# C6.1 Numerical Linear Algebra 

Yuji Nakatsukasa*

Welcome to numerical linear algebra (NLA)! NLA is a beautiful subject that combines mathematical rigor, amazing algorithms, and an extremely rich variety of applications.

What is NLA? In a sentence, it is a subject that deals with the numerical solution (i.e., using a computer) of linear systems $A x=b$ (given $A \in \mathbb{R}^{n \times n}$ (i.e., a real $n \times n$ matrix) and $b \in \mathbb{R}^{n}$ (real $n$-vector), find $x \in \mathbb{R}^{n}$ ) and eigenvalue problems $A x=\lambda x$ (given $A \in \mathbb{R}^{n \times n}$, find $\lambda \in \mathbb{C}$ and $x \in \mathbb{C}^{n}$ ), for problems that are too large to solve by hand ( $n \geq 4$ is already large; we aim for $n$ in the thousands or even millions). This can rightfully sound dull, and some mathematicians (those purely oriented?) tend to get turned off after hearing this how could such a course be interesting compared with other courses offered by the Oxford Mathematical Institute? I hope and firmly believe that at the end of the course you will all agree that there is more to the subject than you imagined. The rapid rise of data science and machine learning has only meant that the importance of NLA is still growing, with a vast number of problems in these fields requiring NLA techniques and algorithms. It is perhaps worth noting also that these fields have had enormous impact on the direction of NLA, in particular the recent and very active field of randomised algorithm was born in light of needs arising from these extremely active fields.

In fact NLA is a truly exciting field that utilises a huge number of ideas from different branches of mathematics (e.g. matrix analysis, approximation theory, and probability) to solve problems that actually matter in real-world applications. Having said that, the number of prerequisites for taking the course is the bare minimum; essentially a basic understanding of the fundamentals of linear algebra would suffice (and the first lecture will briefly review the basic facts). If you've taken the Part A Numerical Analysis course you will find it helpful, but again, this is not necessary.

The field NLA has been blessed with many excellent books on the subject. These notes will try to be self-contained, but these references will definitely help. There is a lot to learn; literally as much as you want to.

- Trefethen-Bau (97) [33]: Numerical Linear Algebra
- covers essentials, beautiful exposition
- Golub-Van Loan (12) [14]: Matrix Computations

[^0]- classic, encyclopedic
- Horn and Johnson (12) [20]: Matrix Analysis (\& Topics in Matrix Analysis (86) [19])
- excellent theoretical treatise, little numerical treatment
- J. Demmel (97) [8]: Applied Numerical Linear Algebra
- impressive content
- N. J. Higham (02) [17]: Accuracy and Stability of Algorithms
- bible for stability, conditioning
- H. C. Elman, D. J. Silvester, A. J. Wathen (14) [11]: Finite Elements and Fast Iterative Solvers
- PDE applications of linear systems, Krylov methods and preconditioning

This course covers the fundamentals of NLA. We first discuss the singular value decomposition (SVD), which is a fundamental matrix decomposition whose importance is only growing. We then turn to linear systems and eigenvalue problems. Broadly, we will cover

- Direct methods ( $n \lesssim 10,000$ ): Sections $5-10$ (except 8 )
- Iterative methods ( $n \lesssim 1,000,000$, sometimes larger): Sections 11-13
- Randomised methods ( $n \gtrsim 1,000,000$ ): Sections $14-16$
in this order. Lectures 1-4 cover the fundamentals of matrix theory, in particular the SVD, its properties and applications.

This document consists of 16 sections. Very roughly speaking, one section corresponds to one lecture (though this will not be followed strictly at all).

## Contents

0 Introduction, why $A x=b$ and $A x=\lambda x$ ? ..... 6
1 Basic LA review ..... 7
1.1 Warmup exercise ..... 8
1.2 Structured matrices ..... 9
1.3 Matrix eigenvalues: basics ..... 10
1.4 Computational complexity (operation counts) of matrix algorithms ..... 11
1.5 Vector norms ..... 11
1.6 Matrix norms ..... 12
1.7 Subspaces and orthonormal matrices ..... 13
2 SVD: the most important matrix decomposition ..... 14
2.1 (Some of the many) applications and consequences of the SVD: rank, col- umn/row space, etc ..... 16
2.2 SVD and symmetric eigenvalue decomposition ..... 17
2.3 Uniqueness etc ..... 17
3 Low-rank approximation via truncated SVD ..... 18
3.1 Low-rank approximation: image compression ..... 21
4 Courant-Fischer minmax theorem ..... 22
4.1 Weyl's inequality ..... 22
4.1.1 Eigenvalues of nonsymmetric matrices are sensitive to perturbation ..... 23
4.2 More applications of C-F ..... 23
4.3 (Taking stock) Matrix decompositions you should know ..... 24
5 Linear systems $\boldsymbol{A} \boldsymbol{x}=\boldsymbol{b}$ ..... 24
5.1 Solving $A x=b$ via LU ..... 26
5.2 Pivoting ..... 26
5.3 Cholesky factorisation for $A \succ 0$ ..... 27
6 QR factorisation and least-squares problems ..... 28
6.1 QR via Gram-Schmidt ..... 28
6.2 Towards a stable QR factorisation: Householder reflectors ..... 29
6.3 Householder QR factorisation ..... 30
6.4 Givens rotations ..... 32
6.5 Least-squares problems via QR ..... 32
6.6 QR-based algorithm for linear systems ..... 34
6.7 Solution of least-squares via normal equation ..... 34
6.8 Application of least-squares: regression/function approximation ..... 35
7 Numerical stability ..... 36
7.1 Floating-point arithmetic ..... 36
7.2 Conditioning and stability ..... 37
7.3 Numerical stability; backward stability ..... 38
7.4 Matrix condition number ..... 39
7.4.1 Backward stable+well conditioned=accurate solution ..... 39
7.5 Stability of triangular systems ..... 40
7.5.1 Backward stability of triangular systems ..... 40
7.5.2 (In)stability of $\boldsymbol{A} \boldsymbol{x}=\boldsymbol{b}$ via LU with pivots ..... 41
7.5.3 Backward stability of Cholesky for $A \succ 0$ ..... 41
7.6 Matrix multiplication is not backward stable ..... 42
7.7 Stability of Householder QR ..... 43
7.7.1 (In)stability of Gram-Schmidt ..... 44
8 Eigenvalue problems ..... 44
8.1 Schur decomposition ..... 45
8.2 The power method for finding the dominant eigenpair $A x=\lambda x$ ..... 46
8.2.1 Digression (optional): Why compute eigenvalues? Google PageRank ..... 47
8.2.2 Shifted inverse power method ..... 47
9 The QR algorithm ..... 48
9.1 QR algorithm for $A x=\lambda x$ ..... 48
9.2 QR algorithm preprocessing: reduction to Hessenberg form ..... 52
9.2.1 The (shifted) QR algorithm in action ..... 53
9.2.2 (Optional) QR algorithm: other improvement techniques ..... 54
10 QR algorithm continued ..... 55
10.1 QR algorithm for symmetric $A$ ..... 55
10.2 Computing the SVD: Golub-Kahan's bidiagonalisation algorithm ..... 56
10.3 (Optional but important) QZ algorithm for generalised eigenvalue problems ..... 56
10.4 (Optional) Tractable eigenvalue problems ..... 57
11 Iterative methods: introduction ..... 58
11.1 Polynomial approximation: basic idea of Krylov ..... 59
11.2 Orthonormal basis for $\mathcal{K}_{k}(A, b)$ ..... 60
11.3 Arnoldi iteration ..... 60
12 Arnoldi and GMRES for $A x=b$ ..... 61
12.1 GMRES convergence: polynomial approximation ..... 62
12.2 When does GMRES converge fast? ..... 64
12.3 Preconditioning for GMRES ..... 65
12.4 Restarted GMRES ..... 66
13 Symmetric case: Lanczos and Conjugate Gradient method for $\boldsymbol{A x}=\boldsymbol{b}$, $A \succ 0$ ..... 66
13.1 Lanczos iteration and Lanczos decomposition ..... 66
13.2 CG algorithm for $A x=b, A \succ 0$ ..... 67
13.3 CG convergence ..... 68
13.3.1 Chebyshev polynomials ..... 69
13.3.2 Properties of Chebyshev polynomials ..... 71
13.4 MINRES: symmetric (indefinite) version of GMRES (nonexaminable) ..... 71
13.4.1 MINRES convergence ..... 72
13.5 Preconditioned CG/MINRES ..... 73
13.6 The Lanczos algorithm for symmetric eigenproblem (nonexaminable) ..... 73
13.7 Arnoldi for nonsymmetric eigenvalue problems (nonexaminable) ..... 75
14 Randomised algorithms in NLA ..... 75
14.1 Gaussian matrices ..... 76
14.1.1 Orthogonal invariance ..... 76
14.1.2 Marchenko-Pastur: Rectangular random matrices are well conditioned ..... 77
14.2 Randomised least-squares ..... 78
14.3 "Fast" algorithm: row subset selection ..... 78
14.4 Sketch-and-solve for least-squares problems ..... 79
14.5 Sketch-to-precondition: Blendenpik ..... 81
14.5.1 Explaining $\kappa_{2}\left(A \hat{R}^{-1}\right)=O(1)$ via Marchenko-Pastur ..... 81
14.5.2 Blendenpik: solving $\min _{x}\|A x-b\|_{2}$ using $\hat{R}$ ..... 82
14.5.3 Blendenpik experiments ..... 82
15 Randomised algorithms for low-rank approximation ..... 83
15.1 Randomised SVD by Halko-Martinsson-Tropp ..... 83
15.2 HMT approximant: analysis (down from 70 pages!) ..... 85
15.3 Precise analysis for HMT (nonexaminable) ..... 86
15.4 Generalised Nyström (nonexaminable) ..... 87
15.5 MATLAB code ..... 88
15.6 Randomised algorithm for $A x=b, A x=\lambda x$ ? ..... 89
16 Conclusion and discussion ..... 89
16.1 Important (N)LA topics not treated ..... 89
16.2 Course summary ..... 90
16.3 Related courses you can take ..... 90

Notation. For convenience below we list the notation that we use throughout the course.

- $\lambda(A)$ : the set of eigenvalues of $A$. If a natural ordering exists (e.g. $A$ is symmetric so $\lambda$ is real), $\lambda_{i}(A)$ is the $i$ th (largest) eigenvalue.
- $\sigma(A)$ : the set of singular values of $A . \sigma_{i}(A)$ always denotes the $i$ th largest singular value. We often just write $\sigma_{i}$.
- $\operatorname{diag}(A)$ : the vector of diagonal entries of $A$.
- We use capital letters for matrices, lower-case for vectors and scalars. Unless otherwise specified, $A$ is a given matrix, $b$ is a given vector, and $x$ is an unknown vector.
- $\|\cdot\|$ denotes a norm for a vector or matrix. $\|\cdot\|_{2}$ denotes the spectral (or 2-) norm, $\|\cdot\|_{F}$ the Frobenius norm. For vectors, to simplify notation we sometimes use $\|\cdot\|$ for the spectral norm (which for vectors is the familiar Euclidean norm).
- $\operatorname{Span}(A)$ denotes the span or range of the column space of $A$. This is the subspace consisting of vectors of the form $A x$.
- We reserve $Q$ for an orthonormal (or orthogonal) matrix. $L,(U)$ are often lower (upper) triangular.
- $I$ always denotes the identity matrix. $I_{n}$ is the $n \times n$ identity when the size needs to be specified.
- $A^{T}$ is the transpose of the matrix; $\left(A^{T}\right)_{i j}=A_{j i} . \quad A^{*}$ is the (complex) conjugate transpose $\left(A^{*}\right)_{i j}=\overline{A_{j i}}$.
- $\succ, \succeq$ denote the positive (semi)definite ordering. That is, $A \succ(\succeq) 0$ means $A$ is positive (semi)definite (abbreviated as PD, PSD), i.e., symmetric and with positive (nonnegative) eigenvalues. $A \succ B$ means $A-B \succ 0$.

