Computing Eigenvectors of Symmetric Tridiagonals with the Correct Number of Sign Changes

Plamen Koev
San José State University

Bulgarian Academy of Sciences

June 10, 2022

Outline

- The symmetric tridiagonal eigenvalue problem is central!
- Every symmetric matrix first reduced to a tridiagonal
- Existing algorithms all have drawbacks
- Focus on the mathematical properties of the eigenvectors, preserve those numerically
- New results

Every symmetric matrix reduces to tridiagonal - 1

$$\begin{bmatrix} * & * & * & * \\ * & * & * & * \\ * & * & * & * \\ * & * & * & * \end{bmatrix}$$

- Start with a symmetric matrix

Every symmetric matrix reduces to tridiagonal - 2

$$\begin{bmatrix} 1 & & & \\ & 1 & & \\ & & c & s \\ & & -s & c \end{bmatrix} \begin{bmatrix} * & * & * & * \\ * & * & * & * \\ * & * & * & * \\ * & * & * & * \end{bmatrix} = \begin{bmatrix} * & * & * & * \\ * & * & * & * \\ * & * & * & * \\ 0 & * & * & * \end{bmatrix}$$

- Start with a symmetric matrix
- Apply an orthogonal matrix to create a zero

Every symmetric matrix reduces to tridiagonal - 3

$$\begin{bmatrix} 1 & & & \\ & 1 & & \\ & & c & s \\ & & -s & c \end{bmatrix} \begin{bmatrix} * & * & * & * \\ * & * & * & * \\ * & * & * & * \\ * & * & * & * \end{bmatrix} \begin{bmatrix} 1 & & & \\ & 1 & & \\ & & c & -s \\ & & s & c \end{bmatrix} = \begin{bmatrix} * & * & * & 0 \\ * & * & * & * \\ * & * & * & * \\ 0 & * & * & * \end{bmatrix}$$

- Start with a symmetric matrix
- Apply an orthogonal matrix to create a zero
- Complete similarity; Result symmetric: $(HAH^T)^T = HAH^T$

Every symmetric matrix reduces to tridiagonal - 4

$$\begin{bmatrix} 1 & & & \\ & 1 & & \\ & & c & s \\ & & -s & c \end{bmatrix} \begin{bmatrix} * & * & * & * \\ * & * & * & * \\ * & * & * & * \\ * & * & * & * \end{bmatrix} \begin{bmatrix} 1 & & & \\ & 1 & & \\ & & c & -s \\ & & s & c \end{bmatrix} = \begin{bmatrix} * & * & * & 0 \\ * & * & * & * \\ * & * & * & * \\ 0 & * & * & * \end{bmatrix}$$

- Start with a symmetric matrix
- Apply an orthogonal matrix to create a zero
- Complete similarity, result symmetric, $(HAH^T)^T = HAH^T$
- Continue until tridiagonal

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

Desired properties of computed eigenvectors

1. Each computed independently in $O(n)$ time
2. Error bound, as accurate as data deserves:

$$\sin \angle(v_i, \hat{v}_i) = O(\varepsilon)/\mathrm{relgap}_i,$$

   - ε = machine precision
   - $\mathrm{relgap}_i = \min_{j \neq i} \frac{|\lambda_j - \lambda_i|}{|\lambda_i|}$
     smallest relative gap between λ_i and rest of spectrum
3. Eigenvectors orthogonal

Algorithms:

- QR iteration delivers 2), 3), but not 1): cost $O(n^3)$
- dqds delivers 1), 2), but not 3) for clusters

Our focus:

- 1) and 2), and ith eigenvector has $i - 1$ changes of sign
- Orthogonality, 3): promising

Preservation of math properties of computed quantities

- Key idea in computing!
- Subnormal numbers in IEEE 754 floating point arithmetic designed so that

$$a \neq b \Longrightarrow a - b \neq 0$$

- Fused multiply-add rejected, kills commutativity: if

$$a = \mathrm{fl}(xy)$$

then

$$\mathrm{fl}(xy - a) \neq 0$$

- Key idea in many accurate eigenvalue algorithms: Sequence of simpler matrix with same structure
  - Demmel: (poly) Vandermonde SVDs, Cauchy, etc.
  - K.: Totally nonegative matrices (including singular)

Key ideas

- Compute eigenvectors with correct mathematical properties:
  - Number of sign changes
  - Hope for orthogonality

  both determined very accurately!
- Under WLOG conditions, ith eigenvector has exactly $i - 1$ changes of sign
- Examine math reasons this happens
- Replicate in algorithm!

