
A Fast Anderson-Chebyshev Acceleration for Nonlinear Optimization

Zhize Li

King Abdullah University of Science and Technology

Jian Li

Tsinghua University

Abstract

Anderson acceleration (or Anderson mixing) is an efficient acceleration method for fixed point iterations $x_{t+1} = G(x_t)$, e.g., gradient descent can be viewed as iteratively applying the operation $G(x) \triangleq x - \alpha \nabla f(x)$. It is known that Anderson acceleration is quite efficient in practice and can be viewed as an extension of Krylov subspace methods for nonlinear problems. In this paper, we show that Anderson acceleration with Chebyshev polynomial can achieve the optimal convergence rate $O(\sqrt{\kappa} \ln \frac{1}{\epsilon})$, which improves the previous result $O(\kappa \ln \frac{1}{\epsilon})$ provided by (Toth and Kelley, 2015) for quadratic functions. Moreover, we provide a convergence analysis for minimizing general nonlinear problems. Besides, if the hyperparameters (e.g., the Lipschitz smooth parameter L) are not available, we propose a *guessing algorithm* for guessing them dynamically and also prove a similar convergence rate. Finally, the experimental results demonstrate that the proposed Anderson-Chebyshev acceleration method converges significantly faster than other algorithms, e.g., vanilla gradient descent (GD), Nesterov’s Accelerated GD. Also, these algorithms combined with the proposed guessing algorithm (guessing the hyperparameters dynamically) achieve much better performance.

1 Introduction

Machine learning problems are usually modeled as optimization problems, ranging from convex optimiza-

tion to highly nonconvex optimization such as deep neural networks, e.g., (Nesterov, 2004; Bubeck, 2015; LeCun et al., 2015; Lei et al., 2017; Li and Li, 2018; Fang et al., 2018; Zhou et al., 2018; Li et al., 2019; Ge et al., 2019; Li, 2019). To solve an optimization problem $\min_x f(x)$, the classical method is gradient descent, i.e., $x_{t+1} = x_t - \alpha_t \nabla f(x_t)$. There exist several techniques to accelerate the standard gradient descent, e.g., momentum (Nesterov, 2004; Allen-Zhu, 2017; Lan and Zhou, 2018; Lan et al., 2019). There are also various vector sequence acceleration methods developed in the numerical analysis literature, e.g., (Brezinski, 2000; Sidi et al., 1986; Smith et al., 1987; Brezinski and Redivo Zaglia, 1991; Brezinski et al., 2018). Roughly speaking, if a vector sequence converges very slowly to its limit, then one may apply such methods to accelerate the convergence of this sequence. Taking gradient descent as an example, the vector sequence are generated by $x_{t+1} = G(x_t) \triangleq x_t - \alpha_t \nabla f(x_t)$, where the limit is the fixed-point $G(x^*) = x^*$ (i.e. $\nabla f(x^*) = 0$). One notable advantage of such acceleration methods is that they usually do not require to know how the vector sequence is actually generated. Thus the applicability of those methods is very wide.

Recently, Scieur et al. (2016) used the minimal polynomial extrapolation (MPE) method (Smith et al., 1987) for convergence acceleration. This is a nice example of using sequence acceleration methods to optimization problems. In this paper, we are interested in another classical sequence acceleration method called *Anderson acceleration* (or *Anderson mixing*), which was proposed by Anderson in 1965 (Anderson, 1965). The method is known to be quite efficient in a variety of applications (Capehart, 1989; Pratapa et al., 2016; Higham and Strabić, 2016; Loffeld and Woodward, 2016). The idea of Anderson acceleration is to maintain m recent iterations for determining the next iteration point, where m is a parameter (typically a very small constant). Thus, it can be viewed as an extension of the existing momentum methods which usually use the last and current points to determine the next iteration point. Anderson acceleration with slight modifications is described in Algorithm 1.

Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics (AISTATS) 2020, Palermo, Italy. PMLR: Volume 108. Copyright 2020 by the author(s).

Algorithm 1: Anderson Acceleration(m)

1 **input:** x_0, T, λ, β_t
2 Define $G(x) \triangleq x + F \triangleq x - \lambda \nabla f(x)$;
3 $x_1 = G(x_0)$, $F_0 = G(x_0) - x_0$;
4 **for** $t = 1, 2, \dots, T$ **do**
5 $m_t = \min\{m, t\}$;
6 $F_t \triangleq G(x_t) - x_t$;
7 Solve $\min_{\alpha^t = (\alpha_0^t, \dots, \alpha_{m_t}^t)^T} \|\sum_{i=0}^{m_t} \alpha_i^t F_{t-i}\|_2$
 subject to $\sum_{i=0}^{m_t} \alpha_i^t = 1$;
8 $x_{t+1} =$
 $(1 - \beta_t) \sum_{i=0}^{m_t} \alpha_i^t x_{t-i} + \beta_t \sum_{i=0}^{m_t} \alpha_i^t G(x_{t-i})$;
9 **return** x_T

Note that the step in Line 7 of Algorithm 1 can be transformed to an equivalent unconstrained least-squares problem:

$$\min_{(\alpha_1^t, \dots, \alpha_{m_t}^t)^T} \left\| F_t - \sum_{i=1}^{m_t} \alpha_i^t (F_t - F_{t-i}) \right\|_2, \quad (1)$$

then let $\alpha_0^t = 1 - \sum_{i=1}^{m_t} \alpha_i^t$. Using QR decomposition, (1) can be solved in time $2m_t^2 d$, where d is the dimension. Moreover, the QR decomposition of (1) at iteration t can be efficiently obtained from that of at iteration $t - 1$ in $O(m_t d)$ (see, e.g. (Golub and Van Loan, 1996)). The constant $m_t \leq m$ is usually very small. We use $m = 3$ and 5 for the numerical experiments in Section 5. Hence, each iteration of Anderson acceleration can be implemented quite efficiently.

Many studies showed the relations between Anderson acceleration and other optimization methods. In particular, for the quadratic case (linear problems), Walker and Ni (2011) showed that it is related to the well-known Krylov subspace method GMRES (generalized minimal residual algorithm) (Saad and Schultz, 1986). Furthermore, Potra and Engler (2013) showed that GMRES is equivalent to Anderson acceleration with any mixing parameters under $m = \infty$ (see Line 5 of Algorithm 1) for linear problems. Concretely, Toth and Kelley (2015) proved the first linear convergence rate $O(\kappa \ln \frac{1}{\epsilon})$ for linear problems with fixed parameter β , where κ is the condition number. Besides, Eyert (1996), and Fang and Saad (2009) showed that Anderson acceleration is related to the multiseccant quasi-Newton methods (more concretely, the generalized Broyden's second method). Despite the above results, the convergence results for this efficient method are still limited (especially for general nonlinear problems and the case where m is small). In this paper, we analyze the convergence for small m which is the typical case in practice and also provide the convergence analysis for general nonlinear problems.