We sometimes use the following shorthand: alg for algorithm, eigval for eigenvalue, eigvec for eigenvector, singval for singular value, and singvec for singular vector, iff for "if and only if".

## 0 Introduction, why $A x=b$ and $A x=\lambda x$ ?

As already stated, NLA is the study of numerical algorithms for problems involving matrices, and there are only two main problems(!):

1. Linear system

$$
A x=b \text {. }
$$

Given a (often square $m=n$ but we will discuss $m>n$ extensively, and $m<n$ briefly at the end) matrix $A \in \mathbb{R}^{m \times n}$ and vector $b \in \mathbb{R}^{m}$, find $x \in \mathbb{R}^{n}$ such that $A x=b$.
2. Eigenvalue problem

$$
A x=\lambda x .
$$

Given a (always! ${ }^{1}$ ) square matrix $A \in \mathbb{R}^{n \times n}$ find $\lambda$ : eigenvalues (eigval), and $x \in \mathbb{R}^{n}$ : eigenvectors (eigvec).

We'll see many variants of these problems; one worthy of particular mention is the SVD, which is related to eigenvalue problems but given its ubiquity has a life of its own. (So if there's a third problem we solve in NLA, it would definitely be the SVD.)

It is worth discussing why we care about linear systems and eigenvalue problems.
The primary reason is that many (in fact most) problems in scientific computing (and even machine learning) boil down to linear problems:

- Because that's often the only way to deal with the scale of problems we face today! (and in future)

[^1]- For linear problems, so much is understood and reliable algorithms are available ${ }^{2}$.

A related important question is where and how these problems arise in real-world problems.

Let us mention a specific context that is relevant in data science: optimisation. Suppose one is interested in minimising a high-dimensional real-valued function $f(x): \mathbb{R}^{n} \rightarrow \mathbb{R}$ where $n \gg 1$.

A successful approach is to try and find critical points, that is, points $x_{*}$ where $\nabla f\left(x_{*}\right)=$ 0. Mathematically, this is a non-linear high-dimensional root-finding problem of finding $x \in \mathbb{R}^{n}$ such that $\nabla f(x)=: F(x)=0$ (the vector $0 \in \mathbb{R}^{n}$ ) where $F: \mathbb{R}^{n} \rightarrow \mathbb{R}^{n}$. One of the most commonly employed methods for this task is Newton's method (which some of you have seen in Prelims Constructive Mathematics). This boils down to

- Newton's method for $F(x)=0, F: \mathbb{R}^{n} \rightarrow \mathbb{R}^{n}$ nonlinear:

1. Start with initial guess $x^{(0)} \in \mathbb{R}^{n}$, set $i=0$
2. Find Jacobian matrix $J \in \mathbb{R}^{n \times n}, J_{i j}=\left.\frac{\partial F_{i}(x)}{\partial x_{j}}\right|_{x=x^{(0)}}$
3. Update $x^{(i+1)}:=x^{(i)}-J^{-1} F\left(x^{(i)}\right), i \leftarrow i+1$, go to step 2 and repeat

Note that the main computational task is to find the vector $y=J^{-1} F\left(x^{(i)}\right)$, which is a linear system $J y=F\left(x^{(i)}\right)$ (which we solve for the vector $y$ )

What about eigenvalue problems $A x=\lambda x$ ? Google's pagerank is a famous application (we will cover this if we have time). Another example the Schrödinger equation of physics and chemistry. Sometimes a nonconvex optimisation problem can be solved by an eigenvalue problem.

Equally important is principal component analysis (PCA), which can be used for data compression. This is more tightly connected to the SVD.

Other sources of linear algebra problems include differential equations, optimisation, regression, data analysis, ...

## 1 Basic LA review

We start with a review of key LA facts that will be used in the course. Some will be trivial to you while others may not. You might also notice that some facts that you have learned in a core LA course will not be used in this course. For example we will never deal with finite fields, and determinants only play a passing role.

[^2]
### 1.1 Warmup exercise

Let $A \in \mathbb{R}^{n \times n}$ ( $n \times n$ square matrix). (or $\mathbb{C}^{n \times n}$; the difference hardly matters in most of this course ${ }^{3}$ ). Try to think of statements that are equivalent to $A$ being nonsingular. Try to come up with as many conditions as possible, before turning the page.

[^3]Here is a list: The following are equivalent.

1. $A$ is nonsingular.
2. $A$ is invertible: $A^{-1}$ exists.
3. The map $A: \mathbb{R}^{n} \rightarrow \mathbb{R}^{n}$ is a bijection.
4. All $n$ eigenvalues of $A$ are nonzero.
5. All $n$ singular values of $A$ are positive.
6. $\operatorname{rank}(A)=n$.
7. The rows of $A$ are linearly independent.
8. The columns of $A$ are linearly independent.
9. $A x=b$ has a solution for every $b \in \mathbb{C}^{n}$.
10. $A$ has no nonzero null vector.
11. $A^{T}$ has no nonzero null vector.
12. $A^{*} A$ is positive definite (not just semidefinite).
13. $\operatorname{det}(A) \neq 0$.
14. An $n \times n$ matrix $A^{-1}$ exists such that $A^{-1} A=I_{n}$. (this, btw, implies (iff) $A A^{-1}=I_{n}$, a nontrivial fact)
15. ... (what did I miss?)

### 1.2 Structured matrices

We will be discussing lots of structured matrices. For square matrices,

- Symmetric: $A_{i j}=A_{j i}$ (Hermitian: $A_{i j}=\overline{A_{j i}}$ )
- The most important property of symmetric matrices is the symmetric eigenvalue decomposition $A=V \Lambda V^{T} ; V$ is orthogonal $V^{T} V=V V^{T}=I_{n}$, and $\Lambda$ is a diagonal matrix of eigenvalues $\Lambda=\operatorname{diag}\left(\lambda_{1}, \ldots, \lambda_{n}\right)$.
- symmetric positive (semi)definite $A \succ(\succeq) 0$ : symmetric and all positive (nonnegative) eigenvalues.
- Orthogonal: $A A^{T}=A^{T} A=I$ (Unitary: $A A^{*}=A^{*} A=I$ ). Note that for square matrices, $A^{T} A=I$ implies $A A^{T}=I$.
- Skew-symmetric: $A_{i j}=-A_{j i}$ (skew-Hermitian: $A_{i j}=-\overline{A_{j i}}$ ).
- Normal: $A^{T} A=A A^{T}$. (Here it's better to discuss the complex case $A^{*} A=A A^{*}$ : this is a necessary and sufficient condition for diagonalisability under a unitary transformation, i.e., $A=U \Lambda U^{*}$ where $\Lambda$ is diagonal and $U$ is unitary.)
- Tridiagonal: $A_{i j}=0$ if $|i-j|>1$.
- Upper triangular: $A_{i j}=0$ if $i>j$.
- Lower triangular: $A_{i j}=0$ if $i<j$.

For (possibly nonsquare) matrices $A \in \mathbb{C}^{m \times n}$, (usually $m \geq n$ ).

- (upper) Hessenberg: $A_{i j}=0$ if $i>j+1$. (we will see this structure often.)
- "orthonormal": $A^{T} A=I_{n}$, and $A$ is (tall) rectangular. (This isn't an established name - we could call it "matrix with orthonormal columns" every time it appearsbut we use these matrices all the time in this course, so we need a consistent shorthand name for it.)
- sparse: most elements are zero. $\mathrm{nnz}(A)$ denotes the number of nonzero elements in $A$. Matrices that are not sparse are called dense.
Other structures: Hankel, Toeplitz, circulant, symplectic,... (we won't use these in this course)


### 1.3 Matrix eigenvalues: basics

$$
A x=\lambda x, \quad A \in \mathbb{R}^{n \times n},(0 \neq) x \in \mathbb{R}^{n}
$$

- Example: $\left[\begin{array}{lll}2 & 1 & 1 \\ 1 & 2 & 1 \\ 1 & 1 & 2\end{array}\right]\left[\begin{array}{l}1 \\ 1 \\ 1\end{array}\right]=4\left[\begin{array}{l}1 \\ 1 \\ 1\end{array}\right]$

This matrix has an eigenvalue $\lambda=4$, with corresponding eigenvector $x=\left[\begin{array}{l}1 \\ 1 \\ 1\end{array}\right]$ (together they are an eigenpair).
An $n \times n$ matrix always has $n$ eigenvalues (not always $n$ linearly independent eigenvectors); In the example above, $(\lambda, x)=\left(1,\left[\begin{array}{c}1 \\ -1 \\ 0\end{array}\right]\right),\left(1,\left[\begin{array}{c}0 \\ 1 \\ -1\end{array}\right]\right)$ are also eigenpairs.

- The eigenvalues are the roots of the characteristic polynomial $\operatorname{det}(\lambda I-A)=0: \operatorname{det}(\lambda I-$ $A)=\prod_{i=1}^{n}\left(\lambda-\lambda_{i}\right)$.
- According to Galois theory, eigenvalues cannot be computed exactly for matrices with $n \geq 5$. But we still want to compute them! In this course we will (among other things) explain how this is done in practice by the $Q R$ algorithm, one of the greatest hits of the field.


### 1.4 Computational complexity (operation counts) of matrix algorithms

Since NLA is a field that aspires to develop practical algorithms for solving matrix problems, it is important to be aware of the computational cost (often referred to as complexity) of the algorithms. We will discuss these as the algorithms are developed, but for now let's examine the costs for basic matrix-matrix multiplication. The cost is measured in terms of flops (floating-point operations), which counts the number of additions, subtractions, multiplications, and divisions (all treated equally) performed.

In NLA the constant in front of the leading term in the cost is (clearly) important. It is customary (for good reason) to only track the leading term of the cost. For example, $n^{3}+10 n^{2}$ is abbreviated to $n^{3}$.

- Multiplying two $n \times n$ matrices $A B$ costs $2 n^{3}$ flops. More generally, if $A$ is $m \times n$ and $B$ is $n \times k$, then computing $A B$ costs $2 m n k$ flops.
- Multiplying a vector to an $m \times n$ matrix $A$ costs $2 m n$ flops.


## Norms

We will need a tool (or metric) to measure how big a vector or matrix is. Norms give us a means to achieve this. Surely you have already seen some norms (e.g. the vector Euclidean norm). We will discuss a number of norms for vectors and matrices that we will use in the upcoming lectures.

### 1.5 Vector norms

For vectors $x=\left[x_{1}, \ldots, x_{n}\right]^{T} \in \mathbb{C}^{n}$

- $p$-norm $\|x\|_{p}=\left(\left|x_{1}\right|^{p}+\left|x_{2}\right|^{p}+\cdots+\left|x_{n}\right|^{p}\right)^{1 / p} \quad(1 \leq p \leq \infty)$
- Euclidean norm $=2$-norm $\|x\|_{2}=\sqrt{\left|x_{1}\right|^{2}+\left|x_{2}\right|^{2}+\cdots+\left|x_{n}\right|^{2}}$
- 1-norm $\|x\|_{1}=\left|x_{1}\right|+\left|x_{2}\right|+\cdots+\left|x_{n}\right|$
$-\infty$-norm $\|x\|_{\infty}=\max _{i}\left|x_{i}\right|$
Of particular importance are the three cases $p=1,2, \infty$. In this course, we will see $p=2$ the most often.

A norm needs to satisfy the following axioms:

- $\|\alpha x\|=|\alpha|\|x\|$ for any $\alpha \in \mathbb{C}$ (homogeneity),
- $\|x\| \geq 0$ and $\|x\|=0 \Leftrightarrow x=0$ (nonnegativity),
- $\|x+y\| \leq\|x\|+\|y\|$ (triangle inequality).