WLOG positive definite, positive, and unreduced

$$A = \begin{bmatrix} \alpha_1 & \beta_1 & & \\ \beta_1 & \alpha_2 & \ddots & \\ & \ddots & \ddots & \beta_{n-1} \\ & & \beta_{n-1} & \alpha_n \end{bmatrix}$$

- Shift $A \to A + \lambda I$ if needed to make positive definite
- Offdiafgonal fixed by diag(± 1):

$$\begin{bmatrix} 1 & & & \\ & 1 & & \\ & & -1 & \\ & & & 1 \end{bmatrix} \begin{bmatrix} 1 & 1 & & \\ 1 & 2 & -1 & \\ & -1 & 4 & 1 \\ & & 1 & 5 \end{bmatrix} \begin{bmatrix} 1 & & & \\ & 1 & & \\ & & -1 & \\ & & & 1 \end{bmatrix} = \begin{bmatrix} 1 & 1 & & \\ 1 & 2 & 1 & \\ & 1 & 4 & 1 \\ & & 1 & 5 \end{bmatrix}$$

- Unreduced, $\beta_i > 0$, (or it splits), so eigenvalues distinct
- Thus Totally Nonnegative, i.e., all minors ≥ 0
  $\Rightarrow$ ith eigenvector has $i - 1$ changes of sign

The ith eigenvector has $i - 1$ changes of sign

- Example, when eigenvalues ordered $\lambda_1 > \cdots > \lambda_n > 0$

$$\begin{bmatrix} -0.0013 & -0.0313 & 0.2674 & 0.7385 & -0.5778 & 0.2144 & 0.0472 \\ -0.0084 & -0.1163 & 0.4827 & 0.3693 & 0.4652 & -0.5833 & -0.2457 \\ -0.0309 & -0.2684 & 0.5473 & -0.1231 & 0.3255 & 0.4597 & 0.5431 \\ -0.0876 & -0.4598 & 0.3342 & -0.3693 & -0.2747 & 0.1901 & -0.6490 \\ -0.2096 & -0.5871 & -0.1263 & -0.1231 & -0.3404 & -0.5196 & 0.4411 \\ -0.4452 & -0.4265 & -0.4714 & 0.3693 & 0.3777 & 0.3043 & -0.1613 \\ -0.8655 & 0.4189 & 0.2146 & -0.1231 & -0.0993 & -0.0611 & 0.0248 \end{bmatrix}$$