1.1 Our Contributions

There has been a growing number of applications of Anderson acceleration method (Pratapa et al., 2016; Higham and Strabić, 2016; Loffeld and Woodward, 2016; Scieur et al., 2018). Towards a better understanding of this efficient method, we make the following technical contributions:

1. We prove the optimal $O(\sqrt{\kappa} \ln \frac{1}{\epsilon})$ convergence rate of the proposed Anderson-Chebyshev acceleration (i.e., Anderson acceleration with Chebyshev polynomial) for minimizing quadratic functions (see Theorem 1). Our result improves the previous result $O(\kappa \ln \frac{1}{\epsilon})$ given by (Toth and Kelley, 2015) and matches the lower bound $\Omega(\sqrt{\kappa} \ln \frac{1}{\epsilon})$ provided by (Nesterov, 2004). Note that for ill-conditioned problems, the condition number κ can be very large.
2. Then, we prove the linear-quadratic convergence of Anderson acceleration for minimizing general nonlinear problems under some standard assumptions (see Theorem 2). Compared with Newton-like methods, it is more attractive since it does not require to compute (or approximate) Hessians, or Hessian-vector products.
3. Besides, we propose a *guessing algorithm* (Algorithm 2) for the case when the hyperparameters (e.g., μ, L) are not available. We prove that it achieves a similar convergence rate $O(\sqrt{\kappa} \ln \frac{1}{\epsilon} + \sqrt{\kappa} (\ln \kappa \ln B)^2)$ (see Theorem 3). This guessing algorithm can also be combined with other algorithms, e.g., Gradient Descent (GD), Nesterov's Accelerated GD (NAGD). The experimental results (see Section 5.1) show that these algorithms combined with the proposed guessing algorithm achieve much better performance.
4. Finally, the experimental results on the real-world UCI datasets and synthetic datasets demonstrate that Anderson acceleration methods converge significantly faster than other algorithms (see Section 5). Combined with our theoretical results, the experiments validate that Anderson acceleration methods (especially Anderson-Chebyshev acceleration) are efficient both in theory and practice.

1.2 Related Work

As aforementioned, Anderson acceleration can be viewed as the extension of the momentum methods (e.g., NAGD) and the potential extension of Krylov subspace methods (e.g., GMRES) for nonlinear problems. In particular, GD is the special case of Anderson

acceleration with $m = 0$, and to some extent NAGD can be viewed as $m = 1$. We also review the equivalence of GMRES and Anderson acceleration without truncation (i.e., $m = \infty$) in Appendix A¹. Besides, Eyert (1996), and Fang and Saad (2009) showed that Anderson acceleration is related to the multiseccant quasi-Newton methods. Note that Anderson acceleration has the advantage over the Newton-like methods since it does not require the computation of Hessians or approximation of Hessians or Hessian-vector products.

There are many sequence acceleration methods in the numerical analysis literatures. In particular, the well-known Aitken's Δ^2 process (Aitken, 1926) accelerated the convergence of a sequence that is converging linearly. Shanks generalized the Aitken extrapolation which was known as Shanks transformation (Shanks, 1955). Recently, Brezinski et al. (2018) proposed a general framework for Shanks sequence transformations which includes many vector sequence acceleration methods. One fundamental difference between Anderson acceleration and other sequence acceleration methods (such as MPE, RRE (reduced rank extrapolation) (Sidi et al., 1986; Smith et al., 1987), etc.) is that Anderson acceleration is a fully dynamic method (Capehart, 1989). Here *dynamic* means all iterations are in the same sequence, and it does not require to restart the procedure. It can be seen from Algorithm 1 that all iterations are applied to the same sequence $\{x_t\}$. In fact, in Capehart's PhD thesis (Capehart, 1989), several experiments were conducted to demonstrate the superior performance of Anderson acceleration over other semi-dynamic methods such as MPE, RRE (semi-dynamic means that the algorithm maintains more than one sequences or needs to restart several times). More recently, Anderson acceleration with different variants and/or under different assumptions are widely studied (see e.g., (Zhang et al., 2018; Evans et al., 2018; Scieur et al., 2019)).

2 The Quadratic Case

In this section, we consider the problem of minimizing a quadratic function (also called least squares, or ridge regression (Boyd and Vandenberghe, 2004; Hoerl and Kennard, 1970)). The formulation of the problem is

$$\min_{x \in \mathbb{R}^d} f(x) = \frac{1}{2} x^T A x - b^T x, \quad (2)$$

where $\mu I_d \preceq \nabla^2 f = A \preceq L I_d$. Note that μ and L are usually called the strongly convex parameter and Lipschitz continuous gradient parameter, respectively (e.g. (Nesterov, 2004; Allen-Zhu, 2017; Lan et al.,

¹All appendices are provided in the Supplementary Material.

2019)). There are many algorithms for optimizing this type of functions. See e.g. (Bubeck, 2015) for more details. We analyze the problem of minimizing a more general function $f(x)$ in the next Section 3.

We prove that Anderson acceleration with Chebyshev polynomial parameters $\{\beta_t\}$ achieves the optimal convergence rate, i.e., it obtains an ϵ -approximate solution using $O(\sqrt{\kappa} \ln \frac{1}{\epsilon})$ iterations. The convergence result is stated in the following Theorem 1.

Theorem 1 *The Anderson-Chebyshev acceleration method achieves the optimal convergence rate $O(\sqrt{\kappa} \ln \frac{1}{\epsilon})$ for obtaining an ϵ -approximate solution of problem (2) for any $0 \leq m \leq k$, where $\kappa = L/\mu$ is the condition number, k is defined in Definition 1 and this method combines Anderson acceleration (Algorithm 1) with the Chebyshev polynomial parameters $\beta_t = 1 / (\frac{L+\mu}{2} + \frac{L-\mu}{2} \cos(\frac{(2t-1)\pi}{2T}))$, for $t = 1, 2, \dots, T$.*

Remark: In this quadratic case, we mention that Toth and Kelley (2015) proved the first convergence rate $O(\kappa \ln \frac{1}{\epsilon})$ for fixed parameter β . Here we use the Chebyshev polynomials to improve the result to the optimal $O(\sqrt{\kappa} \ln \frac{1}{\epsilon})$ which matches the lower bound $\Omega(\sqrt{\kappa} \ln \frac{1}{\epsilon})$. Note that for ill-conditioned problems, the condition number κ can be very large. Also note that in practice the constant m is usually very small. Particularly, $m = 3$ has already achieved a remarkable performance from our experimental results (see Figures 2-5 in Section 5).