The vector $p$-norm satisfies all these, for any $p$.
Here are some useful inequalities for vector norms. A proof is left for your exercise and is highly recommended. (Try to think when each equality is satisfied.) For $x \in \mathbb{C}^{n}$,

- $\frac{1}{\sqrt{n}}\|x\|_{2} \leq\|x\|_{\infty} \leq\|x\|_{2}$
- $\frac{1}{\sqrt{n}}\|x\|_{1} \leq\|x\|_{2} \leq\|x\|_{1}$
- $\frac{1}{n}\|x\|_{1} \leq\|x\|_{\infty} \leq\|x\|_{1}$

Note that with the 2-norm, $\|U x\|_{2}=\|x\|_{2}$ for any unitary $U$ and any $x \in \mathbb{C}^{n}$. Norms with this property are called unitarily invariant.

The 2-norm is also induced by the inner product $\|x\|_{2}=\sqrt{x^{T} x}$. An important property of inner products is the Cauchy-Schwarz inequality $\left|x^{T} y\right| \leq\|x\|_{2}\|y\|_{2}$ (which can be directly proved but is perhaps best to prove in general setting) ${ }^{4}$. When we just say $\|x\|$ for a vector we mean the 2-norm.

### 1.6 Matrix norms

We now turn to norms of matrices. As you will see, many (but not the Frobenius and trace norms) are defined via the vector norms (these are called induced norms).

- $p$-norm $\|A\|_{p}=\max _{x} \frac{\|A x\|_{p}}{\|x\|_{p}}$
- 2-norm=spectral norm(=Euclidean norm) $\|A\|_{2}=\sigma_{\max }(A)$ (largest singular value; see Section 2)
- 1-norm $\|A\|_{1}=\max _{i} \sum_{j=1}^{m}\left|A_{j i}\right|$
$-\infty$-norm $\|A\|_{\infty}=\max _{i} \sum_{j=1}^{n}\left|A_{i j}\right|$
- Frobenius norm $\|A\|_{F}=\sqrt{\sum_{i} \sum_{j}\left|A_{i j}\right|^{2}}$ (2-norm of vectorisation)
- trace norm=nuclear norm $\|A\|_{*}=\sum_{i=1}^{\min (m, n)} \sigma_{i}(A)$. (this is the maximum trace of $Q^{T} A$ where $Q$ is orthonormal, hence the name)

Colored in red are unitarily invariant norms $\|A\|_{*}=\|U A V\|_{*},\|A\|_{F}=\|U A V\|_{F},\|A\|_{2}=$ $\|U A V\|_{2}$ for any unitary/orthogonal $U, V$.

Norm axioms hold for each of these. Useful inequalities include: For $A \in \mathbb{C}^{m \times n}$, (exercise; it is instructive to study the cases where each of these equalities holds)

- $\frac{1}{\sqrt{n}}\|A\|_{\infty} \leq\|A\|_{2} \leq \sqrt{m}\|A\|_{\infty}$

[^4]- $\frac{1}{\sqrt{m}}\|A\|_{1} \leq\|A\|_{2} \leq \sqrt{n}\|A\|_{1}$
- $\|A\|_{2} \leq\|A\|_{F} \leq \sqrt{\min (m, n)}\|A\|_{2}$

A useful property of $p$-norms is that they are subordinate, i.e., $\|A B\|_{p} \leq\|A\|_{p}\|B\|_{p}$ (problem sheet). Note that not all norms satisfy this, e.g. with the max norm $\|A\|_{\max }=$ $\max _{i, j}\left|A_{i j}\right|$, with $A=[1,1]$ and $B=\left[\begin{array}{l}1 \\ 1\end{array}\right]$ one has $\|A B\|_{\max }=2$ but $\|A\|_{\max }=\|B\|_{\max }=1$.

### 1.7 Subspaces and orthonormal matrices

A key notion that we will keep using throughout is a subspace $\mathcal{S}$. In this course we will almost exclusively confine ourselves to subspaces of $\mathbb{R}^{n}$, even though they generalize to more abstract vector spaces. A subspace is the set of vectors that can be written as a linear combination basis vectors $v_{1}, \ldots, v_{d}$, which are assumed to be linearly independent (otherwise there is a basis with fewer vectors). That is, $x \in \mathcal{S}$ iff $\sum_{i=1}^{d} c_{i} v_{i}$ (where $c_{i}$ are scalars, $\mathbb{R}$ or $\mathbb{C}$ ). The integer $d$ is called the dimension of the subspace. We also say the subspace is spanned by the vectors $v_{1}, \ldots, v_{d}$, or that $v_{1}, \ldots, v_{d}$ spans the subpsace.

How does one represent a subspace? An obvious answer is to use the basis vectors $v_{1}, \ldots, v_{d}$. This sometimes becomes cumbersome, and a common and convenient way to represent the subspace is to use a (tall-skinny) rectangular matrix $V \in \mathbb{R}^{n \times d}=\left[v_{1}, v_{2}, \ldots, v_{d}\right]$, as $\mathcal{S}=\operatorname{span}(V)$ (or sometimes just "subspace $V$ ") which means the subspace of vectors that can be written as $V c$, where $c$ is a 'coefficient' vector $c \in \mathbb{R}^{d}$.

It will be (not necessary but) convenient to represent subspaces using an orthonormal matrix $Q \in \mathbb{R}^{n \times d}$. (once we cover the QR factorisation, you'll see that there is no loss of generality in doing so).

An important fact about subspaces of $\mathbb{R}^{n}$ is the following:
Lemma 1.1 Let $V_{1} \in \mathbb{R}^{n \times d_{1}}$ and $V_{2} \in \mathbb{R}^{n \times d_{2}}$ each have linearly independent column vectors. If $d_{1}+d_{2}>n$, then there is a nonzero intersection between two subspaces $\mathcal{S}_{1}=\operatorname{span}\left(V_{1}\right)$ and $\mathcal{S}_{2}=\operatorname{span}\left(V_{2}\right)$, that is, there is a nonzero vector $x \in \mathbb{R}^{n}$ such that $x=V_{1} c_{1}=V_{2} c_{2}$ for some vectors $c_{1}, c_{2}$.

This is straightforward but important enough to warrant a proof.
Proof: Consider the matrix $M:=\left[V_{1}, V_{2}\right]$, which is of size $n \times\left(d_{1}+d_{2}\right)$. Since $d_{1}+d_{2}>n$ by assumption, this matrix has a right null vector ${ }^{5} c \neq 0$ such that $M c=0$. Splitting $c=\left[\begin{array}{c}c_{1} \\ -c_{2}\end{array}\right]$ we have the required result.

Let us conclude this review with a list of useful results that will be helpful. Proofs (or counterexample) should be straightforward.

[^5]- $(A B)^{T}=B^{T} A^{T}$
- If $A, B$ invertible, $(A B)^{-1}=B^{-1} A^{-1}$
- If $A, B$ square and $A B=I$, then $B A=I$
- $\left[\begin{array}{cc}I_{m} & X \\ 0 & I_{n}\end{array}\right]^{-1}=\left[\begin{array}{cc}I_{m} & -X \\ 0 & I_{n}\end{array}\right]$
- Neumann series: if $\|X\|<1$ in any norm,

$$
(I-X)^{-1}=I+X+X^{2}+X^{3}+\cdots
$$

- For a square $n \times n$ matrix $A$, the trace is $\operatorname{Trace}(A)=\sum_{i=1}^{n} A_{i, i}$ (sum of diagonals). For any $X, Y$ such that $X Y$ is square, $\operatorname{Trace}(X Y)=\operatorname{Trace}(Y X)$ (quite useful). For $B \in \mathbb{R}^{m \times n}$, we have $\|B\|_{F}^{2}=\sum_{i} \sum_{j}\left|B_{i j}\right|^{2}=\operatorname{Trace}\left(B^{T} B\right)$.
- Triangular structure (upper or lower) is invariant under addition, multiplication, and inversion. That is, triangular matrices form a ring (in abstract algebra; don't worry if this is foreign to you).
- Symmetry is invariant under addition and inversion, but not multiplication; $A B$ is usually not symmetric even if $A, B$ are.


## 2 SVD: the most important matrix decomposition

We now start the discussion of the most important topic of the course: the singular value decomposition (SVD). The SVD exists for any matrix, square or rectangular, real or complex.

We will prove its existence and discuss its properties and applications in particular in low-rank approximation, which can immediately be used for compressing the matrix, and therefore data.

The SVD has many intimate connections to symmetric eigenvalue problems. Let's start with a review.

- Symmetric eigenvalue decomposition: Any symmetric matrix $A \in \mathbb{R}^{n \times n}$ has the decomposition

$$
\begin{equation*}
A=V \Lambda V^{T} \tag{1}
\end{equation*}
$$

where $V$ is orthogonal, $V^{T} V=I_{n}=V V^{T}$, and $\Lambda=\operatorname{diag}\left(\lambda_{1}, \ldots, \lambda_{n}\right)$ is a diagonal matrix of eigenvalues.
$\lambda_{i}$ are the eigenvalues, and $V$ is the matrix of eigenvectors (its columns are the eigenvectors).

The decomposition (1) makes two remarkable claims: the eigenvectors can be taken to be orthogonal (which is true more generally of normal matrices s.t. $A^{*} A=A A^{*}$ ), and the eigenvalues are real.

It is worth reminding you that eigenvectors are not uniquely determined: (assuming they are normalised s.t. the 2-norm is 1) their signs can always be flipped. (And actually more, when there are eigenvalues that are multiple $\lambda_{i}=\lambda_{j}$; in this case the eigenvectors span a subspace whose dimension matches the multiplicity. For example, any vector is an eigenvector of the identity matrix $I$ ).

Now here is the protagonist of this course: Be sure to spend dozens (if not hundreds) of hours thinking about it!

Theorem 2.1 (Singular Value Decomposition (SVD)) Any matrix $A \in \mathbb{R}^{m \times n}$ has the decomposition :

$$
\begin{equation*}
A=U \Sigma V^{T} \tag{2}
\end{equation*}
$$

Here $U^{T} U=V^{T} V=I_{n}$ (assuming $m \geq n$ for definiteness), $\Sigma=\operatorname{diag}\left(\sigma_{1}, \ldots, \sigma_{n}\right), \sigma_{1} \geq \sigma_{2} \geq$ $\cdots \geq \sigma_{n} \geq 0$.
$\sigma_{i}$ (always nonnegative) are called the singular values of $A$. The rank of $A$ is the number of positive singular values. The columns of $U$ are called the left singular vectors, and the columns of $V$ are the right singular vectors.

Writing $U=\left[u_{1}, \ldots, u_{n}\right]$ and $V=\left[v_{1}, \ldots, v_{n}\right]$, we have an important alternative expression for $A: A=\sum_{i=1}^{n} \sigma_{i} u_{i} v_{i}^{T}$. We will use this expression repeatedly in what follows. Also note that we always order the singular values in nonincreasing order, so $\sigma_{1}$ is always the largest singular value.

The SVD tells us that any (tall) matrix can be written as orthonormal-diagonal-orthogonal. Roughly, ortho-normal/gonal matrices can be thought of as rotations or reflection, so the SVD says the action of a matrix can be thought of as a rotation/reflection followed by magnification (or shrinkage), followed by another rotation/reflection.

Proof: For $\mathrm{SVD}^{6}$ ( $m \geq n$ and assume full-rank $\sigma_{n}>0$ for simplicity): Take Gram matrix $A^{T} A$ (symmetric) and its eigendecomposition $A^{T} A=V \Lambda V^{T}$ with $V$ orthogonal. $\Lambda$ is nonnegative, and $(A V)^{T}(A V)=: \Sigma^{2}$ is diagonal, so $A V \Sigma^{-1}=: U$ is orthonormal. Right-multiply by $\Sigma V^{T}$ to get $A=U \Sigma V^{T}$.

