Global convergence of the Hessenberg QR algorithm

Jorge Garza-Vargas

Joint work with Jess Banks and Nikhil Srivastava

BIRS
March of 2023

Caltech

About this talk

- *The eigenvalue problem:* "accurately" compute the eigenvalues and eigenvectors of an input matrix $A \in \mathbb{C}^{n\times n}$.
- *The QR algorithm:* The go-to method for obtaining the full eigendecomposition when no particular structure of A is known.
- *Rigorous guarantees:* We show that (with high probability) the QR algorithm can be used solve the eigenvalue problem in $\tilde{O}(n^3)$ arithmetic operations.

Table of Contents

Backward approximate eigenvalue problem

We will focus on the following version of the eigenvalue problem:

Problem (Backward approximate Schur form): Given a matrix $A \in \mathbb{C}^{n \times n}$ and $\delta > 0$ find $T, U \in \mathbb{C}^{n \times n}$ with T upper triangular and U unitary such that

$$\|A - UTU^*\| \leq \delta \|A\|.$$

In general, the quality of the forward approximation is given by

$$\delta \cdot \text{Condition number of the problem}$$

Caltech

Eigenvector condition number

If $A \in \mathbb{C}^{n\times n}$ is diagonalizable, define its eigenvector condition number as

$$\kappa_V(A) = \inf_{V : A = VDV^{-1}} \|V\|\|V^{-1}\|.$$

- When A is normal $\kappa_V(A) = 1$. When A is non-diagonalizable (e.g. a Jordan block) $\kappa_V(A) = \infty$.

Theorem (Bauer-Fike 60) For any $A, E \in \mathbb{C}^{n\times n}$, with $\|E\| \leq \epsilon$

$$\mathrm{Spec}(A+E) \subset \bigcup_{\lambda_i \in \mathrm{Spec}(A)} D(\lambda_i, \epsilon\kappa_V(A))$$

Caltech

Table of Contents

Caltech

The QR algorithm

(Francis 61, Kublanovskaya 62) On an input $A \in \mathbb{C}^{n\times n}$:

0. Put A in Hessenberg form, that is, compute a Hessenberg H_0 with

$$H_0 := U^* A U \quad \text{for } U \text{ unitary.}$$

(Hessenberg matrices) An upper Hessenberg matrix $H \in \mathbb{C}^{n\times n}$ is a matrix with $H(i,j) = 0$ for all $i > j+1$. E.g.

$$\begin{pmatrix} * & * & * & * \\ * & * & * & * \\ 0 & * & * & * \\ 0 & 0 & * & * \end{pmatrix}$$

(Hessenberg form) For any $A \in \mathbb{C}^{n\times n}$ one can compute in $O(n^3)$ operations a Hessenberg matrix H that is unitarily equivalent to A, that is $H = U^* A U$.

The QR algorithm

(Francis 61, Kublanovskaya 62) On an input $A \in \mathbb{C}^{n \times n}$:

0. Put A in Hessenberg form, that is, compute a Hessenberg H_0 with
$$H_0 := U^* A U \quad \text{for } U \text{ unitary.}$$
1. Generate a sequence $H_0, H_1, \ldots$ of Hessenberg matrices:
$$\text{If} \quad p_t(H_t) = Q_t R_t \quad \text{then} \quad H_{t+1} = Q_t^* H_t Q_t$$

where $p_t = \mathrm{Sh}(H_t)$. The roots of $p_t(z)$ are "guesses" for $\mathrm{Spec}(H_t)$, and the recipe for choosing the p_t is the *shifting strategy*.

- *(Unitary equivalence)* $A = U_t^* H_t U_t$ where $U_t = U Q_0 \cdots Q_t$.
- *(The hope)* The shifting strategy leads to rapid convergence of H_t to a triangular T, and therefore:
$$\lim_{t \to \infty} U_t^* H_t U_t = U_\infty^* T U_\infty = A.$$

Caltech

Decoupling and deflation

The following quantitative notion of convergence proves useful.

- *δ-Decoupling:* We say that $H \in \mathrm{Hess}(n)$ is δ-decoupled if $|H(i, i-1)| < \delta \|H\|$ for *some* $i = 1, \ldots, n$.

- *Deflation:* Once a matrix is decoupled we can deflate it into smaller subproblems:

$$\begin{pmatrix} * & * & * & * & * \\ * & * & * & * & * \\ 0 & \text{small} & * & * & * \\ 0 & 0 & * & * & * \\ 0 & 0 & 0 & * & * \end{pmatrix} \longrightarrow \begin{pmatrix} * & * & * & * & * \\ * & * & * & * & * \\ 0 & 0 & * & * & * \\ 0 & 0 & * & * & * \\ 0 & 0 & 0 & * & * \end{pmatrix}$$

Gold standard: Devise a shifting strategy (with moderate $\deg(p_t) = k$) that guarantees δ-decoupling in $\mathrm{polylog}(1/\delta) \implies O(\mathrm{polylog}(1/\delta) k n^3)$ diagonalization algorithm.