Before proving Theorem 1, we first define k and then briefly review some properties of the Chebyshev polynomials. We refer to (Rivlin, 1974; Olshanskii and Tyrtysnikov, 2014; Hageman and Young, 2012) for more details of Chebyshev polynomials.

Definition 1 *Let v_i 's be the unit eigenvectors of A , where A is defined in (2). Consider a unit vector $c \triangleq \sum_{i=1}^d c_i v_i$ and let $c' \triangleq \text{Proj}_{B_k^\perp} c = \sum_{i=1}^d c'_i v_i$, where $\text{Proj}_{B_k^\perp}$ denotes the projection to the orthogonal complement of the column space of $B_k \triangleq A[x_{t-k} - x_t, \dots, x_{t-1} - x_t] \in \mathbb{R}^{d \times k}$. Define k to be the maximum integer such that $c'_i \leq (1 + \frac{1}{\sqrt{\kappa+1}}) c_i$ for any $i \in [d]$.*

Obviously, $k \geq 0$ since $c' = c$ due to $B_0 = 0$ and $\text{Proj}_{B_0^\perp} = I$.

Now we review the Chebyshev polynomials. The *Chebyshev polynomials* are polynomials $P_k(x)$, where $k \geq 0$, $\deg(P_k) = k$, which is defined by the recursive relation:

$$\begin{aligned} P_0(x) &= 1, \\ P_1(x) &= x, \\ P_{k+1}(x) &= 2xP_k(x) - P_{k-1}(x). \end{aligned} \quad (3)$$

The key property is that $P_k(x)$ has minimal deviation from 0 on $[-1, 1]$ among all polynomials Q_k with $\deg(Q_k) = k$ and coefficient $\alpha_k = 2^{k-1}$ for the largest degree term x^k , i.e.,

$$\max_{x \in [-1, 1]} |P_k(x)| \leq \max_{x \in [-1, 1]} |Q_k(x)| \quad \text{for all } Q_k. \quad (4)$$

In particular, for $|x| \leq 1$, Chebyshev polynomials can be written in an equivalent way:

$$P_k(x) = \cos(k \arccos x). \quad (5)$$

In our proof, we use this equivalent form (5) instead of (3). The equivalence can be verified as follows:

$$\begin{aligned} P_k(x) &= 2x \cos((k-1) \arccos x) - \cos((k-2) \arccos x) \\ &= 2 \cos \theta \cos((k-1)\theta) - \cos((k-2)\theta) \end{aligned} \quad (6)$$

$$\begin{aligned} &= \cos(k\theta) + \cos((k-2)\theta) - \cos((k-2)\theta) \\ &= \cos(k \arccos x), \end{aligned} \quad (7)$$

where (6) and (7) use the transformation $x = \cos \theta$ due to $|x| \leq 1$. According to (5), $\max_{x \in [-1, 1]} |P_k(x)| = 1$ and the k roots of P_k are as follows:

$$x_i = \cos\left(\frac{(2i-1)\pi}{2k}\right), \quad i = 1, 2, \dots, k. \quad (8)$$

To demonstrate it more clearly, we provide an example for $P_4(x)$ (W-shape curve) in Figure 1. Since $k = 4$ in this polynomial $P_4(x)$, the first root $x_1 = \cos\left(\frac{(2i-1)\pi}{2k}\right) = \cos\left(\frac{\pi}{8}\right) \approx 0.92$. The remaining three roots for $P_4(x)$ can be easily computed too.




Figure 1: The Chebyshev polynomial $P_4(x)$

Proof of Theorem 1. For iteration $t + 1$, the residual $F_{t+1} \triangleq -\lambda \nabla f(x_{t+1}) = -(Ax_{t+1} - b)$ (let $\lambda = 1$) can

be deduced as follows:

$$\begin{aligned} F_{t+1} &= b - Ax_{t+1} \\ &= b - A \left[(1 - \beta_t) \sum_{i=0}^{m_t} \alpha_i^t x_{t-i} + \beta_t \sum_{i=0}^{m_t} \alpha_i^t G(x_{t-i}) \right] \\ &= b - A \left[\sum_{i=0}^{m_t} \alpha_i^t x_{t-i} + \beta_t \sum_{i=0}^{m_t} \alpha_i^t (b - Ax_{t-i}) \right] \end{aligned} \quad (9)$$

$$\begin{aligned} &= b - \beta_t Ab - A \left[\sum_{i=0}^{m_t} \alpha_i^t ((I - \beta_t A)x_{t-i}) \right] \\ &= (I - \beta_t A) \sum_{i=0}^{m_t} \alpha_i^t (b - Ax_{t-i}) \\ &= (I - \beta_t A) \sum_{i=0}^{m_t} \alpha_i^t F_{t-i}, \end{aligned} \quad (10)$$

where (9) uses $G(x_t) = x_t + F_t$.

To bound $\|F_{t+1}\|_2$ (i.e., $\|\nabla f(x_{t+1})\|_2$), we first obtain the following lemma by using Singular Value Decomposition (SVD) to solve the least squares problem (1) and then using several transformations. We defer the proof of Lemma 1 to Appendix B.2.

Lemma 1 *Let $F_1 = b - Ax_1$ and $F_{t+1} = b - Ax_{t+1}$, then*

$$\|F_{t+1}\|_2 / \|F_1\|_2 \leq \sqrt{2 \min_{\beta} \max_{\lambda \in [\mu, L]} |H_t(\lambda)|} \quad (11)$$

where $H_t(\lambda) = (1 - \beta_t \lambda) \cdots (1 - \beta_1 \lambda)$ is a degree t polynomial.