It is also worth mentioning the "full" SVD: $A=\left[U, U_{\perp}\right]\left[\begin{array}{l}\Sigma \\ 0\end{array}\right] V^{T}$ where $\left[U, U_{\perp}\right] \in \mathbb{R}^{m \times m}$ is square and orthogonal. $U_{\perp} \in \mathbb{R}^{m \times m-n}$ (pronounced 'U-perp', for perpendicular) is an orthonormal matrix such that $U^{T} U_{\perp}=0$ (whose existence and construction can be established via the Householder QR factorsation in Section 6). Essentially the full SVD follows from the (thin) SVD (2) by filling in $U$ in (2) with its orthogonal complement $U_{\perp}$.

[^6]Example 2.1 Let's find the $S V D$ of $A=\left[\begin{array}{cc}-1 & -2 \\ 2 & 1 \\ 1 & 0 \\ 0 & 1\end{array}\right]$. We can follow the proof and compute the Gram matrix $A^{T} A=\left[\begin{array}{ll}6 & 4 \\ 4 & 6\end{array}\right]$. Its eigenvalues are 10 and 2 (via the characteristic polynomial $(\lambda-6)(\lambda-6)-4^{2}=0$-please note though that this isn't how they should be computed on a computer (the QR algorithm), which we'll cover later.) An eigenvector matrix can be found as $V=\frac{1}{\sqrt{2}}\left[\begin{array}{cc}1 & -1 \\ 1 & 1\end{array}\right]$ (degree of freedom in the signs of the columns), computed e.g. via the null vector of $(A-\lambda I)$. Hence $A^{T} A=V \Sigma^{2} V^{T}$ where $\Sigma^{2}=\left[\begin{array}{ll}10 & \\ & 2\end{array}\right]$. Let $U=A V \Sigma^{-1}=\left[\begin{array}{cc}-3 / \sqrt{20} & -1 / 2 \\ 3 / \sqrt{20} & -1 / 2 \\ 1 / \sqrt{20} & -1 / 2 \\ 1 / \sqrt{20} & 1 / 2\end{array}\right]$, which is orthonormal. We thus have $A=U \Sigma V^{T}$.

## 2.1 (Some of the many) applications and consequences of the SVD: rank, column/row space, etc

From the SVD one can immediately read off a number of important properties of the matrix. For example:

- The rank $r$ of $A \in \mathbb{R}^{m \times n}$, often denoted $\operatorname{rank}(A)$ : this is the number of nonzero (positive) singular values $\sigma_{i}(A) \cdot \operatorname{rank}(A)$ is also equal to the number of linearly independent rows or the number of independent columns, as you probably learnt in your first course on linear algebra (exercise).
- We can always write $A=\sum_{i=1}^{\operatorname{rank}(A)} \sigma_{i} u_{i} v_{i}^{T}$.

Important: An $m \times n$ matrix $A$ that is of rank $r$ can be written as an (outer) product of $m \times r$ and $r \times n$ matrices:

$$
A_{r}=\boxed{U_{r}} \boxed{\Sigma_{r}} \begin{array}{|}
V_{r}^{T} \\
& \\
\end{array}
$$

To see this, note from the SVD that $A=\sum_{i=1}^{n} \sigma_{i} u_{i} v_{i}^{T}=\sum_{i=1}^{r} \sigma_{i} u_{i} v_{i}^{T}$ (since $\sigma_{r+1}=0$ ), and so $A=U_{r} \Sigma_{r} V_{r}^{T}$ where $U_{r}=\left[u_{1}, \ldots, u_{r}\right], V_{r}=\left[v_{1}, \ldots, v_{r}\right]$, and $\Sigma_{r}$ is the leading $r \times r$ submatrix of $\Sigma$.

- Column space (Span(A), linear subspace spanned by vectors $A x)$ : span of $U=\left[u_{1}, \ldots, u_{r}\right]$, often denoted $\operatorname{Span}(U)$.
- Row space $\left(\operatorname{Span}\left(A^{T}\right)\right)$ : row span of $v_{1}^{T}, \ldots, v_{r}^{T}$. $\operatorname{Span}(V)$.
- Null space $\operatorname{Null}(\mathrm{A}):$ span of $v_{r+1}, \ldots, v_{n}$; as $A v=0$ for these vectors. Null(A) is empty (or just the 0 vector) if $m \geq n$ and $r=n$. (When $m<n$ the full SVD is needed to describe $\operatorname{Null(A).)~}$
- We have $A v_{i}=\sigma_{i} u_{i}$, and $u_{i}^{T} A=\sigma_{i} v_{i}$. If $\sigma_{i}, u_{i}, v_{i}$ satisfy this pair of equations, they are called a (the $i$ th) singular triplet of $A$.

Aside from these and other applications, the SVD is also a versatile theoretical tool. Very often, a good place to start in proving a fact about matrices is to first consider its SVD; you will see this many times in this course. For example, the SVD can give solutions immediately for linear systems and least-squares problems, though there are more efficient ways to solve these problems.

### 2.2 SVD and symmetric eigenvalue decomposition

As mentioned above, the SVD and the eigenvalue decomposition for symmetric matrices are closely connected. Here are some results that highlight the connections between $A=U \Sigma V^{T}$ and symmetric eigenvalue decomposition. (We assume $m \geq n$ for definiteness)

- $V$ is an eigvector matrix of $A^{T} A$. (To verify, see proof of SVD)
- $U$ is an eigvector matrix (for nonzero eigvals) of $A A^{T}$ (be careful with sign flips; see below)
- $\sigma_{i}=\sqrt{\lambda_{i}\left(A^{T} A\right)}$ for $i=1, \ldots, n$
- If $A$ is symmetric, its singular values $\sigma_{i}(A)$ are the absolute values of its eigenvalues $\lambda_{i}(A)$, i.e., $\sigma_{i}(A)=\left|\lambda_{i}(A)\right|$.
Exercise: What if $A$ is unitary, skew-symmetric, normal matrices, triangular? (problem sheet)
- Jordan-Wieldant matrix $\left[\begin{array}{cc}0 & A \\ A^{T} & 0\end{array}\right]$ : This matrix has eigenvalues $\pm \sigma_{i}(A)$, and $m-n$ copies of eigenvalues at 0 . Its eigenvector matrix is $\left[\begin{array}{ccc}U & U & U_{\perp} \\ V & -V & 0\end{array}\right]$, where $A^{T} U_{\perp}=0\left(U_{\perp}\right.$ is empty when $m=n$ ). This matrix, along with the Gram matrix $A^{T} A$, is a very useful tool when one tries to extend a result on symmetric eigenvalues to an analogue in terms of the SVD (or vice versa).


### 2.3 Uniqueness etc

We have established the existence of the SVD $A=U \Sigma V^{T}$. A natural question is: is it unique? In other words, are the factors $U, \Sigma, V$ uniquely determined by $A$ ?

It is straightforward to see that the singular vectors are not uniquely determined. Most obviously, the singular vectors can be flipped in signs, just like eigenvectors. However, note
that the signs of $u_{i}, v_{i}$ are not entirely arbitrary: if we replace $u_{i}$ by $-u_{i}$, the same needs to be done for $v_{i}$ in order to satisfy $A=\sum_{i=1}^{n} \sigma_{i} u_{i} v_{i}^{T}$. Essentially, once the sign (or rotation) of $u_{i}$ (or $v_{i}$ ) is fixed, $v_{i}$ (or $u_{i}$ ) is determined uniquely.

More generally, in the presence of multiple singular values (i.e., $\sigma_{i}=\sigma_{i+1}$ for some $i$ ), there is a higher degree of freedom in the SVD. Again this is very much analogous to eigenvalues (recall the discussion on eigenvectors of the identity). Here think of what the SVD is for an orthogonal matrix: there is an enormous amount of degrees of freedom in the choice of $U$ and $V$. The singular values $\sigma_{i}$, on the other hand, are always unique, as they are the eigenvalues of the Gram matrix.

## 3 Low-rank approximation via truncated SVD

While the SVD has a huge number of applications, undoubtedly the biggest reason that makes it so important in computational mathematics is its optimality for low-rank approximation.

To discuss this topic we need some preparation. We will make heavy use of the spectral norm of matrices. We start with an important characterisation of $\|A\|_{2}$ in terms of the singular value(s), which we previously stated but did not prove.

## Proposition 3.1

$$
\|A\|_{2}=\max _{x} \frac{\|A x\|_{2}}{\|x\|_{2}}=\max _{\|x\|_{2}=1}\|A x\|_{2}=\sigma_{1}(A)
$$

Proof: Use the SVD: For any $x$ with unit norm $\|x\|_{2}=1$,

$$
\begin{aligned}
\|A x\|_{2} & =\left\|U \Sigma V^{T} x\right\|_{2} \\
& =\left\|\Sigma V^{T} x\right\|_{2} \quad \text { by unitary invariance } \\
& =\|\Sigma y\|_{2} \quad \text { with }\|y\|_{2}=1 \\
& =\sqrt{\sum_{i=1}^{n} \sigma_{i}^{2} y_{i}^{2}} \\
& \leq \sqrt{\sum_{i=1}^{n} \sigma_{1}^{2} y_{i}^{2}}=\sigma_{1}\|y\|_{2}^{2}=\sigma_{1}
\end{aligned}
$$

Finally, note that taking $x=v_{1}$ (the leading right singular vector), we have $\left\|A v_{1}\right\|_{2}=\sigma_{1}$.
Similarly, the Frobenius norm can be expressed as $\|A\|_{F}=\sqrt{\sum_{i} \sum_{j}\left|A_{i j}\right|^{2}}=\sqrt{\sum_{i=1}^{n}\left(\sigma_{i}(A)\right)^{2}}$, and the trace norm is (by definition) $\|A\|_{*}=\sum_{i=1}^{\min (m, n)} \sigma_{i}(A)$. (exercise) In general, norms that are unitarily invariant can be characterised by the singular values [20].

Now to the main problem: Given $A \in \mathbb{R}^{m \times n}$, consider the problem of finding a rank-r matrix (remember; these are matrices with $r$ nonzero singular values) $A_{r} \in \mathbb{R}^{m \times n}$ that best approximates $A$. That is, find the minimiser $A_{r}$ for

$$
\begin{equation*}
\operatorname{argmin}_{\operatorname{rank}\left(A_{r}\right) \leq r}\left\|A-A_{r}\right\|_{2} \tag{3}
\end{equation*}
$$

It is definitely worth visualizing the situation:


We immediately see that a low rank approximation (when possible) is beneficial in terms of the storage cost when $r \ll m, n$. Instead of storing $m n$ entries for $A$, we can store entries for $U_{r}, \Sigma_{r}, V_{r}((m+n+1) r$ entries $)$ to keep the low-rank factorisation without losing information.

Low-rank approximation can also bring computational benefits. For example, in order to compute $A x$ for a vector $x$, by noting that $A_{r} x=U_{r}\left(\Sigma_{r}\left(V_{r}^{T} x\right)\right)$, one needs only $O((m+n) r)$ operations ${ }^{7}$ instead of $O(m n)$. The utility and prevalence of low-rank matrices in data science is remarkable.

