ for $x \in[0,1]$ and $\sigma=0.85$. The maximum value of $T_{k}$ on $[0, \sigma]$ decreases geometrically fast when $k$ grows.

## Approximate Minimal Polynomial Extrapolation

## Proposition [Scieur, d'Aspremont, and Bach, 2016]

AMPE convergence. Let $A$ be symmetric, $0 \preceq A \preceq \sigma I$ with $\sigma<1$ and $c^{*}$ be the solution of (AMPE). Then

$$
\begin{equation*}
\left\|\sum_{i=0}^{k} c_{i}^{*} x_{i}-x^{*}\right\|_{2} \leq \kappa(A-I) \frac{2 \zeta^{k}}{1+\zeta^{2 k}}\left\|x_{0}-x^{*}\right\|_{2} \tag{1}
\end{equation*}
$$

where $\kappa(A-I)$ is the condition number of the matrix $A-I$ and $\zeta$ is given by

$$
\begin{equation*}
\zeta=\frac{1-\sqrt{1-\sigma}}{1+\sqrt{1-\sigma}}<\sigma \tag{2}
\end{equation*}
$$

Typically, $\sigma=1-\mu / L$ (gradient method) so the convergence rate is

$$
\left\|\sum_{i=0}^{k} c_{i}^{*} x_{i}-x^{*}\right\|_{2} \leq \kappa(A-I)\left(\frac{1-\sqrt{\mu / L}}{1+\sqrt{\mu / L}}\right)^{k}\left\|x_{0}-x^{*}\right\|_{2}
$$

## Approximate Minimal Polynomial Extrapolation

## AMPE versus Nesterov, conjugate gradient.

- Key difference with conjugate gradient: we do not observe $A$. . .
- Chebyshev polynomials satisfy a two-step recurrence. For quadratic minimization using the gradient method:

$$
\left\{\begin{array}{l}
z_{k-1}=y_{k-1}-\frac{1}{L}\left(B y_{k-1}-b\right) \\
y_{k}=\frac{\alpha_{k-1}}{\alpha_{k}}\left(\frac{2 z_{k-1}}{\sigma}-y_{k-1}\right)-\frac{\alpha_{k-2}}{\alpha_{k}} y_{k-2}
\end{array}\right.
$$

where $\alpha_{k}=\frac{2-\sigma}{\sigma} \alpha_{k-1}-\alpha_{k-2}$

- Nesterov's acceleration recursively computes a similar polynomial with

$$
\left\{\begin{array}{l}
z_{k-1}=y_{k-1}-\frac{1}{L}\left(B y_{k-1}-b\right) \\
y_{k}=z_{k-1}+\beta_{k}\left(z_{k-1}-z_{k-2}\right)
\end{array}\right.
$$

see also [Hardt, 2013].

## Approximate Minimal Polynomial Extrapolation

Accelerating optimization algorithms. For gradient descent, we have

$$
\tilde{x}_{k+1}:=\tilde{x}_{k}-\frac{1}{L} \nabla f\left(\tilde{x}_{k}\right)
$$

- This means $\tilde{x}_{k+1}-x^{*}:=A\left(\tilde{x}_{k}-x^{*}\right)+O\left(\left\|\tilde{x}_{k}-x^{*}\right\|_{2}^{2}\right)$ where

$$
A=I-\frac{1}{L} \nabla^{2} f\left(x^{*}\right),
$$

meaning that $\|A\|_{2} \leq 1-\frac{\mu}{L}$, whenever $\mu I \preceq \nabla^{2} f(x) \preceq L I$.

- Approximation error is a sum of three terms

$$
\left\|\sum_{i=0}^{k} \tilde{c}_{i} \tilde{x}_{i}-x^{*}\right\|_{2} \leq \underbrace{\left\|\sum_{i=0}^{k} c_{i} x_{i}-x^{*}\right\|_{2}}_{\text {AMPE }}+\underbrace{\left\|\sum_{i=0}^{k}\left(\tilde{c}_{i}-c_{i}\right) x_{i}\right\|_{2}}_{\text {Stability }}+\underbrace{\left\|\sum_{i=0}^{k} \tilde{c}_{i}\left(\tilde{x}_{i}-x_{i}\right)\right\|_{2}}_{\text {Nonlinearity }}
$$

Stability is key here.