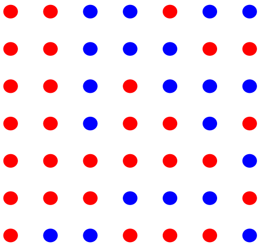

Example: Roundoff errors kill the sign pattern, 30×30

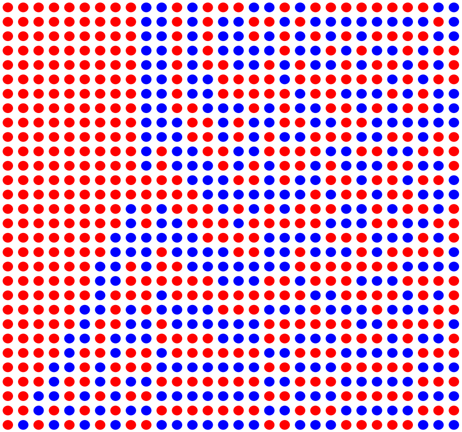

Example: Roundoff errors kill the sign pattern, 50 $\times$ 50

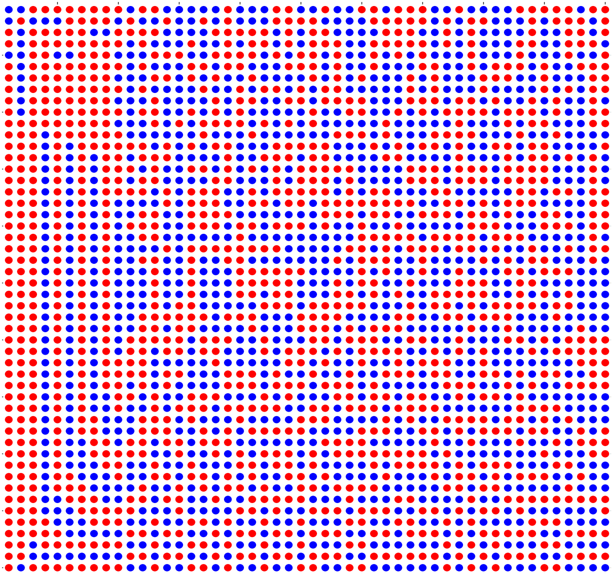

Where is sign pattern coming from? Interlacing!

Define $\quad v(x) = \begin{bmatrix} \chi_0(x) \\ \chi_1(x) \\ \vdots \\ \chi_{n-1}(x) \end{bmatrix}$; $\quad$ ith eigenvector $\sim v(\lambda_i)$

- χ_{k-1} = char poly of leading $k \times k$ principal submatrix
- Cauchy Interlace Thm: roots of χ_{k-1} interlace those of χ_k
- Consider sequence

$$\chi_{k-1}(x), \chi_k(x), \chi_{k+1}(x) \tag{1}$$

  around some root z of χ_k (i.e., when $\chi_k(x)$ switches signs)
- If χ_{k-1} has t roots $< z$, then χ_{k+1} has $t+1$ roots $< z$

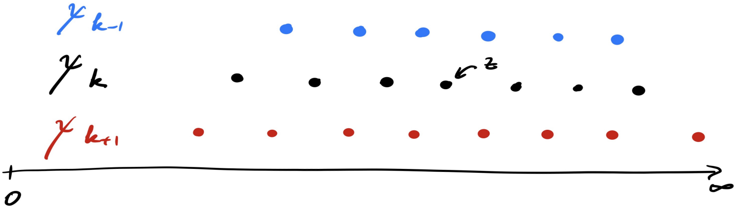

Where is sign pattern coming from? Interlacing!

$$\text{Define} \quad v(x) = \begin{bmatrix} \chi_0(x) \\ \chi_1(x) \\ \vdots \\ \chi_{n-1}(x) \end{bmatrix}; \quad i\text{th eigenvector} \sim v(\lambda_i)$$

- χ_{k-1} = char poly of leading $k \times k$ principal submatrix
- Cauchy Interlace Thm: roots of χ_{k-1} interlace those of χ_k
- Consider sequence

$$\chi_{k-1}(x), \chi_k(x), \chi_{k+1}(x) \tag{2}$$

  around some root z of χ_k (i.e., when $\chi_k(x)$ switches signs)
- If χ_{k-1} has t roots $< z$, then χ_{k+1} has $t+1$ roots $< z$
- $\Rightarrow$ $\chi_{k-1}(z)$ and $\chi_{k+1}(z)$ are of different signs
- As x passes through z, # sign changes in (2) is exactly 1.
- Sign change in $v(x)$ only occurs when x passes through root of χ_{n-1}
- $i-1$ roots of χ_{n-1} to left of $\lambda_i \Rightarrow i-1$ changes of sign in v_i

Interlacing in practice?

- Nonexistent:

9.174243420830803e-01
7.847392948263211e-01
7.701597286355866e-01
7.655000166675910e-01
7.384268399820760e-01
7.217580334780621e-01
6.073892138195941e-01
4.734859929761793e-01
4.713571537120613e-01
4.609163660873046e-01
4.243490397612005e-01
3.411246072099285e-01
3.224718071611237e-01
2.690615866769628e-01
2.428495983120915e-01
1.917452554462249e-01
1.886619767839813e-01
1.758744157991657e-01
1.527212003495465e-01
3.576273324260682e-02

9.174243420830800e-01
7.847392948263211e-01
7.701597286355862e-01
7.655000166675913e-01
7.384268399820757e-01
7.217580334219912e-01
6.073892138195942e-01
4.713571537120613e-01
4.609163660873047e-01
4.243490397612006e-01
3.411246072099282e-01
3.224718071611238e-01
2.690615866769628e-01
2.428495983120914e-01
1.917452554462248e-01
1.886619767839813e-01
1.758744157991657e-01
1.527212003649134e-01
3.576273324260682e-02