Here is the solution for (3): Truncated SVD, defined via $A_{r}=\sum_{i=1}^{r} \sigma_{i} u_{i} v_{i}^{T}\left(=U_{r} \Sigma_{r} V_{r}^{T}\right)$. This is the matrix obtained by truncating (removing) the trailing terms in the expression $A=\sum_{i=1}^{n} \sigma_{i} u_{i} v_{i}^{T}$. Pictorially,

$$
\begin{aligned}
& A=\underbrace{\left[\begin{array}{c}
* \\
* \\
\vdots \\
* \\
*
\end{array}\right]\left[\begin{array}{lllll}
* & * & \cdots & * & *
\end{array}\right]}_{\sigma_{1} u_{1} v_{1}}+\underbrace{\left[\begin{array}{c}
* \\
* \\
\vdots \\
* \\
*
\end{array}\right]\left[\begin{array}{lllll}
* & * & \cdots & * & *
\end{array}\right]}_{\sigma_{2} u_{2} v_{2}}+\cdots+\underbrace{\left[\begin{array}{c}
* \\
* \\
\vdots \\
* \\
*
\end{array}\right]\left[\begin{array}{lllll}
* & * & \cdots & * & *
\end{array}\right]}_{\sigma_{n} u_{n} v_{n}} \text {, } \\
& A_{r}=\underbrace{\left[\begin{array}{c}
* \\
* \\
\vdots \\
* \\
*
\end{array}\right]\left[\begin{array}{lllll}
* & * & \cdots & * & *
\end{array}\right]}_{\sigma_{1} u_{1} v_{1}}+\cdots+\underbrace{\left[\begin{array}{c}
* \\
* \\
\vdots \\
* \\
*
\end{array}\right]\left[\begin{array}{lllll}
* & * & \cdots & * & *
\end{array}\right] .}_{\sigma_{r} u_{r} v_{r}}
\end{aligned}
$$

In particular, we have
Theorem 3.1 For any $A \in \mathbb{R}^{m \times n}$ with singular values $\sigma_{1}(A) \geq \sigma_{2}(A) \geq \cdots \geq \sigma_{n}(A) \geq 0$, and any nonnegative integer ${ }^{8} r<\min (m, n)$,

$$
\begin{equation*}
\left\|A-A_{r}\right\|_{2}=\sigma_{r+1}(A)=\min _{\operatorname{rank}(B) \leq r}\|A-B\|_{2} \tag{4}
\end{equation*}
$$

Before proving this result, let us make some observations.

- Good approximation $A \approx A_{r}$ is obtained iff $\sigma_{r+1} \ll \sigma_{1}$.

[^7]- Optimality holds for any unitarily invariant norm: that is, the norms in (3) can be replaced by e.g. the Frobenius norm. This is surprising, as the low-rank approximation problem $\min _{\operatorname{rank}(B) \leq r}\|A-B\|$ does depend on the choice of the norm (and for many problems, including the least-squares problem in Section 6.5, the norm choice has a significant effect on the solution). The proof for this fact is nonexaminable, but if curious see [20] for a complete proof.
- A prominent application of low-rank approximation is PCA (principal component analysis) in statistics and data science.
- Many matrices have explicit or hidden low-rank structure (nonexaminable, but see e.g. [34]).

Proof: of Theorem 3.1:

1. Since $\operatorname{rank}(B) \leq r$, we can write $B=B_{1} B_{2}^{T}$ where $B_{1}, B_{2}$ have $r$ columns.
2. It follows that $B_{2}^{T}$ (and hence $B$ ) has a null space of dimension at least $n-r$. That is, there exists an orthonormal matrix $W \in \mathbb{C}^{n \times(n-r)}$ s.t. $B W=0$. Then $\|A-B\|_{2} \geq$ $\|(A-B) W\|_{2}=\|A W\|_{2}=\left\|U \Sigma\left(V^{T} W\right)\right\|_{2}$. (Why does the first inequality hold?)
3. Now since $W$ is $(n-r)$-dimensional, by Lemma 1.1 there is an intersection between $\operatorname{Span}(W)$ and $\operatorname{Span}\left(\left[v_{1}, \ldots, v_{r+1}\right]\right)$, the $(r+1)$-dimensional subspace spanned by the leading $r+1$ right singular vectors.

We will use this type of argument again, so to be more precise: the matrix $\left[W, v_{1}, \ldots, v_{r+1}\right.$ ] is "fat" rectangular, so must have a null vector. That is, $\left[W, v_{1}, \ldots, v_{r+1}\right]\left[\begin{array}{l}x_{1} \\ x_{2}\end{array}\right]=0$ has a nonzero solution $x_{1}, x_{2}$; then $W x_{1}$ is such a vector ${ }^{9}$. We scale it so that it has unit norm $\left\|W x_{1}\right\|_{2}=\left\|\left[v_{1}, \ldots, v_{r+1}\right] x_{2}\right\|_{2}=1$, that is, $\left\|x_{1}\right\|_{2}=\left\|x_{2}\right\|_{2}=1$.
4. Note that $U \Sigma V^{T} W x_{1}=U \Sigma V^{T}\left[v_{1}, \ldots, v_{r+1}\right] x_{2}$ and $V^{T}\left[v_{1}, \ldots, v_{r+1}\right]=\left[\begin{array}{c}I_{r+1} \\ 0\end{array}\right]$, so $\left\|U \Sigma V^{T} W x_{1}\right\|_{2}=\left\|U\left[\begin{array}{c}\Sigma_{r+1} \\ 0\end{array}\right] x_{2}\right\|_{2}$, where $\Sigma_{r+1}$ is the leading (top-left) $(r+1) \times(r+1)$ part of $\Sigma$. As $U$ is orthogonal this is equal to $\left\|\Sigma_{r+1} y_{1}\right\|_{2}$, and $\left\|\Sigma_{r+1} y_{1}\right\|_{2} \geq \sigma_{r+1}$ can be verified by direct calculation.

Finally, for the reverse direction, take $B=A_{r}$.

[^8]
### 3.1 Low-rank approximation: image compression

A classical and visually pleasing illustration of low-rank approximation by the truncated SVD is image compression. (Admittedly this is slightly outdated-it is not the state-of-the-art way of compressing images.)

The idea is that greyscale images can be represented by a matrix where each entry indicates the intensity of a pixel. The matrix can then be compressed by finding a low-rank approximation ${ }^{10}$, resulting in image compression. (Of course there are generalisations for color images.)

Below in Figure 1 we take the Oxford logo, represent it as a matrix and find its rank$r$ approximation, for varying $r$. (The matrix being a mere $500 \times 500$, its SVD is easy to compute in a fraction of a second; see Section 10.2 for how this is done.) We then reconstruct the image from the low-rank approximations to visualise them.


Figure 1: Image compression by low-rank approximation via the truncated SVD.
We see that as the rank is increased the image becomes finer and finer. At rank 50 it is fair to say the image looks almost identical to the original. The original matrix is $500 \times 500$, so we still achieve a significant amount of data compression in the matrix with $r=50$.

[^9]
## 4 Courant-Fischer minmax theorem

Continuing on SVD-related topics, we now discuss a very important and useful result with far-reaching ramifications: the Courant-Fischer (C-F) minimax characterisation.

Theorem 4.1 The ith largest ${ }^{11}$ eigenvalue $\lambda_{i}$ of a symmetric matrix $A \in \mathbb{R}^{n \times n}$ is (below $x \neq 0$ )

$$
\begin{equation*}
\lambda_{i}(A)=\max _{\operatorname{dim} \mathcal{S}=i} \min _{x \in \mathcal{S}} \frac{x^{T} A x}{x^{T} x}\left(=\min _{\operatorname{dim} \mathcal{S}=n-i+1} \max _{x \in \mathcal{S}} \frac{x^{T} A x}{x^{T} x}\right) \tag{5}
\end{equation*}
$$

Analogously, for any rectangular $A \in \mathbb{C}^{m \times n}(m \geq n)$, we have

$$
\begin{equation*}
\sigma_{i}(A)=\max _{\operatorname{dim} \mathcal{S}=i} \min _{x \in \mathcal{S}} \frac{\|A x\|_{2}}{\|x\|_{2}}\left(=\min _{\operatorname{dim} \mathcal{S}=n-i+1} \max _{x \in \mathcal{S}} \frac{\|A x\|_{2}}{\|x\|_{2}}\right) . \tag{6}
\end{equation*}
$$

It would take some time to get a hang of what the statements mean. One helpful way to look at it is perhaps to note that inside the maximum in (6) the expression is $\min _{x \in \mathcal{S},\|x\|_{2}=1}\|A x\|_{2}=$ $\min _{Q^{T} Q=I_{i},\|y\|_{2}=1}\|A Q y\|_{2}=\sigma_{\min }(A Q)=\sigma_{i}(A Q)$, where $\operatorname{span}(Q)=\mathcal{S}$. The C-F theorem says $\sigma_{i}(A)$ is equal to the maximum possible value of this over all subspaces $\mathcal{S}$ of dimension $i$.

Proof: We will prove (6). A proof for (5) is analogous and a recommended exercise.

1. Fix $S$ and let $V_{i}=\left[v_{i}, \ldots, v_{n}\right]$. We have $\operatorname{dim}(\mathcal{S})+\operatorname{dim}\left(\operatorname{span}\left(V_{i}\right)\right)=i+(n-i+1)=n+1$, so $\exists$ intersection $w \in S \cap V_{i},\|w\|_{2}=1$.
2. For this $w$, we have $\|A w\|_{2}=\left\|\operatorname{diag}\left(\sigma_{i}, \ldots, \sigma_{n}\right)\left(V_{i}^{T} w\right)\right\|_{2} \leq \sigma_{i}$; thus $\sigma_{i} \geq \min _{x \in \mathcal{S}} \frac{\|A x\|_{2}}{\|x\|_{2}}$.
3. For the reverse inequality, take $S=\left[v_{1}, \ldots, v_{i}\right]$, for which $w=v_{i}$.

### 4.1 Weyl's inequality

As an example of the many significant ramifications of the C-F theorem, we present Weyl's theorem ${ }^{12}$ (or Weyl's inequality), an important perturbation result for singular values and eigenvalues of symmetric matrices.

Theorem 4.2 Weyl's inequality

- For the singular values of any matrix $A$,
$-\sigma_{i}(A+E) \in \sigma_{i}(A)+\left[-\|E\|_{2},\|E\|_{2}\right]$ for all $i$.
- Special case: $\|A\|_{2}-\|E\|_{2} \leq\|A+E\|_{2} \leq\|A\|_{2}+\|E\|_{2}$

[^10]- For eigenvalues of a symmetric matrix $A, \lambda_{i}(A+E) \in \lambda_{i}(A)+\left[-\|E\|_{2},\|E\|_{2}\right]$ for all $i$.
(Proof: exercise; almost a direct consequence of C-F.)
The upshot is that singular values and eigenvalues of symmetric matrices are insensitive to perturbation; a property known as being well conditioned.

This is important because this means a backward stable algorithm (see Section 7) computes these quantities with essentially full precision.

### 4.1.1 Eigenvalues of nonsymmetric matrices are sensitive to perturbation

It is worth remarking that eigenvalues of nonsymmetric matrices can be far from well conditioned! Consider for example the Jordan block $\left[\begin{array}{ll}1 & 1 \\ 0 & 1\end{array}\right]$. By perturbing this to $\left[\begin{array}{ll}1 & 1 \\ \epsilon & 1\end{array}\right]$ one gets eigenvalues that are perturbed by $\sqrt{\epsilon}$, a magnification factor of $1 / \sqrt{\epsilon} \gg 1$. More generally, consider the eigenvalues of a Jordan block and its perturbation

$$
J=\left[\begin{array}{cccc}
1 & 1 & & \\
& 1 & \ddots & \\
& & \ddots & 1 \\
& & & 1
\end{array}\right] \in \mathbb{R}^{n \times n}, \quad J+E=\left[\begin{array}{cccc}
1 & 1 & & \\
& 1 & \ddots & \\
& & \ddots & 1 \\
\epsilon & & & 1
\end{array}\right]
$$

$\lambda(J)=1$ ( $n$ copies), but we have $|\lambda(J+E)-1| \approx \epsilon^{1 / n}$. For example when $n=100$, an $10^{-100}$ perturbation in $J$ would result in a 0.1 perturbation in all the eigenvalues!
(nonexaminable) This is pretty much the worst-case situation. In the generic case where the matrix is diagonalizable, $A=X \Lambda X^{-1}$, with an $\epsilon$-perturbation the eigenvalues get perturbed by $O(c \epsilon)$, where the constant $c$ depends on the so-called condition number $\kappa_{2}(X)$ of the eigenvector matrix $X$ (see Section 7.2, and [8]).