## Approximate Minimal Polynomial Extrapolation

## Stability.

■ The iterations span a Krylov subspace

$$
\mathcal{K}_{k}=\operatorname{span}\left\{U_{0}, A U_{0}, \ldots, A^{k-1} U_{0}\right\}
$$

so the matrix $U$ in AMPE is a Krylov matrix.

- Similar to Hankel or Toeplitz case. $U^{T} U$ has a condition number typically growing exponentially with dimension [Tyrtyshnikov, 1994].
- In fact, the Hankel, Toeplitz and Krylov problems are directly connected, hence the link with Levinson-Durbin [Heinig and Rost, 2011].
- For generic optimization problems, eigenvalues are perturbed by deviations from the linear model, which can make the situation even worse.

Be wise, regularize . . .

## Outline

- Introduction
- Minimal Polynomial Extrapolation
- Regularized MPE
- Numerical results


## Regularized Minimal Polynomial Extrapolation

Regularized AMPE. Add a regularization term to AMPE.

- Regularized formulation of problem (AMPE),

$$
\begin{array}{ll}
\operatorname{minimize} & c^{T}\left(U^{T} U+\lambda I\right) c \\
\text { subject to } & \mathbf{1}^{T} c=1 \tag{RMPE}
\end{array}
$$

- Solution given by a linear system of size $k+1$.

$$
\begin{equation*}
c_{\lambda}^{*}=\frac{\left(U^{T} U+\lambda I\right)^{-1} \mathbf{1}}{\mathbf{1}^{T}\left(U^{T} U+\lambda I\right)^{-1} \mathbf{1}} \tag{3}
\end{equation*}
$$

## Regularized Minimal Polynomial Extrapolation

## Regularized AMPE.

## Proposition [Scieur et al., 2016]

Stability Let $c_{\lambda}^{*}$ be the solution of problem (RMPE). Then the solution of problem (RMPE) for the perturbed matrix $\tilde{U}=U+E$ is given by $c_{\lambda}^{*}+\Delta c_{\lambda}$ where

$$
\left\|\Delta c_{\lambda}\right\|_{2} \leq \frac{\|P\|_{2}}{\lambda}\left\|c_{\lambda}^{*}\right\|_{2}
$$

with $P=\tilde{U}^{T} \tilde{U}-U^{T} U$ the perturbation matrix.

## Regularized Minimal Polynomial Extrapolation

## RMPE algorithm.

Input: Sequence $\left\{x_{0}, x_{1}, \ldots, x_{k+1}\right\}$, parameter $\lambda>0$
1: Form $U=\left[x_{1}-x_{0}, \ldots, x_{k+1}-x_{k}\right]$
2: Solve the linear system $\left(U^{T} U+\lambda I\right) z=\mathbf{1}$
3: Set $c=z /\left(z^{T} \mathbf{1}\right)$
Output: Return $\sum_{i=0}^{k} c_{i} x_{i}$, approximating the optimum $x^{*}$

## Regularized Minimal Polynomial Extrapolation

Regularized AMPE. Define

$$
S(k, \alpha) \triangleq \min _{\left\{q \in \mathbb{R}_{k}[x]: q(1)=1\right\}}\left\{\max _{x \in[0, \sigma]}((1-x) q(x))^{2}+\alpha\|q\|_{2}^{2}\right\},
$$

## Proposition [Scieur et al., 2016]

Error bounds Let matrices $X=\left[x_{0}, x_{1}, \ldots, x_{k}\right], \tilde{X}=\left[x_{0}, \tilde{x}_{1}, \ldots, \tilde{x}_{k}\right]$ and scalar $\kappa=\left\|(A-I)^{-1}\right\|_{2}$. Suppose $\tilde{c}_{\lambda}^{*}$ solves problem (RMPE) and assume $A=g^{\prime}\left(x^{*}\right)$ symmetric with $0 \preceq A \preceq \sigma I$ where $\sigma<1$. Let us write the perturbation matrices $P=\tilde{U}^{T} \tilde{U}-U^{T} U$ and $\mathcal{E}=(X-\tilde{X})$. Then