Must compute eigenvalues of A and $A_{1:n-1,1:n-1}$ simultaneously to ensure interlacing

- Work with bidiagonal Cholesky factor instead (Parlett)

$$\underbrace{\begin{bmatrix} a_1 & & & \\ b_1 & a_2 & & \\ & \ddots & \ddots & \\ & & b_{n-1} & a_n \end{bmatrix}}_{B} \underbrace{\begin{bmatrix} a_1 & b_1 & & \\ & a_2 & \cdot & \\ & & \ddots & b_{n-1} \\ & & & a_n \end{bmatrix}}_{B^T} = \underbrace{\begin{bmatrix} \alpha_1 & \beta_1 & & \\ \beta_1 & \alpha_2 & \cdot & \\ & \ddots & \ddots & \beta_{n-1} \\ & & \beta_{n-1} & \alpha_n \end{bmatrix}}_{A}$$

- Eigenvalues of A are squares of singular values of B
- Use differential quotient-difference with shift (dqds)
- Run dqds on B and $B_{1:n-1,1:n-1}$ simultaneously!

The dqds algorithm

- Input:

$$B = \begin{bmatrix} a_1 & & & \\ b_1 & a_2 & & \\ & \ddots & \ddots & \\ & & b_{n-1} & a_n \end{bmatrix}.$$

  and shift $\sigma < \sigma_{\min}(B)$
- Output: Cholesky factor, B' : $\quad B'(B')^T = B^T B - \sigma^2 I$

$$\begin{array}{l} d_1 = q_1 - \sigma^2 \\ \text{for } i = 1 \text{ to } n-1 \\ \quad q_i' = d_i + e_i \\ \quad e_i' = e_i q_{i+1}/q_i' \\ \quad d_i = q_{i+1}(d_{i-1}/q_i') - \sigma^2 \\ \text{end} \\ q_n' = d_n + e_n \end{array}$$

(where $q_i = a_i^2$, $e_i = b_i^2$, typical)

dqds computes singular values to high rel accuracy

- 2-3 steps max reveal a singular value
- Accumulate shifts to produce $\sigma_{\min}^2 = \lambda_n$

$$B = \begin{bmatrix} a_1 & & & & \\ b_1 & a_2 & & & \\ & \ddots & \ddots & & \\ & & b_{n-2} & a_{n-1} & \\ & & & b_{n-1} & a_n \end{bmatrix} \longrightarrow B' = \begin{bmatrix} a'_1 & & & & \\ b'_1 & a'_2 & & & \\ & \ddots & \ddots & & \\ & & b'_{n-2} & a'_{n-1} & \\ & & & b'_{n-1} & a'_n \end{bmatrix}$$

where a'_n is *tiny* and is thus *deflated*:

$$B'' = \begin{bmatrix} a'_1 & & & \\ b'_1 & a'_2 & & \\ & \ddots & \ddots & \\ & & b'_{n-2} & a'_{n-1} \\ & & & b'_{n-1} \end{bmatrix}$$

- Continue the same way to produce $(\sigma_{\min}(B''))^2 = \lambda_{n-1} - \lambda_n$
- ... until all eigenvalues computed, accurately

New results on dqds

- Monotonicity: locates $\sigma_{\min}$ between two consecutive floating point numbers
- Perfect shift $=$ floating point number to left of $\sigma_{\min}$
- dqds can be run simultaneously on B and $B_{1:n-1,1:n-1}$
- Ensures interlacing, and in turn, number of sign changes!

New algorithm: Double dqds

- Runs simultaneously on B and $B_{1:n-1,1:n-1}$ (eigenvalues $\lambda_1 > \cdots > \lambda_n > 0$ and $\mu_1 > \cdots > \mu_{n-1} > 0$)
- Pick the perfect shift σ^2 for B. This is $\hat{\lambda}_n$! Run dqds

$$B: \begin{bmatrix} a_1 & & & & \\ b_1 & a_2 & & & \\ & \ddots & \ddots & & \\ & & b_{n-2} & a_{n-1} & \\ & & & b_{n-1} & a_n \end{bmatrix} \longrightarrow \begin{bmatrix} a'_1 & & & & \\ b'_1 & a'_2 & & & \\ & \ddots & \ddots & & \\ & & b'_{n-2} & a'_{n-1} & \\ & & & b'_{n-1} & a'_n \end{bmatrix}$$

$$B_{1:n-1,1:n-1}: \begin{bmatrix} a_1 & & & \\ b_1 & a_2 & & \\ & \ddots & \ddots & \\ & & b_{n-2} & a_{n-1} \end{bmatrix} \longrightarrow \begin{bmatrix} a'_1 & & & \\ b'_1 & a'_2 & & \\ & \ddots & \ddots & \\ & & b'_{n-2} & \bar{a}'_{n-1} \end{bmatrix}$$