### 4.2 More applications of C-F

(Somewhat optional) Let's explore more applications of C-F. A lot more can be proved; see [19] for many more results and examples along these lines.

Example 4.1 - $\sigma_{i}\left(\left[\begin{array}{l}A_{1} \\ A_{2}\end{array}\right]\right) \geq \max \left(\sigma_{i}\left(A_{1}\right), \sigma_{i}\left(A_{2}\right)\right)$
$\operatorname{Proof}$ (sketch): LHS $=\max _{\operatorname{dim} \mathcal{S}=i} \min _{x \in \mathcal{S},\|x\|_{2}=1}\left\|\left[\begin{array}{l}A_{1} \\ A_{2}\end{array}\right] x\right\|_{2}$, and for any $x,\left\|\left[\begin{array}{l}A_{1} \\ A_{2}\end{array}\right] x\right\|_{2} \geq$ $\max \left(\left\|A_{1} x\right\|_{2},\left\|A_{2} x\right\|_{2}\right)$.

- $\sigma_{i}\left(\left[\begin{array}{ll}A_{1} & A_{2}\end{array}\right]\right) \geq \max \left(\sigma_{i}\left(A_{1}\right), \sigma_{i}\left(A_{2}\right)\right)$

Proof: LHS $=\max _{\mathrm{dim} \mathcal{S}=i} \min _{\left[\begin{array}{ll}x_{1} \\ x_{2}\end{array}\right] \in \mathcal{S},\left\|\left[\begin{array}{l}x_{1} \\ x_{2}\end{array}\right]\right\|_{2}=1}\left\|\left[\begin{array}{ll}A_{1} & A_{2}\end{array}\right]\left[\begin{array}{l}x_{1} \\ x_{2}\end{array}\right]\right\|_{2}$, while
$\sigma_{i}\left(A_{1}\right)=\max _{\left.\operatorname{dim} \mathcal{S}=i, \text { range(S) }) \in \operatorname{range}\left(\begin{array}{c}I_{n} \\ 0\end{array}\right]\right)} \min _{\left[\begin{array}{c}x_{1} \\ x_{2}\end{array}\right] \in \mathcal{S},\left\|\left[\begin{array}{l}x_{1} \\ x_{2}\end{array}\right]\right\|_{2}=1}\left\|\left[\begin{array}{ll}A_{1} & \left.A_{2}\right]\end{array}\right]\left[\begin{array}{l}x_{1} \\ x_{2}\end{array}\right]\right\|_{2}$. Since
the latter imposes restrictions on $\mathcal{S}$ to take the maximum over, the former is at least as big.

## 4.3 (Taking stock) Matrix decompositions you should know

Let us now take stock to review the matrix decompositions that we have covered, along with those that we will discuss next.

- SVD $A=U \Sigma V^{T}$
- Eigenvalue decomposition $A=X \Lambda X^{-1}$
- Normal: $X$ unitary $X^{*} X=I$
- Symmetric: $X$ unitary and $\Lambda$ real
- Jordan decomposition: $A=X J X^{-1}, J=\operatorname{diag}\left(\left[\begin{array}{cccc}\lambda_{i} & 1 & & \\ & \lambda_{i} & \ddots & \\ & & \ddots & 1 \\ & & & \lambda_{i}\end{array}\right]\right)$
- Schur decomposition $A=Q T Q^{*}: T$ upper triangular
- QR: $Q$ orthonormal, $U$ upper triangular
- LU: $L$ lower triangular, $U$ upper triangular

Red: Orthogonal decompositions, stable computation available
The Jordan decomposition is mathematically the ultimate form that any matrix can be reduced to by a similarity transformation, but numerically it is not very useful-one of the problems is that Jordan decompositions are very difficult to compute, as an arbitrarily small perturbation can change the eigenvalues and block sizes by a large amount (recall the discussion in Section 4.1.1).

## 5 Linear systems $\boldsymbol{A x}=\boldsymbol{b}$

We now (finally) start our discussion on direct algorithms in NLA. The fundamental idea is to factorisa a matrix into a product of two (or more) simpler matrices. This foundational idea has been named one of the top 10 algorithms of the 20th centry [9].

The sophistication of the state-of-the-art implementation of direct methods is simply astonishing. For instance, a $100 \times 100$ dense linear system or eigenvalue problem can be solved in less than a milisecond on a standard laptop. Imagine solving it by hand!

We start with solving linear systems, unquestionably the most important problem in NLA (for applications). The classical way to solve it is via the (pivoted) LU decomposition. In some sense we needn't spend too much time here, as you must have seen much of the material (e.g. Gaussian elimination) before. However, the description of the LU decomposition given below will likely be different from the one that you have seen before. We have chosen a nonstandard description as it reveals its connection to low-rank approximation.

Let $A \in \mathbb{R}^{n \times n}$. Suppose we can decompose (or factorise) $A$ into (here and below, $*$ denotes entries that are possibly nonzero).

$$
A=\left[\begin{array}{lllll}
* & * & * & * & * \\
* & * & * & * & * \\
* & * & * & * & * \\
* & * & * & * & * \\
* & * & * & * & *
\end{array}\right]=\left[\begin{array}{lllll}
* & & & \\
* & * & & \\
* & * & * & \\
* & * & * & * \\
* & * & * & * & *
\end{array}\right]\left[\begin{array}{llll}
* & * & * & * \\
& * & * & * \\
& * & * & * \\
& & * & * \\
& & & *
\end{array}\right]=L U
$$

Here $L$ is lower triangular, and $U$ is upper triangular. How can we find $L, U$ ?
To get started, consider rewriting $A$ as

$$
\begin{aligned}
A & =\left[\begin{array}{l}
* \\
* \\
* \\
* \\
*
\end{array}\right]\left[\begin{array}{llll}
* & * & * & * \\
*
\end{array}\right]+\left[\begin{array}{llll}
* & * & * & * \\
* & * & * \\
* & * & * \\
* & * & * & * \\
*
\end{array}\right] \\
& =\underbrace{\left[\begin{array}{llll}
* \\
* \\
* \\
* \\
*
\end{array}\right]}_{L_{1} U_{1}}\left[\begin{array}{lllll}
* & * & * & * & *
\end{array}\right] \\
L_{2} & {\left[\begin{array}{lllll}
* \\
* \\
* \\
*
\end{array}\right]\left[\begin{array}{lllll}
0 & * & * & * & *
\end{array}\right]+\left[\begin{array}{ccc}
* & * & * \\
* & * & * \\
* & * & *
\end{array}\right]=\cdots }
\end{aligned}
$$

Namely, the first step finds $L_{1}, U_{1}$ such that:

$$
A=\underbrace{\left[\begin{array}{l}
* \\
* \\
* \\
*
\end{array}\right]\left[\begin{array}{lllll}
* & * & * & * & *
\end{array}\right]}_{L_{1} U_{1}}+\left[\begin{array}{lllll}
* & * & * & * \\
* & * & * & * & * \\
* & * & * & * \\
* & * & * & *
\end{array}\right] .
$$

Specifying the elements, the algorithm can be described as (taking $a=A_{11}$ )

$$
\begin{aligned}
& {\left[\begin{array}{lllll}
A_{11} & A_{12} & A_{13} & A_{14} & A_{15} \\
A_{21} & & & & \\
A_{31} & & & & \\
A_{41} & & & & \\
A_{51} & &
\end{array}\right]=\left[\begin{array}{lllll}
L_{11} \\
L_{21} \\
L_{31} \\
L_{41} \\
L_{51}
\end{array}\right]\left[\begin{array}{lllll}
U_{11} & U_{12} & U_{13} & U_{14} & U_{15} \\
& & & & \\
& & & &
\end{array}\right]+\left[\begin{array}{llll}
* & * & * & * \\
* & * & * & * \\
* & * & * & * \\
* & * & * & *
\end{array}\right]} \\
& =\underbrace{\left[\begin{array}{c}
1 \\
A_{21} / a \\
A_{31} / a \\
A_{41} / a \\
A_{51} / a
\end{array}\right.}_{=L_{1} U_{1}}]\left[\begin{array}{lllll}
A_{11} & A_{12} & A_{13} & A_{14} & A_{15} \\
& & & & \\
& & & &
\end{array}\right]+\left[\begin{array}{llll}
* & & * & * \\
* & * & * & * \\
* & * & * & * \\
* & * & * & *
\end{array}\right]
\end{aligned}
$$

Here we've assumed $a \neq 0$; we'll discuss the case $a=0$ later. Repeating the process gives

$$
\begin{aligned}
& =\left[L_{1}, L_{2}, \ldots, L_{5}\right]\left[\begin{array}{c}
U_{1} \\
U_{2} \\
\vdots \\
U_{5}
\end{array}\right]=\left[\begin{array}{cccccc}
* & & & & \\
* & * & & & \\
* & * & * & & \\
* & * & * & * & \\
* & * & * & * & *
\end{array}\right]\left[\begin{array}{lllll}
* & * & * & * & * \\
& * & * & * & * \\
& * & * & * \\
& & & * & * \\
& & & & *
\end{array}\right],
\end{aligned}
$$

an LU decomposition as required.
Note the above expression for $A$; clearly the $L_{i} U_{i}$ factors are rank- 1 matrices; the LU decomposition can be thought of as writing $A$ as a sum of (structured) rank-1 matrices.

### 5.1 Solving $A x=b$ via LU

Having found an LU decomposition $A=L U$, one can efficiently solve an $n \times n$ linear system $A x=b:$ First solve $L y=b$, then $U x=y$. Then $b=L y=L U x=A x$.

- These are triangular linear systems, which are easy to solve and can be done in $O\left(n^{2}\right)$ flops.
- Triangular solve is always backward stable: e.g. $(L+\Delta L) \hat{y}=b$ (see Higham's book)

The computational cost is

- For LU: $\frac{2}{3} n^{3}$ flops (floating-point operations).
- Triangular solve is $O\left(n^{2}\right)$.

Note that once we have an LU factorisation we can solve another linear system with respect to the same matrix with only $O\left(n^{2}\right)$ additional operations.

### 5.2 Pivoting

Above we've assumed the diagonal element (pivot) $a \neq 0$-when $a=0$ we are in trouble! In fact not every matrix has an LU decomposition. For example, there is no LU for $\left[\begin{array}{ll}0 & 1 \\ 1 & 0\end{array}\right]$. We need a remedy. In practice, a remedy is needed whenever $a$ is small.

The idea is to permute the rows, so that the largest element of the (first active) column is brought to the pivot. This process is called pivoting (sometimes partial pivoting, to emphasize the difference from complete pivoting wherein both rows and columns are permuted ${ }^{13}$ ). This

[^11]results in $P A=L U$, where $P$ is a permutation matrix: orthogonal matrices with only 1 and 0s (every row/column has exactly one 1 ); applying $P$ would reorder the rows (with $P A$ ) or columns (AP)).