According to Lemma 1, to bound $\|F_{t+1}\|_2$, it is sufficient to bound the right-hand-side (RHS) of (11) (i.e., $\min_{\beta} \max_{\lambda \in [\mu, L]} |H_t(\lambda)|$). So we want to choose parameter β in order to make $\max_{\lambda \in [\mu, L]} |H_t(\lambda)|$ as small as possible. According to (4) (the minimal deviation property of standard Chebyshev polynomials), hence a natural idea is to choose β such that $H_t(\lambda) = (1 - \beta_t \lambda) \cdots (1 - \beta_1 \lambda)$ is a kind of modified Chebyshev polynomials. In order to do this, we first transform $[\mu, L]$ into $[-1, 1]$, i.e., let $\lambda = \frac{L+\mu}{2} + \frac{L-\mu}{2}x$, where $x \in [-1, 1]$. Also note that polynomial $H_t(\lambda) = (1 - \beta_t \lambda) \cdots (1 - \beta_1 \lambda)$ has (only) one constraint, i.e., $H_t(0) = 1$. Thus we choose β such that

$$\begin{aligned} H_t(\lambda) &= P_t\left(\frac{2\lambda - (L + \mu)}{L - \mu}\right) / P_t\left(-\frac{L + \mu}{L - \mu}\right) \\ &= P_t(x) / P_t\left(-\frac{L + \mu}{L - \mu}\right), \end{aligned} \quad (12)$$

where $P_t(\cdot)$ is the standard Chebyshev polynomials.

Now, the RHS of (11) can be bounded as follows:

$$\begin{aligned} & \min_{\beta} \max_{\lambda \in [\mu, L]} |H_t(\lambda)| \\ & \leq \max_{x \in [-1, 1]} \left| P_t(x) / P_t\left(-\frac{L+\mu}{L-\mu}\right) \right| \end{aligned} \quad (13)$$

$$\leq 1 / \left| P_t\left(-\frac{L+\mu}{L-\mu}\right) \right|, \quad (14)$$

where (13) uses (12), and (14) uses $\max_{x \in [-1, 1]} |P_t(x)| = 1$ (see (5)). According to (8), it is not hard to see that $H_t(\lambda)$ is defined by the mixing parameters $\beta_i = 1 / \left(\frac{L+\mu}{2} + \frac{L-\mu}{2} \cos\left(\frac{(2i-1)\pi}{2t}\right) \right)$ according to $\lambda = \frac{L+\mu}{2} + \frac{L-\mu}{2}x$, where $i = 1, 2, \dots, t$. Note that the roots of standard Chebyshev polynomials (i.e., (8)) can be found from many textbooks, e.g., Section 1.2 of (Rivlin, 1974). Now, we only need to bound $|P_t(-\frac{L+\mu}{L-\mu})|$. First, we need to transform the form (5) of Chebyshev polynomials $P_t(x)$ as follows:

$$\begin{aligned} P_t(x) &= \cos(t \arccos x) \\ &= \cos(t\theta) \quad \text{Define } x \triangleq \cos \theta \\ &= (e^{i\theta t} + e^{-i\theta t}) / 2 \\ &= ((\cos \theta + i \sin \theta)^t + (\cos \theta - i \sin \theta)^t) / 2 \\ &= \left((x + \sqrt{x^2 - 1})^t + (x - \sqrt{x^2 - 1})^t \right) / 2. \end{aligned}$$

Let $x = -\frac{L+\mu}{L-\mu}$, we get $\sqrt{x^2 - 1} = \sqrt{\frac{(L+\mu)^2 - (L-\mu)^2}{(L-\mu)^2}} = \sqrt{\frac{4L\mu}{(L-\mu)^2}} = \frac{2\sqrt{L\mu}}{L-\mu}$. So we have

$$\begin{aligned} \left| P_t\left(-\frac{L+\mu}{L-\mu}\right) \right| &\geq \frac{1}{2} \left(\frac{L+\mu}{L-\mu} + \frac{2\sqrt{L\mu}}{L-\mu} \right)^t \\ &\geq \frac{1}{2} \left(\frac{\sqrt{L} + \sqrt{\mu}}{\sqrt{L} - \sqrt{\mu}} \right)^t. \end{aligned} \quad (15)$$

Now, the RHS of (11) can be bounded as

$$\begin{aligned} \sqrt{2 \min_{\beta} \max_{\lambda \in [\mu, L]} |H_t(\lambda)|} &\leq \sqrt{2 / \left| P_t\left(-\frac{L+\mu}{L-\mu}\right) \right|} \quad (16) \\ &\leq 2 \left(\frac{\sqrt{L} - \sqrt{\mu}}{\sqrt{L} + \sqrt{\mu}} \right)^{t/2}, \end{aligned} \quad (17)$$

where (16) follows from (14), and (17) follows from (15). Then, according to (11), the gradient norm is bounded as $\|\nabla f(x_{t+1})\|_2 = \|F_{t+1}\|_2 \leq 2 \left(\frac{\sqrt{L} - \sqrt{\mu}}{\sqrt{L} + \sqrt{\mu}} \right)^{t/2} \|F_1\|_2 = 2 \left(\frac{\sqrt{\kappa} - 1}{\sqrt{\kappa} + 1} \right)^{t/2} \|\nabla f(x_1)\|_2$, where $\kappa = L/\mu$. Note that if the number of iterations $t = (\sqrt{\kappa} + 1) \ln \frac{1}{\epsilon}$, then

$$\left(\frac{\sqrt{\kappa} - 1}{\sqrt{\kappa} + 1} \right)^{t/2} = \left(1 - \frac{2}{\sqrt{\kappa} + 1} \right)^{t/2} \leq \epsilon.$$

Thus the Anderson-Chebyshev acceleration method achieves the optimal convergence rate $O(\sqrt{\kappa} \ln \frac{1}{\epsilon})$ for obtaining an ϵ -approximate solution. \square

3 The General Case

In this section, we analyze the Anderson Acceleration (Algorithm 1) in the general nonlinear case:

$$\min_{x \in \mathbb{R}^d} f(x). \quad (18)$$

We prove that Anderson acceleration method achieves the linear-quadratic convergence rate under the following standard Assumptions 1 and 2, where $\|\cdot\|$ denotes the Euclidean norm. Let \mathcal{B}_t denote the small matrix of the least-square problem in Line 7 of Algorithm 1, i.e., $\mathcal{B}_t \triangleq [F_t - F_{t-1}, \dots, F_t - F_{t-m}] \in \mathbb{R}^{d \times m}$ (see problem (1)). Then, we define its condition number $\kappa_t \triangleq \|\nabla f(x_t)\| / \tilde{\mu}_t$ and $\tilde{\kappa} \triangleq \max_t \{\kappa_t\}$, where $\tilde{\mu}_t$ denotes the least non-zero singular value of \mathcal{B}_t .