Table of Contents

Caltech

Previous work

- *Hermitian case (Wilkinson 68, Dekker-Traub 71).* The Wilkinson shift achieves δ-decoupling in $\log(1/\delta)$ iterations on any Hermitian input.
- *Unitary case (Eberlein-Huang 75, Wang-Gragg 2002).* A mixture of the Wilkinson shift and an exceptional shift achieve δ-decoupling in $\log(1/\delta)$ iterations on any unitary input.
- *General case:* A complex but empirically reliable and practical version of the algorithm has been obtained over the decades by addressing non-convergent cases with heuristic modifications and improvements.

Main result (Controlled $\kappa_V(H)$)

Theorem (Banks, GV, Srivastava 2021-2022). For every k, we devise a shifting strategy of degree k, which achieves δ-decoupling in $\log(1/\delta)$ iterations, provided that the input H satisfies

$$k \geq C \log \kappa_V(H) \log\log \kappa_V(H).$$

- Computing each shift has a cost of at most $O(k^2 n^2)$ arithmetic operations.
- This allows to solve the eigenvalue problem, with accuracy δ, in $O(\log(1/\delta) k^2 n^3)$ operations.

Main result (Arbitrary inputs)

Random matrix theory (Armentano et al. 2015, Banks et al. 2019, Banks et al. 2020, Jain et al. 2020, Erdös et al. 2023)
Let G_n be a normalized $n \times n$ Ginibre matrix. For any $A \in \mathbb{C}^{n\times n}$ with $\|A\| \leq 1$ and $\gamma > 0$, with high probability

$$\kappa_V(A + \gamma G_n) \leq \frac{n^4}{\gamma}.$$

Preprocessing: Rather than running ShiftedQR on A, run it on $\tilde{A} = A + \gamma G_n$, say for $\gamma = \frac{\delta}{10}$. So with high probability

$$C \log \kappa_V(\tilde{A}) \log\log \kappa_V(\tilde{A}) = O(\log(n/\delta) \log\log(n/\delta))$$

Conclusion: We get an algorithm which WHP runs in $O(n^3 \log(n/\delta)^3 \log\log(n/\delta)^2) = \tilde{O}(n^3)$ arithmetic operations on any input.

Caltech

Table of Contents

Caltech

The algorithm: Mixed strategy for normal matrices

$H_t - H_t(n,n)I_n = Q_t R_t, \quad \hat{H}_t = Q_t^* H_t Q_t$ Rayleigh shift

If $|\hat{H}_t(n,n-1)| < .8|H_t(n,n-1)|$, put $H_{t+1} = \hat{H}_t$

Else: Take $\mathcal{N} \subset \mathcal{A}_{H_t(n,n-1)}$ with 20 points Exceptional shift

For $\alpha \in \mathcal{N}$ $\quad H_t - \alpha I_n = Q_t R_t, \quad \hat{H}_t = Q_t^* H_t Q_t$

If $|\hat{H}_t(n,n-1)| < .8|H_t(n,n-1)|$, put $H_{t+1} = \hat{H}_t$

Claim: WHP δ-decoupling is attained in $O(\log(1/\delta))$ iterations.

$$\begin{pmatrix} * & * & * & * \\ * & * & * & * \\ 0 & * & * & * \\ 0 & 0 & * & * \end{pmatrix} \longrightarrow \begin{pmatrix} * & * & * & * \\ * & * & * & * \\ 0 & * & * & * \\ 0 & 0 & \text{small} & * \end{pmatrix}$$

Caltech

The shift design: insight

- We use a *potential function* to track progress of the algorithm.
- We use a *main shift* that avoids transient behavior and guarantees that *progress* is not lost. This is where the assumption $k \geq \log \kappa_V(H) \log \log \kappa_V(H)$ is necessary.
- We use an *exceptional shift* to avoid *stagnation* when little to no progress is made.

Conclusions

- *Theory:* We prove that a relatively simple shifting strategy can achieve rapid decoupling and we have a clear conceptual explanation of how it works.
- *Practice:* Our theoretical algorithm is not a prescription for what should be done in practice, and does not seek to replace the current fine-tuned LAPACK routines.
- *The dream:* Our work suggests that there might be a simple, efficient, and infallible shifting strategy for the QR algorithm.

Thanks!

Caltech