Thus solving $A x_{i}=b_{i}$ for $i=1, \ldots, k$ requires $\frac{2}{3} n^{3}+O\left(k n^{2}\right)$ operations instead of $O\left(k n^{3}\right)$.

$$
\left[\begin{array}{lllll}
A_{11} & A_{12} & A_{13} & A_{14} & A_{15} \\
A_{21} & & & & \\
A_{31} \\
A_{11} & & & & \\
A_{51} & & &
\end{array}\right]=\left[\begin{array}{c}
1 \\
A_{21} / a \\
A_{31} / a \\
A_{41}, a \\
A_{51} / a
\end{array}\right]\left[\begin{array}{lllll}
A_{11} & A_{12} & A_{13} & A_{14} & A_{15} \\
& & & & \\
& & & &
\end{array}\right]+\left[\begin{array}{cccc}
* & * & * & * \\
* & * & * & * \\
* & * & * & * \\
* & * & * & *
\end{array}\right]
$$

When $a=0$, remedy: pivot, permute rows such that the largest entry of first (active) column is at the top. $\Rightarrow P A=L U, P$ : permutation matrix

- for $A x=b$, solve $P A x=P b \Leftrightarrow L U x=P b$
- cost still $\frac{2}{3} n^{3}+O\left(n^{2}\right)$

In fact, one can show that any nonsingular matrix $A$ has a pivoted LU decomposition. (proof: exercise). This means that any linear system that is computationally feasible can be solved by pivoted LU decomposition.

- Even with pivoting, unstable examples exist (Section 7), but almost always stable in practice and used everywhere.
- Stability here means $\hat{L} \hat{U}=P A+\Delta A$ with small $\|\Delta A\|$; see Section 7.


### 5.3 Cholesky factorisation for $A \succ 0$

If $A \succ 0$ (symmetric positive definite ${ }^{14}(\mathrm{~S}) \mathrm{PD} \Leftrightarrow \lambda_{i}(A)>0$ for all $i$ ), two simplifications happen in LU:

- We can take $U_{i}=L_{i}^{T}=: R_{i}$ by symmetry
- No pivot needed as long as $A$ is PD

$$
A=\underbrace{\left[\begin{array}{l}
*  \tag{7}\\
* \\
* \\
* \\
*
\end{array}\right]\left[\begin{array}{llll}
* & * & * & * \\
*
\end{array}\right]}_{R_{1} R_{1}^{T}}+\underbrace{\left[\begin{array}{ccccc}
* & * & * & * \\
* & * & * & * \\
* & * & * & * \\
*
\end{array}\right]}_{\text {also PD }}
$$

Notes:

- $\frac{1}{3} n^{3}$ flops, half as many as LU

[^12]- $\operatorname{diag}(R)$ no longer 1's (clearly)
- $A$ can be written as $A=R^{T} R$ for some $R \in \mathbb{R}^{n \times n}$ iff $A \succeq 0\left(\lambda_{i}(A) \geq 0\right)$
- Indefinite case: when $A=A^{*}$ but $A$ not PSD (i.e., negative eigenvalues are present), $\exists A=L D L^{*}$ where $D$ diagonal (when $A \in \mathbb{R}^{n \times n}, D$ can have $2 \times 2$ diagonal blocks), $L$ has 1's on diagonal. This is often called the "LDLT factorisation".
- It's not easy at this point to see why the second (red) term above is also PD; one way to see this is via the uniqueness of the Cholesky factorisation; see next section.

Therefore, roughly speaking, symmetric linear systems can be solved with half the effort of a general non-symmetric system.

## 6 QR factorisation and least-squares problems

We've seen that the LU decomposition is a key step towards solving $A x=b$. For an overdetemined problem (least-squares problems, the subject of Section 6.5), we will need the $Q R$ factorisation. For any $A \in \mathbb{R}^{m \times n}$, there exists a factorisation

where $Q \in \mathbb{R}^{m \times n}$ is orthonormal $Q^{T} Q=I_{n}$, and $R \in \mathbb{R}^{n \times n}$ is upper triangular.

- Many algorithms available: Gram-Schmidt, Householder QR, CholeskyQR, ...
- various applications: least-squares, orthogonalisation, computing SVD, manifold retraction...
- With Householder, pivoting $A=Q R P$ not needed for numerical stability.
- but pivoting gives rank-revealing QR (nonexaminable).


### 6.1 QR via Gram-Schmidt

No doubt you have seen the Gram-Schmidt (G-S) process. What you might not know is that when applied to the columns of a matrix $A$, it gives you a QR factorisation $A=Q R$.

Gram-Schmidt: Given $A=\left[a_{1}, a_{2}, \ldots, a_{n}\right] \in \mathbb{R}^{m \times n}$ (assume full-rank $\operatorname{rank}(A)=n$ ), find orthonormal $\left[q_{1}, \ldots, q_{n}\right]$ s.t. $\operatorname{span}\left(q_{1}, \ldots, q_{n}\right)=\operatorname{span}\left(a_{1}, \ldots, a_{n}\right)$

More precisely, the algorithm performs the following: $q_{1}=\frac{a_{1}}{\left\|a_{1}\right\|}$, then $\tilde{q}_{2}=a_{2}-q_{1} q_{1}^{T} a_{2}$, $q_{2}=\frac{\tilde{q}_{2}}{\left\|\tilde{q}_{2}\right\|}$, (orthogonalise and normalise)
repeat for $j=3, \ldots, n: \tilde{q}_{j}=a_{j}-\sum_{i=1}^{j-1} q_{i} q_{i}^{T} a_{j}, q_{j}=\frac{\tilde{q}_{j}}{\left\|\tilde{q}_{j}\right\|}$.
This gives a QR factorisation! To see this, let $r_{i j}=q_{i}^{T} a_{j}(i \neq j)$ and $r_{j j}=\| a_{j}-$ $\sum_{i=1}^{j-1} r_{i j} q_{i} \|$,

$$
\begin{aligned}
& q_{1}=\frac{a_{1}}{r_{11}} \\
& q_{2}=\frac{a_{2}-r_{12} q_{1}}{r_{22}} \\
& q_{j}=\frac{a_{j}-\sum_{i=1}^{j-1} r_{i j} q_{i}}{r_{j j}}
\end{aligned}
$$

which can be written equivalently as

$$
\begin{aligned}
a_{1} & =r_{11} q_{1} \\
a_{2} & =r_{12} q_{1}+r_{22} q_{2} \\
a_{j} & =r_{1 j} q_{1}+r_{2 j} q_{2}+\cdots+r_{j j} q_{j}
\end{aligned}
$$



- But this isn't the recommended way to compute the QR factorisation, as it's numerically unstable; see Section 7.7.1 and [17, Ch. 19,20].


### 6.2 Towards a stable QR factorisation: Householder reflectors

There is a beautiful alternative algorithm for computing the QR factorisation: Householder $Q R$ factorisation. In order to describe it, let us first introduce Householder reflectors. These are the class of matrices $H$ that are symmetric, orthogonal and can be written as a rank one update of the identity

$$
H=I-2 v v^{T}, \quad\|v\|=1
$$

- $H$ is orthogonal and symmetric: $H^{T} H=$ $H^{2}=I$. Its eigenvalues are 1 ( $n-1$ copies) and -1 (1 copy).
- For any given $u, w \in \mathbb{R}^{n}$ s.t. $\|u\|=\|w\|$ and $u \neq v, H=I-2 v v^{T}$ with $v=\frac{w-u}{\|w-u\|}$ gives $H u=w(\Leftrightarrow u=H w$, thus 'reflector')


It follows that by choosing the vector $v$ appropriately, one can perform a variety of operations to a given vector $x$. The primary example is the particular Householder reflector

$$
H=I-2 v v^{T}, \quad v=\frac{x-\|x\| e}{\|x-\| x\|e\|}, \quad e=[1,0, \ldots, 0]^{T},
$$

which satisfies $H x=[\|x\|, 0, \ldots, 0]^{T}$. That is, an arbitrary vector $x$ is mapped by $H$ to a multiple of $e=[1,0, \ldots, 0]^{T}$.
(I hope the picture above is helpful-the reflector reflects vectors about the hyperplane described by $v^{T} x=0$.)

In summary, we have the useful result
Lemma 6.1 For any $x \in \mathbb{R}^{n}$ not in the form $[ \pm\|x\|, 0, \ldots, 0]^{T}$, there exists a Householder reflector $H=I-2 v v^{T}$ where $\|v\|=1$ such that $H x=[\|x\|, 0, \ldots, 0]^{T}$.

### 6.3 Householder QR factorisation

Now we describe how to use the Householder reflectors in order to compute a QR factorisation of a given matrix $A \in \mathbb{R}^{m \times n}$.

The first step to obtain QR is to find $H_{1}$ s.t. $H_{1} a_{1}=\left[\begin{array}{c}\left\|a_{1}\right\| \\ 0 \\ \vdots \\ 0\end{array}\right]$,
and repeat to get $H_{n} \cdots H_{2} H_{1} A=R$ upper triangular, then $A=\left(H_{1} \cdots H_{n-1} H_{n}\right) R=$ $Q R$

Here is a pictorial illustration: start with

$$
A=\left[\begin{array}{llll}
* & * & * & * \\
* & * & * & * \\
* & * & * & * \\
* & * & * & * \\
* & * & * & *
\end{array}\right] .
$$

We apply a sequence of Householder reflectors

$$
\begin{aligned}
H_{1} A=\left(I-2 v_{1} v_{1}^{T}\right) A=\left[\begin{array}{llll}
* & * & * & * \\
* & * & * \\
* & * & * \\
* & * & * \\
* & * & *
\end{array}\right], & H_{2} H_{1} A=\left(I-2 v_{2} v_{2}^{T}\right) H_{1} A=\left[\begin{array}{llll}
* & * & * & * \\
& * & * & * \\
& * & * \\
& * & * \\
& * & *
\end{array}\right], \\
H_{3} H_{2} H_{1} A=\left[\begin{array}{llll}
* & * & * \\
& * & * & * \\
& * & * \\
& & * \\
& & *
\end{array}\right], & H_{n} \cdots H_{3} H_{2} H_{1} A=\left[\begin{array}{cccc}
* & * & * & * \\
& * & * & * \\
& & * & * \\
& & & *
\end{array}\right] .
\end{aligned}
$$

Note the zero pattern $v_{k}=\underbrace{0,0, \ldots, 0}_{k-10^{\prime} \text { 's }}, *, *, \ldots, *]^{T}$. We have

$$
H_{n} \cdots H_{2} H_{1} A=\left[\begin{array}{llll}
* & * & * & * \\
& * & * & * \\
& & * & * \\
& & &
\end{array}\right]=\left[\begin{array}{c}
R \\
0
\end{array}\right]
$$

To obtain a QR factorisation of $A$ we simply invert the Householder reflectors, noting that they are both orthogonal and symmetric; therefore the inverse is itself. This yields
$\Leftrightarrow A=\left(H_{1}^{T} \cdots H_{n-1}^{T} H_{n}^{T}\right)\left[\begin{array}{c}R \\ 0\end{array}\right]=: Q_{F}\left[\begin{array}{c}R \\ 0\end{array}\right]$; which is a full QR , wherein $Q_{F}$ is square orthogonal.

Moreover, writing $Q_{F}=\left[Q Q_{\perp}\right]$ where $Q \in \mathbb{R}^{m \times n}$ is orthonormal, we also have $A=Q R$ ('thin' QR or just QR ); this is more economical especially when $A$ is tall-skinny. In the majority of cases in computational mathematics, this is the object that we wish to compute).

We note some properties of Householder QR.

- Cost $\frac{4}{3} n^{3}$ flops with Householder-QR (twice that of LU).
- Unconditionally backward stable: the computed version satisfies $\hat{Q} \hat{R}=A+\Delta A$, $\left\|\hat{Q}^{T} \hat{Q}-I\right\|_{2}=\epsilon($ Section 7$)$.
- The algorithm gives a constructive proof for the existence of a full $\mathrm{QR} A=Q R$. It also gives, for example, a proof of the existence of the orthogonal complement of the column space of an orthonormal matrix $U$.
- To solve $A x=b$, solve $R x=Q^{T} b$ via triangle solve.
$\rightarrow$ Excellent method, but twice slower than LU (so it is rarely used)
- The process is aptly called orthogonal triangularisation. (By contrast, Gram-Schmidt and CholeskyQR ${ }^{15}$ are triangular orthogonalisation).