Assumption 1 *The Hessian $\nabla^2 f$ satisfies $\mu \leq \|\nabla^2 f\| \leq L$, where $0 \leq \mu \leq L$.*

Assumption 2 *The Hessian $\nabla^2 f$ is γ -Lipschitz continuous, i.e.,*

$$\|\nabla^2 f(x) - \nabla^2 f(y)\| \leq \gamma \|x - y\|. \quad (19)$$

Theorem 2 *Suppose Assumption 1 and 2 hold. Let step-size $\lambda = \frac{2}{L+\mu}$. The convergence rate of Anderson Acceleration(m) (Algorithm 1) is linear-quadratic for problem (18), i.e.,*

$$\|\nabla f(x_{t+1})\| \leq c_1 \Delta_t^2 + c_2 \Delta_t \|\nabla f(x_t)\| + (1 - c_3) \|\nabla f(x_t)\|, \quad (20)$$

where $c_1 = \frac{3\tilde{\kappa}^2 \gamma m}{(L+\mu)^2}$, $c_2 = \frac{2\tilde{\kappa} \beta_t \gamma \sqrt{m}}{(L+\mu)^2}$, $c_3 = \beta_t \frac{2\mu}{L+\mu}$ and $\Delta_t \triangleq \max_{i \in [m]} \|x_t - x_{t-i}\|$.

Remark:

1. The constant $m \geq 0$ is usually very small. Particularly, we use $m = 3$ and 5 for the numerical experiments in Section 5. Hence Δ_t is very small and also decreases as the algorithm converges.
2. Besides, one can also use $\|\nabla f(x_t)\|$ instead of Δ_t in (20) according to the property of f (Assumption 1), i.e., $\mu \|x_t - x^*\| \leq \|\nabla f(x_t) - \nabla f(x^*)\| = \|\nabla f(x_t)\|$, and $\|x_t - x_{t-i}\| = \|x_t - x^* + x^* - x_{t-i}\| \leq \|x_t - x^*\| + \|x_{t-i} - x^*\|$.
3. Note that the first two terms in RHS of (20) converge quadratically and the last term converges linearly. Due to the fully dynamic property of Anderson acceleration as we discussed in Section 1.2, it turns out the exact convergence rate of Anderson acceleration in the general case is not easy to obtain. But we note that the convergence rate is roughly linear, i.e., $O(\frac{1}{c_3} \log \frac{1}{\epsilon})$ since the first two quadratic terms converge much faster

than the last linear term in some neighborhood of optimum. In particular, if f is a quadratic function, then $\gamma = 0$ (Assumption 2) and thus $c_1 = c_2 = 0$ in (20). Only the last linear term remained, thus it converges linearly (see the following corollary).

Corollary 1 *If f is a quadratic function, let step-size $\lambda = \frac{2}{L+\mu}$ and $\beta_t = 1$. Then the convergence rate of Anderson Acceleration is linear, i.e., $O(\kappa \ln \frac{1}{\epsilon})$, where $\kappa = L/\mu$ is the condition number.*

Note that this corollary recovers the previous result (i.e., $O(\kappa \ln \frac{1}{\epsilon})$) obtained by (Toth and Kelley, 2015), and we use *Chebyshev polynomial* to improve this result to the optimal convergence rate $O(\sqrt{\kappa} \ln \frac{1}{\epsilon})$ in our previous Section 2 (see Theorem 1). Concretely, we transfer the weight of step-size λ to the parameters β_t 's and use Chebyshev polynomial parameters β_t 's in our Theorem 1 instead of using fixed parameter $\beta \equiv 1$.

Now, we provide a proof sketch for Theorem 2. The detailed proof can be found in Appendix B.1.

Proof Sketch of Theorem 2. Consider the iteration $t+1$, we have $F_t = -\frac{2}{L+\mu} \nabla f(x_t)$ according to $\lambda = \frac{2}{L+\mu}$. First, we need to demonstrate several useful forms of x_{t+1} as follows:

$$\begin{aligned}
x_{t+1} &= (1 - \beta_t) \sum_{i=0}^{m_t} \alpha_i^t x_{t-i} + \beta_t \sum_{i=0}^{m_t} \alpha_i^t G(x_{t-i}) \\
&= \sum_{i=0}^{m_t} \alpha_i^t x_{t-i} + \beta_t \sum_{i=0}^{m_t} \alpha_i^t (G(x_{t-i}) - x_{t-i}) \\
&= \sum_{i=0}^{m_t} \alpha_i^t x_{t-i} + \beta_t \sum_{i=0}^{m_t} \alpha_i^t F_{t-i} \quad (21) \\
&= x_t - \sum_{i=1}^{m_t} \alpha_i^t (x_t - x_{t-i}) \\
&\quad + \beta_t \left(F_t - \sum_{i=1}^{m_t} \alpha_i^t (F_t - F_{t-i}) \right), \quad (22)
\end{aligned}$$

where (21) holds due to the definition $G_t = G(x_t) = x_t + F_t$, and (22) holds since $\sum_{i=0}^{m_t} \alpha_i^t = 1$.

Then, to bound $\|F_{t+1}\|_2$ (i.e., $\|\nabla f(x_{t+1})\|_2$), we deduce F_{t+1} as follows:

$$\begin{aligned}
F_{t+1} &= G_{t+1} - x_{t+1} \\
&= G_{t+1} - \sum_{i=0}^{m_t} \alpha_i^t x_{t-i} - \beta_t \sum_{i=0}^{m_t} \alpha_i^t F_{t-i} \\
&= G_{t+1} - \sum_{i=0}^{m_t} \alpha_i^t (G_{t-i} - F_{t-i}) - \beta_t \sum_{i=0}^{m_t} \alpha_i^t F_{t-i} \\
&= G_{t+1} - \sum_{i=0}^{m_t} \alpha_i^t G_{t-i} + (1 - \beta_t) \mathcal{F}, \quad (23)
\end{aligned}$$

where (23) uses the definition $\mathcal{F} \triangleq \sum_{i=0}^{m_t} \alpha_i^t F_{t-i}$. Now, we bound the first two terms of (23) as follows:

$$\begin{aligned}
G_{t+1} - \sum_{i=0}^{m_t} \alpha_i^t G_{t-i} &= G_{t+1} - \left(G_t - \sum_{i=1}^{m_t} \alpha_i^t (G_t - G_{t-i}) \right) \\
&= \int_0^1 G'(x_t + u(x_{t+1} - x_t))(x_{t+1} - x_t) du \\
&\quad - \sum_{i=1}^{m_t} \alpha_i^t \int_0^1 G'(x_t + u(x_{t-i} - x_t))(x_{t-i} - x_t) du \\
&= \sum_{i=1}^{m_t} \alpha_i^t \int_0^1 G'(x_t + u(x_{t+1} - x_t))(x_{t-i} - x_t) du \\
&\quad + \int_0^1 G'(x_t + u(x_{t+1} - x_t)) \beta_t \mathcal{F} du \\
&\quad - \sum_{i=1}^{m_t} \alpha_i^t \int_0^1 G'(x_t + u(x_{t-i} - x_t))(x_{t-i} - x_t) du, \quad (24)
\end{aligned}$$