$$
\left\|\tilde{X} \tilde{c}_{\lambda}^{*}-x^{*}\right\|_{2} \leq C(\mathcal{E}, P, \lambda) S\left(k, \lambda /\left\|x_{0}-x^{*}\right\|_{2}^{2}\right)^{\frac{1}{2}}\left\|x_{0}-x^{*}\right\|_{2}
$$

where

$$
C(\mathcal{E}, P, \lambda)=\left(\kappa^{2}+\frac{1}{\lambda}\left(1+\frac{\|P\|_{2}}{\lambda}\right)^{2}\left(\|\mathcal{E}\|_{2}+\kappa \frac{\|P\|_{2}}{2 \sqrt{\lambda}}\right)^{2}\right)^{\frac{1}{2}}
$$

## Regularized Minimal Polynomial Extrapolation

On the gradient method. Setting for instance $L=100, \mu=10, M=10^{-1}$, $\left\|x_{0}-x^{*}\right\|_{2}=10^{-4}$ and finally $\lambda=\|P\|_{2}$.



Left: Relative value for the regularization parameter $\lambda$ used in the theoretical bound. Right: Convergence speedup relative to the gradient method, for Nesterov's accelerated method and the theoretical RMPE.

## Regularized Minimal Polynomial Extrapolation

## Proposition [Scieur et al., 2016]

Asymptotic acceleration Using the gradient method with stepsize in $] 0, \frac{2}{L}[$ on a L-smooth, $\mu$-strongly convex function $f$ with Lipschitz-continuous Hessian of constant $M$.

$$
\left\|\tilde{X} \tilde{c}_{\lambda}^{*}-x^{*}\right\|_{2} \leq \kappa\left(1+\frac{\left(1+\frac{1}{\beta}\right)^{2}}{4 \beta^{2}}\right)^{1 / 2} \frac{2 \zeta^{k}}{1+\zeta^{2 k}}\left\|x_{0}-x^{*}\right\|
$$

with

$$
\zeta=\frac{1-\sqrt{\mu / L}}{1+\sqrt{\mu / L}}
$$

for $\left\|x_{0}-x^{*}\right\|$ small enough, where $\lambda=\beta\|P\|_{2}$ and $\kappa=\frac{L}{\mu}$ is the condition number of the function $f(x)$.

We (asymptotically) recover the accelerated rate in [Nesterov, 1983].

## Regularized Minimal Polynomial Extrapolation

## Complexity, online mode.

- Cholesky updates. Given the Cholesky factorization $L L^{T}=\tilde{U}^{T} \tilde{U}+\lambda I$ and a new vector $u_{+}$,

$$
L_{+} L_{+}^{T}=\left[\begin{array}{cc}
L & 0 \\
a^{T} & b
\end{array}\right]\left[\begin{array}{cc}
L^{T} & a \\
0 & b
\end{array}\right]=\left[\begin{array}{cc}
\tilde{U}^{T} \tilde{U}+\lambda I & \tilde{U}^{T} u_{+} \\
\left(\tilde{U}^{T} u_{+}\right)^{T} & u_{+}^{T} u_{+}+\lambda
\end{array}\right] .
$$

the solutions $a$ and $b$ are

$$
a=L^{-1} \tilde{U}^{T} u_{+}, \quad b=a^{T} a+\lambda
$$

- The complexity of an update at iteration $i$ is $O\left(i n+i^{2}\right)$, so the overall complexity after $k$ iterations is

$$
O\left(n k^{2}+k^{3}\right)
$$

In the experiments that follow, $k$ is typically 5. . .

## Regularized Minimal Polynomial Extrapolation

Smooth functions. Suppose $f$ is not strongly convex.

- The function

$$
\min _{x \in \mathbb{R}^{n}} f_{\varepsilon}(x) \triangleq f(x)+\frac{\varepsilon}{2 D^{2}}\|x\|_{2}^{2}
$$

has a Lipschitz continuous gradient with parameter $L+\varepsilon / D^{2}$ and is strongly convex with parameter $\varepsilon / D^{2}$.