- $\bar{a}'_{n-1}$ only different entry, provably $\bar{a}'_{n-1} < a'_{n-1}$
- a'_n is *deflated*, obtaining new pair of matrices

$$\underbrace{\begin{bmatrix} a'_1 & & & \\ b'_1 & a'_2 & & \\ & \ddots & \ddots & \\ & & b'_{n-2} & a'_{n-1} \\ & & & b'_{n-1} \end{bmatrix}}_{B'} \quad \text{and} \quad \underbrace{\begin{bmatrix} a'_1 & & & \\ b'_1 & a'_2 & & \\ & \ddots & \ddots & \\ & & b'_{n-2} & \bar{a}'_{n-1} \end{bmatrix}}_{B''}$$

New algorithm: Double dqds

- New pair of matrices

$$\underbrace{\begin{bmatrix} a'_1 & & & \\ b'_1 & a'_2 & & \\ & \ddots & \ddots & \\ & & b'_{n-2} & a'_{n-1} \\ & & & b'_{n-1} \end{bmatrix}}_{B'} \quad \text{and} \quad \underbrace{\begin{bmatrix} a'_1 & & & \\ b'_1 & a'_2 & & \\ & \ddots & \ddots & \\ & & b'_{n-2} & \bar{a}'_{n-1} \end{bmatrix}}_{B''}$$

- $b'_{n-1} > 0, \bar{a}'_{n-1} < a'_{n-1}$ ensure $\sigma_{\min}(B'') < \sigma_{\min}(B'')$
- Run dqds again on B', B'' with perfect shift for B'' :

$$\hat{\mu}_{n-1} - \hat{\lambda}_n > 0$$

- The process produces

$$\hat{\lambda}_n > 0,\ \hat{\mu}_{n-1} - \hat{\lambda}_n > 0,\ \hat{\lambda}_{n-1} - \hat{\mu}_n > 0, \ldots$$

- Interlacing guaranteed, and in turn, number of sign changes

Number of sign changes guaranteed, orthogonality?

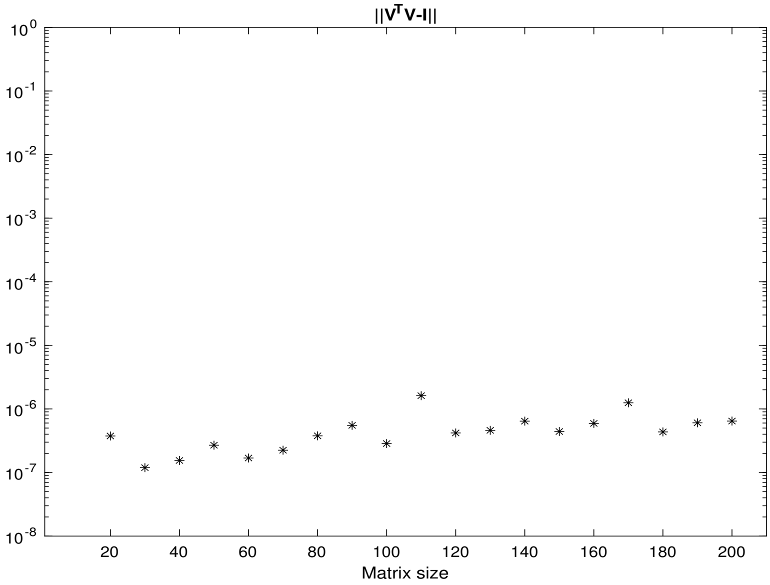

Conclusions

- Established new properties of dqds
- Eigenvectors with correct number of sign changes
- Orthogonality, what property determines that numerically?
- Joint work with Anthony Mendoza of Gavilan College
- Thank you!