where (24) is obtained by using (22) to replace x_{t+1} . To bound (24), we use Assumptions 1, 2, and the equation

$$G'_t = I + F'_t = I - \frac{2}{L + \mu} \nabla^2 f(x_t).$$

After some non-trivial calculations (details can be found in Appendix B.1), we obtain

$$\begin{aligned}
\|F_{t+1}\| &\leq \frac{\gamma(m\|\alpha\|^2 + \sqrt{m}\|\alpha\|)\Delta_t^2}{L + \mu} + \frac{\gamma\sqrt{m}\|\alpha\|\beta_t\Delta_t\|\mathcal{F}\|}{L + \mu} \\
&\quad + \left(1 - \frac{2\mu}{L + \mu}\beta_t\right)\|\mathcal{F}\|,
\end{aligned}$$

where $\|\alpha\|$ denotes the Euclidean norm of $\alpha = (\alpha_1^t, \dots, \alpha_{m_t}^t)^T$. Then, according to the problem (1) and the definition of $\mathcal{F} \triangleq \sum_{i=0}^{m_t} \alpha_i^t F_{t-i}$, we have $\|\mathcal{F}\| \leq \|F_t\|$. Finally, we bound $\|\alpha\| \leq \frac{2\bar{\kappa}}{L+\mu}$ using QR decomposition of problem (1) and recall $F_t = -\frac{2}{L+\mu} \nabla f(x_t)$ to finish the proof of Theorem 2. \square

4 Guessing Algorithm

In this section, we provide a *Guessing Algorithm* (described in Algorithm 2) which guesses the parameters (e.g., μ, L) dynamically. Intuitively, we guess the parameter μ and the condition number κ in a doubling way. Note that in general these parameters are not available, since the time for computing these parameters is almost the same as (or even longer than) solving the original problem. Also note that the condition in Line 14 of Algorithm 2 depends on the algorithm used in Line 12.

Algorithm 2: Guessing Algorithm

```

1 input:  $x_0, T, \delta, B$ 
2 Let  $t = 0$ ;
3 for  $i = 1, 2, \dots$  do
4      $\kappa_i = e^{i+2}$ ;
5     for  $j = 1, 2, \dots, \ln B$  do
6          $\mu_i = e^j \delta, L_i = \mu_i \kappa_i, t_i = 1$ ;
7         do
8              $t_i = \lfloor et_i \rfloor$ ;
9             if  $t + t_i > T$  then
10                 break;
11              $x_{t-1} = x_t$ ;
12              $x = \text{Anderson Acceleration}(x_t, t_i, \mu_i, L_i)$ 
13                 //can be replaced by other algorithms;
14              $t = t + t_i, x_t = x$ ;
15         while  $\frac{\|\nabla f(x_t)\|}{\|\nabla f(x_{t-1})\|} \leq 2 \left( \frac{\sqrt{\kappa_i} - 1}{\sqrt{\kappa_i} + 1} \right)^{t_i}$ ;
16         if  $\|\nabla f(x_t)\| > \|\nabla f(x_{t-1})\|$  then
17              $x_t = x_{t-1}$ ;
18 return  $x_t$ 
    
```

The convergence result of our Algorithm 2 is stated in the following Theorem 3. The detailed proof is deferred to Appendix B.3. Note that we only prove the quadratic case for Theorem 3, but it is possible to extend it to the general case.

Theorem 3 *Without knowing the parameters μ and L , Algorithm 2 achieves $O(\sqrt{\kappa} \ln \frac{1}{\epsilon} + \sqrt{\kappa} (\ln \kappa \ln B)^2)$ convergence rate for obtaining an ϵ -approximate solution of problem (2), where $\kappa = L/\mu$, and B can be any number as long as the eigenvalue spectrum belongs to $[\delta, B\delta]$.*

Remark: We provide a simple example to show why this guessing algorithm is useful. Note that algorithms usually need the (exact) parameters μ and L to set the step size. Without knowing the exact values μ and L , one needs to approximate these parameters once at the beginning. Let $\mu' = \frac{1}{c_1} \mu$ and $L' = c_2 L$ denote the approximated values, where $c_1, c_2 \geq 1$. Without guessing them dynamically, one fixes μ' and L' all the time in its algorithm. According to the lower bound $\Omega(\sqrt{\kappa} \ln \frac{1}{\epsilon})$, we know that its convergence rate cannot be better than $O(\sqrt{\kappa'} \ln \frac{1}{\epsilon}) = O(\sqrt{c_1 c_2 \kappa} \ln \frac{1}{\epsilon})$, where $\kappa' = L'/\mu'$. However, if one combines with our Algorithm 2 (guessing the parameters dynamically), the convergence rate can be improved to $O(\sqrt{\kappa} \ln \frac{1}{\epsilon} + \sqrt{\kappa} (\ln \kappa \ln(c_1 c_2 \kappa))^2)$ according to our Theorem 3 by letting $\delta = \mu'$ and $B\delta = L'$ (hence $B = c_1 c_2 \kappa$). Note that there is no ϵ (accuracy) in the second term $\sqrt{\kappa} (\ln \kappa \ln(c_1 c_2 \kappa))^2$. Thus the rate turns to the optimal $O(\sqrt{\kappa} \ln \frac{1}{\epsilon})$ when $\epsilon \rightarrow 0$. To achieve an ϵ -approximate solution, our

guessing algorithm can improve the convergence a lot especially for an imprecise estimate at the beginning (i.e., c_1 and c_2 are very large). The corresponding experimental results in Section 5.1 (see Figure 6) indeed validate our theoretical results.

5 Experiments

In this section, we conduct the numerical experiments on the real-world UCI datasets² and synthetic datasets. We compare the performance among these five algorithms: Anderson Acceleration (AA), Anderson-Chebyshev acceleration (AA-Cheby), vanilla Gradient Descent (GD), Nesterov’s Accelerated Gradient Descent (NAGD) (Nesterov, 2004) and Regularized Minimal Polynomial Extrapolation (RMPE) with $k = 5$ (same as (Scieur et al., 2016)).