- Accelerated algorithm converge with a linear rate, with a bound equivalent to

$$
\sqrt{1+\frac{L D^{2}}{\varepsilon}}
$$

which matches the optimal complexity bound for smooth functions.

Handling the strongly convex case, allows us to produce bounds in the smooth case, on paper. . .

## Outline

- Introduction
- Minimal Polynomial Extrapolation
- Regularized MPE
- Numerical results


## Numerical Results



Logistic regression with $\ell_{2}$ regularizartion, on Madelon Dataset (500 features, 2000 data points), solved using several algorithms. The penalty parameter has been set to $10^{2}$ in order to have a condition number equal to $1.2 \times 10^{9}$.

## Numerical Results



Logistic regression on Sido0 Dataset (4932 features, 12678 data points). Penalty parameter $\tau=10^{2}$, so the condition number is equal to $1.5 \times 10^{5}$.

## Numerical Results



Logistic regression on Madelon UCI Dataset, solved using the gradient method, Nesterov's method and AMPE (i.e. RMPE with $\lambda=0$ ). The condition number is equal to $1.2 \times 10^{9}$. We see that without regularization, AMPE becomes unstable as $\left\|\left(\tilde{U}^{T} \tilde{U}\right)^{-1}\right\|_{2}$ gets too large.

## Conclusion

## Postprocessing works.

- Simple postprocessing step.
- Marginal complexity, can be performed in parallel.
- Significant convergence speedup over optimal methods.
- Adaptive. Does not need knowledge of smoothness parameters.

Work in progress. . .

- Extrapolating accelerated methods.
- Constrained problems.
- Better handling of smooth functions.


## Open problems

- Regularization. How do we account for the fact that we are estimating the limit of a VAR sequence with a fixed point?
- The VAR matrix $A$ is formed implicitly, but we have some information on its spectrum through smoothness.
- Explicit bounds on the regularized Chebyshev problem,

$$
S(k, \alpha) \triangleq \min _{\left\{q \in \mathbb{R}_{k}[x]: q(1)=1\right\}}\left\{\max _{x \in[0, \sigma]}((1-x) q(x))^{2}+\alpha\|q\|_{2}^{2}\right\} .
$$

Preprint on ArXiv:1606.04133 and NIPS 2016.

## References

Alexander Craig Aitken. On Bernoulli's numerical solution of algebraic equations. Proceedings of the Royal Society of Edinburgh, 46:289-305, 1927.

C Brezinski. Accélération de la convergence en analyse numérique. Lecture notes in mathematics (ISSN 0075-8434, (584), 1977.
RP Eddy. Extrapolating to the limit of a vector sequence. Information linkage between applied mathematics and industry, pages 387-396, 1979.
M. Hardt. The zen of gradient descent. Mimeo, 2013.

Georg Heinig and Karla Rost. Fast algorithms for Toeplitz and Hankel matrices. Linear Algebra and its Applications, 435(1):1-59, 2011.
M Mešina. Convergence acceleration for the iterative solution of the equations $x=a x+f$. Computer Methods in Applied Mechanics and Engineering, 10(2):165-173, 1977.
Y. Nesterov. A method of solving a convex programming problem with convergence rate $O\left(1 / k^{2}\right)$. Soviet Mathematics Doklady, 27(2): 372-376, 1983.
D. Scieur, A. d'Aspremont, and F. Bach. Regularized Nonlinear Acceleration. NIPS, 2016.

Avram Sidi, William F Ford, and David A Smith. Acceleration of convergence of vector sequences. SIAM Journal on Numerical Analysis, 23 (1):178-196, 1986.

David A Smith, William F Ford, and Avram Sidi. Extrapolation methods for vector sequences. SIAM review, 29(2):199-233, 1987.
Evgenij E Tyrtyshnikov. How bad are Hankel matrices? Numerische Mathematik, 67(2):261-269, 1994.
Peter Wynn. On a device for computing the $e_{m}\left(s_{n}\right)$ transformation. Mathematical Tables and Other Aids to Computation, 10(54):91-96, 1956.