Regarding the hyperparameters, we directly set them from their corresponding theoretical results. See Proposition 1 of (Lessard et al., 2016) for GD and NAGD. For RMPE5, we follow the same setting as in (Scieur et al., 2016). For our AA/AA-Cheby, we set them according to our Theorem 1 and 2.

Figure 2 demonstrates the convergence performance of these algorithms in general nonlinear case and Figures 3–5 demonstrate the convergence performance in quadratic case. The last Figure 6 demonstrates the convergence performance of these algorithms combined with our guessing algorithm (Algorithm 2). The values of m in the caption of figures denote the mixing parameter of Anderson acceleration algorithms (see Line 5 of Algorithm 1).







Figure 2: Logistic regression, $m = 3$

In Figure 2, we use the negative log-likelihood as the loss function f (logistic regression), i.e., $f(\theta) = -\sum_{i=1}^n (y_i \log \phi(\theta^T x_i) + (1 - y_i) \log(1 - \phi(\theta^T x_i)))$, where $\phi(z) = 1/(1 + \exp(-z))$. We run these five algorithms on real-world *diabetes* and *cancer* datasets which are standard UCI datasets. The x-axis and y-axis represent the number of iterations and the norm of the gradient of loss function respectively.

²The UCI datasets can be downloaded from <https://archive.ics.uci.edu/ml>


Figure 3: $\kappa \in [0, 500]$; $m = 3$ (left), $m = 5$ (right)

Figure 4: $\kappa \in [500, 2000]$; $m = 3$ (left), $m = 5$ (right)

Figure 5: $\kappa \in [2000, 5000]$; $m = 3$ (left), $m = 5$ (right)

Figures 3–5 demonstrate the convergence performance for the quadratic case, where $f(x) = \frac{1}{2}x^T Ax - b^T x$. Concretely, we compared the convergence performance among these algorithms when the condition number $\kappa(A)$ and the mixing parameter m are varied, e.g., the left figure in Figure 3 is the case $\kappa \in [0, 500]$ and $m = 3$. Recall that m is the mixing parameter for Anderson acceleration algorithms (see Line 5 of Algorithm 1). We run these five algorithms on the synthetic datasets in which we randomly generate the A and b for the loss function f . Note that for randomly generated A satisfying the property of $A \in \mathcal{S}_{++}^d$, we randomly generate B instead and let $A \triangleq B^T B$.

In conclusion, Anderson acceleration methods converge the fastest no matter it is a quadratic function or general function in all of our experiments. The efficient Anderson acceleration methods can be viewed as the extension of momentum methods (e.g., NAGD) since GD is the special case of Anderson Acceleration with $m = 0$, and to some extent NAGD can be viewed as $m = 1$. Combined with our theoretical results (i.e., optimal convergence rate in quadratic case and linear-quadratic convergence in general case), the experimental results validate that Anderson acceleration methods are efficient both in theory and practice.

5.1 Experiments for Guessing Algorithm

In this section, we conduct the experiments for guessing the hyperparameters (i.e., μ, L) dynamically using our Algorithm 2.

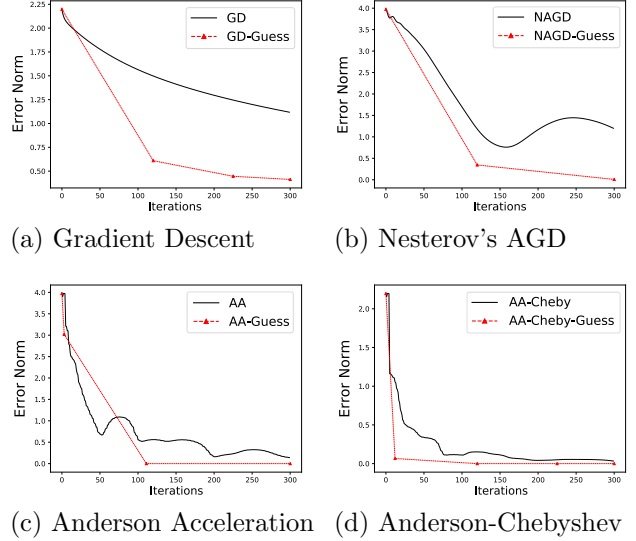


Figure 6: Algorithms with/without guessing algorithm

In Figure 6, we separately consider these algorithms. For each of them, we compare its convergence performance between its original version and the one combined with our guessing algorithm (Algorithm 2). The experimental results show that all these four algorithms combined with our guessing algorithm achieve much better performance than their original versions. Thus it validates our theoretical results (see Theorem 3 and its following Remark).

6 Conclusion

In this paper, we prove that Anderson acceleration with Chebyshev polynomial can achieve the optimal convergence rate $O(\sqrt{\kappa} \ln \frac{1}{\epsilon})$, which improves the previous result $O(\kappa \ln \frac{1}{\epsilon})$ provided by (Toth and Kelley, 2015). Thus it can deal with ill-conditioned problems (condition number κ is large) more efficiently. Furthermore, we also prove the linear-quadratic convergence of Anderson acceleration for minimizing general nonlinear problems. Besides, if the hyperparameters (e.g., the Lipschitz smooth parameter L) are not available, we propose a guessing algorithm for guessing them dynamically and also prove a similar convergence rate. Finally, the experimental results demonstrate that the efficient Anderson acceleration methods converge significantly faster than other algorithms. This validates that Anderson-Chebyshev acceleration is efficient both in theory and practice.

Acknowledgements

Zhize was supported by the Office of Sponsored Research of KAUST, through the Baseline Research Fund of Prof. Peter Richtárik. Jian was supported in part by the National Natural Science Foundation of China Grant 61822203, 61772297, 61632016, 61761146003, and the Zhongguancun Haihua Institute for Frontier Information Technology and Turing AI Institute of Nanjing. The authors also would like to thank Francis Bach, Claude Brezinski, Rong Ge, Damien Scieur, Le Zhang and anonymous reviewers for useful discussions and suggestions.

References

- A Aitken. On bernoulli’s numerical solution of algebraic equations. *Proceedings of the Royal Society of Edinburgh*, 46:289–305, 1926.
- Zeyuan Allen-Zhu. Katyusha: the first direct acceleration of stochastic gradient methods. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 1200–1205. ACM, 2017.
- Donald G Anderson. Iterative procedures for nonlinear integral equations. *Journal of the ACM*, 12(4):547–560, 1965.
- Stephen Boyd and Lieven Vandenbergh. *Convex optimization*. Cambridge university press, 2004.
- Claude Brezinski. Convergence acceleration during the 20th century. *Journal of Computational and Applied Mathematics*, 122:1–21, 2000.
- Claude Brezinski and M Redivo Zaglia. Extrapolation methods: theory and practice. 1991.
- Claude Brezinski, Michela Redivo-Zaglia, and Yousef Saad. Shanks sequence transformations and anderson acceleration. *SIAM Review*, 60(3):646–669, 2018.
- Sébastien Bubeck. Convex optimization: Algorithms and complexity. *Foundations and Trends® in Machine Learning*, 8(3-4):231–357, 2015.
- Steven Russell Capehart. *Techniques for accelerating iterative methods for the solution of mathematical problems*. PhD thesis, Oklahoma State University, 1989.
- Claire Evans, Sara Pollock, Leo G Rebholz, and Mengying Xiao. A proof that anderson acceleration increases the convergence rate in linearly converging fixed point methods (but not in quadratically converging ones). *arXiv preprint arXiv:1810.08455*, 2018.
- V Eyert. A comparative study on methods for convergence acceleration of iterative vector sequences. *Journal of Computational Physics*, 124(2):271–285, 1996.
- Cong Fang, Chris Junchi Li, Zhouchen Lin, and Tong Zhang. Spider: Near-optimal non-convex optimization via stochastic path-integrated differential estimator. In *Advances in Neural Information Processing Systems*, pages 687–697, 2018.
- Haw-ren Fang and Yousef Saad. Two classes of multi-secant methods for nonlinear acceleration. *Numerical Linear Algebra with Applications*, 16(3):197–221, 2009.
- Rong Ge, Zhize Li, Weiyao Wang, and Xiang Wang. Stabilized svrg: Simple variance reduction for non-convex optimization. In *Conference on Learning Theory*, 2019.
- GH Golub and CF Van Loan. Matrix computations. 3rd ed., *The John Hopkins University Press, Baltimore, MD*, 1996.
- Louis A Hageman and David M Young. *Applied iterative methods*. Courier Corporation, 2012.
- Nicholas J Higham and Nataša Strabić. Anderson acceleration of the alternating projections method for computing the nearest correlation matrix. *Numerical Algorithms*, 72(4):1021–1042, 2016.
- Arthur E Hoerl and Robert W Kennard. Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics*, 12(1):55–67, 1970.
- Guanghui Lan and Yi Zhou. Random gradient extrapolation for distributed and stochastic optimization. *SIAM Journal on Optimization*, 28(4):2753–2782, 2018.
- Guanghui Lan, Zhize Li, and Yi Zhou. A unified variance-reduced accelerated gradient method for convex optimization. In *Advances in Neural Information Processing Systems*, 2019.
- Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *nature*, 521(7553):436, 2015.
- Lihua Lei, Cheng Ju, Jianbo Chen, and Michael I Jordan. Non-convex finite-sum optimization via scsg methods. In *Advances in Neural Information Processing Systems*, pages 2348–2358, 2017.
- Laurent Lessard, Benjamin Recht, and Andrew Packard. Analysis and design of optimization algorithms via integral quadratic constraints. *SIAM Journal on Optimization*, 26(1):57–95, 2016.
- Zhize Li. Ssrgd: Simple stochastic recursive gradient descent for escaping saddle points. In *Advances in Neural Information Processing Systems*, 2019.
- Zhize Li and Jian Li. A simple proximal stochastic gradient method for nonsmooth nonconvex optimization. In *Advances in Neural Information Processing Systems*, pages 5569–5579, 2018.

- Zhize Li, Tianyi Zhang, Shuyu Cheng, Jun Zhu, and Jian Li. Stochastic gradient hamiltonian monte carlo with variance reduction for bayesian inference. *Machine Learning*, 108(8-9):1701–1727, 2019.
- John Loffeld and Carol S Woodward. Considerations on the implementation and use of anderson acceleration on distributed memory and gpu-based parallel computers. *Advances in the Mathematical Sciences*, page 417, 2016.
- Yurii Nesterov. *Introductory Lectures on Convex Optimization: A Basic Course*. Kluwer, 2004.
- Maxim A Olshanskii and Eugene E Tyrtysnikov. *Iterative methods for linear systems: theory and applications*. SIAM, 2014.
- Florian A Potra and Hans Engler. A characterization of the behavior of the anderson acceleration on linear problems. *Linear Algebra and its Applications*, 438(3):1002–1011, 2013.
- Phanisri P Pratapa, Phanish Suryanarayana, and John E Pask. Anderson acceleration of the jacobi iterative method: An efficient alternative to krylov methods for large, sparse linear systems. *Journal of Computational Physics*, 306:43–54, 2016.
- Theodore J Rivlin. *The Chebyshev polynomials*. Wiley, 1974.
- Yousef Saad and Martin H Schultz. Gmres: A generalized minimal residual algorithm for solving nonsymmetric linear systems. *SIAM Journal on scientific and statistical computing*, 7(3):856–869, 1986.
- Damien Scieur, Alexandre d’Aspremont, and Francis Bach. Regularized nonlinear acceleration. In *Advances in Neural Information Processing Systems*, pages 712–720, 2016.
- Damien Scieur, Edouard Oyallon, Alexandre d’Aspremont, and Francis Bach. Nonlinear acceleration of deep neural networks. *arXiv preprint arXiv:1805.09639v1*, 2018.
- Damien Scieur, Edouard Oyallon, Alexandre d’Aspremont, and Francis Bach. Online regularized nonlinear acceleration. *arXiv preprint arXiv:1805.09639v2*, 2019.
- Daniel Shanks. Non-linear transformations of divergent and slowly convergent sequences. *Studies in Applied Mathematics*, 34(1-4):1–42, 1955.
- 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.
- Alex Toth and CT Kelley. Convergence analysis for anderson acceleration. *SIAM Journal on Numerical Analysis*, 53(2):805–819, 2015.
- Homer F Walker and Peng Ni. Anderson acceleration for fixed-point iterations. *SIAM Journal on Numerical Analysis*, 49(4):1715–1735, 2011.
- Junzi Zhang, Brendan O’Donoghue, and Stephen Boyd. Globally convergent type-i anderson acceleration for non-smooth fixed-point iterations. *arXiv preprint arXiv:1808.03971*, 2018.
- Dongruo Zhou, Pan Xu, and Quanquan Gu. Stochastic nested variance reduction for nonconvex optimization. *arXiv preprint arXiv:1806.07811*, 2018.