[^13]
### 6.4 Givens rotations

Householder QR is an excellent method for computing the QR factorisation of a general matrix, and it is widely used in practice. However, each Householder reflector acts globallyit affects all the entries of the (active part of) the matrix. For structured matrices - such as sparse matrices-sometimes there is a better tool to reduce the matrix to triangular form (and other forms) by working more locally. Givens rotations give a convenient tool for this. They are matrices of the form

$$
G=\left[\begin{array}{cc}
c & s \\
-s & c
\end{array}\right], \quad c^{2}+s^{2}=1
$$

Designed to 'zero' one element at a time. For example to compute the QR factorisation for an upper Hessenberg matrix, one can perform

$$
\begin{aligned}
& A=\left[\begin{array}{lllll}
* & * & * & * & * \\
* & * & * & * & * \\
& * & * & * & * \\
& & * & * & * \\
& & & * & *
\end{array}\right], \quad G_{1} A=\left[\begin{array}{llll}
* & * & * & *
\end{array}\right] \\
& G_{3} G_{2} G_{1} A=\left[\begin{array}{ccccc}
* & * & * & * & * \\
& * & * & * & * \\
& & * & * & * \\
& & & * & * \\
& & *
\end{array}\right], G_{4} G_{3} G_{2} G_{1} A=\left[\begin{array}{lllll}
* & * & * & * & * \\
& * & * & * & * \\
& & * & * & * \\
& & & * & * \\
& & & & *
\end{array}\right]=: R .
\end{aligned}
$$

This means $A=G_{1}^{T} G_{2}^{T} G_{3}^{T} G_{4}^{T} R$ is the QR factorisation. (note that Givens rotations are orthogonal but not symmetric-so its inverse is $G^{T}$, not $G$ ).

- $G$ acts locally on two rows (when left-multiplied; two columns if right-multiplied)
- Non-neighboring rows/cols allowed. For example, a rotation acting on the $i, j$ th columns would have $c, s$ values in the $(i, i),(i, j),(j, i),(j, j)$ entries. Visually,

$$
G_{i, j}=\left[\begin{array}{llllllllll}
1 & & & & & & & & & \\
& \ddots & & & & & & & & \\
& & 1 & & & & & & & \\
& & & \cos (\theta) & & & & \sin (\theta) & & \\
& & & & 1 & & & & & \\
& & & & \ddots & & & & & \\
& & & -\sin (\theta) & & & 1 & & & \\
& & & & & & & & & \\
& & & & & & & & & \ddots
\end{array}\right]
$$

### 6.5 Least-squares problems via QR

So far we have discussed linear systems wherein the coefficient matrix is square. However, in many situations in data science and beyond, there is a good reason to over-sample to obtain a
robust solution (for instance, in the presence of noise in measurements). For example, when there is massive data and we would like to fit the data with a simple model, we will have many more equations than the degrees of freedom. This leads to the so-called least-squares problem which we will discuss here.

Given $A \in \mathbb{R}^{m \times n}, m \geq n$ and $b \in \mathbb{R}^{m}$, a least-squares problem seeks to find $x \in \mathbb{R}^{n}$ such that the residual is minimised:

$$
\begin{equation*}
\min _{x} \| \boxed{A} \boxed{\boxed{x}-\square \|_{2}} \tag{8}
\end{equation*}
$$

- 'Overdetermined' linear system; attaining equality $A x=b$ is usually impossible
- Thus the goal is to try minimise the residual $\|A x-b\|$; usually $\|A x-b\|_{2}$ but sometimes e.g. $\|A x-b\|_{1}$ is of interest. Here we focus on $\|A x-b\|_{2}$.
- Throughout we assume full rank condition $\operatorname{rank}(A)=n$; this makes the solution unique and is generically satisfied. If not, the problem will have infinitely many minimisers (and a standard practice is to look for the minimum-norm solution).

Here is how we solve the least-squares problem (8).
Theorem 6.1 Let $A \in \mathbb{R}^{m \times n}, m>n$ and $b \in \mathbb{R}^{m}$, with $\operatorname{rank}(A)=n$. The least-squares problem $\min _{x}\|A x-b\|_{2}$ has solution given by $x=R^{-1} Q^{T} b$, where $A=Q R$ is the (thin) $Q R$ factorisation.

Proof: Let $A=\left[\begin{array}{ll}Q & Q_{\perp}\end{array}\right]\left[\begin{array}{c}R \\ 0\end{array}\right]=Q_{F}\left[\begin{array}{c}R \\ 0\end{array}\right]$ be 'full' QR factorisation. Then

$$
\|A x-b\|_{2}=\left\|Q_{F}^{T}(A x-b)\right\|_{2}=\left\|\left[\begin{array}{c}
R \\
0
\end{array}\right] x-\left[\begin{array}{c}
Q^{T} b \\
Q_{\perp}^{T} b
\end{array}\right]\right\|_{2}
$$

so $x=R^{-1} Q^{T} b$ is solution.
This also gives an algorithm (which is essentially the workhorse algorithm used in practice):

1. Compute thin QR factorisation $A=Q R$ (using Householder QR )
2. Solve linear system $R x=Q^{T} b$.

- This process is backward stable. That is, the computed $\hat{x}$ solution for $\min _{x} \|(A+$ $\Delta A) x+(b+\Delta b) \|_{2}$ (see Higham's book Ch.20)
- Unlike square system $A x=b$, one really needs QR : LU won't do the job at all.

One might wonder why we chose the 2-norm in the least-squares formulation (8). Unlike for low-rank approximation (where the truncated SVD is a solution for any unitarily invariant norm ${ }^{16}$ ) the choice of the norm does matter and affects the properties of the solution $x$ significantly. For example, an increasingly popular choice of norm is the 1-norm, which tends to promote sparsity in the quantity to be minimised. In particular, if we simply replace the tune alarm with the 1 norm, the solution tends to give a residual that is sparse. (nonexaminable)

### 6.6 QR-based algorithm for linear systems

It is straightforward to see that the exact same algorithm can be applied for solving square linear systems. Is this algorithm good? Absolutely! It turns out that it is even better than the LU-based method in that backward stability can be guaranteed (which isn't the case with pivoted LU). However, it is unfortunately twice expensive; which is the reason LU is used in the vast majority of cases for solving linear systems.

Another very stable algorithm is to compute the SVD $A=U \Sigma V^{T}$ and take $x=$ $V \Sigma^{-1} U^{T} b$. This is even more expensive than via QR (by $\approx \mathrm{x} 10$ ).

### 6.7 Solution of least-squares via normal equation

There is another way to solve the least-squares problem, by the so-called normal equation. We've seen that

$$
\min _{x}\|A x-b\|_{2}, \quad A \in \mathbb{R}^{m \times n}, m \geq n
$$

$x=R^{-1} Q^{T} b$ is the solution $\Leftrightarrow x$ solution for $n \times n$ normal equation

$$
\left(A^{T} A\right) x=A^{T} b
$$

- $A^{T} A \succeq 0$ (always) and $A^{T} A \succ 0$ if $\operatorname{rank}(A)=n$; then PD linear system; use Cholesky to solve.
- This is fast! but NOT backward stable; $\kappa_{2}\left(A^{T} A\right)=\left(\kappa_{2}(A)\right)^{2}$ where $\kappa_{2}(A)=\frac{\sigma_{\max }(A)}{\sigma_{\min }(A)}$ condition number (next topic)

In fact, more generally, given a linear least-squares approximation problem $\min _{p \in \mathcal{P}}\|f-p\|$ in an inner-product space (of which (8) is a particular example; other examples include polynomial approximation of a function with the inner product e.g. $\left.\langle f, g\rangle=\int_{-1}^{1} f(x) g(x) d x\right)$ the solution is characterised by the property that the residual is orthogonal to the subspace from which a solution is sought, that is, $\langle q, f-p\rangle=0$ for all $q \in \mathcal{P}$. To see this, consider the problem of approximating $f$ with $p$ in a subspace $\mathcal{P}$. Let $p_{*} \in \mathcal{P}$ be such that the residual $f-p_{+}$is orthogonal to any element in $\mathcal{P}$. Then for any $q \in \mathcal{P}$, we have $\left\|f-\left(p_{*}+q\right)\right\|^{2}=$

[^14]$\left\|f-p_{*}\right\|^{2}-2\left\langle f-p_{*}, q\right\rangle+\|q\|^{2}=\left\|f-p_{*}\right\|^{2}+\|q\|^{2} \geq\left\|f-p_{*}\right\|^{2}$, proving $p_{*}$ is a minimiser (it is actually unique).

Since we mentioned Cholesky, let us now revisit (7) and show why the second term there must be PSD. A PD matrix has an eigenvalue decomposition $A=V D^{2} V^{T}=\left(V D V^{T}\right)^{2}=$ $\left(V D V^{T}\right)^{T}\left(V D V^{T}\right)$. Now let $V D V^{T}=Q R$ be the QR factorisation. Then $\left(V D V^{T}\right)^{T}\left(V D V^{T}\right)=$ $R^{T} R$ (this establishes the existence of Cholesky). But now the 0 -structure in (7) means the first term must be $r r^{T}$ where $r^{T}$ is the first row of $R$, and hence the second term must be $R_{2}^{T} R_{2}$, which is PSD. Here $R^{T}=\left[\begin{array}{rl}r & R_{2}^{T}\end{array}\right]$.

### 6.8 Application of least-squares: regression/function approximation

To illustrate the usefulness of least-squares problems as compared with linear systems here let's consider a function approximation problem.

Given function $f:[-1,1] \rightarrow \mathbb{R}$,
Consider approximating via polynomial $f(x) \approx p(x)=\sum_{i=0} c_{i} x^{i}$.
Very common technique: Regression: this is a very widely applicable problem in statistics and data science.

1. Sample $f$ at points $\left\{z_{i}\right\}_{i=1}^{m}$, and
2. Find coefficients $c$ defined by Vandermonde system $A c \approx f$,

$$
\left[\begin{array}{cccc}
1 & z_{1} & \cdots & z_{1}^{n} \\
1 & z_{2} & \cdots & z_{2}^{n} \\
\vdots & \vdots & & \vdots \\
1 & z_{m} & \cdots & z_{m}^{n}
\end{array}\right]\left[\begin{array}{c}
c_{0} \\
\vdots \\
c_{n}
\end{array}\right] \approx\left[\begin{array}{c}
f\left(z_{1}\right) \\
f\left(z_{2}\right) \\
\vdots \\
f\left(z_{m}\right)
\end{array}\right] .
$$

- Numerous applications, e.g. in statistics, numerical analysis, approximation theory, data analysis!

Illustration We illustrate this with an example where we approximate the function $f(x)=$ $1+\sin (10 x) \exp (x)$ (which we suppose we don't know but we can sample it).



We observe that with 11 (equispaced) sample points, the degree-10 polynomial is deviating from the 'true function' quite a bit. With many more sample points the situation significantly improves. This is not a cherry-picked example but a phenomenon that can be mathematically proved; look for "Lebesgue constant" if interested (nonexaminable).

## 7 Numerical stability

An important aspect that is very often overlooked in numerical computing is numerical stability. Very roughly, it concerns the quality of a solution obtained by a numerical algorithm, given that computation on computers is done not exactly but with rounding errors. So far we have mentioned stability here and there in passing, but in this section it will be our focus.

Let us first look at an example where roundoff errors play a visible role to affect the computed solution of a linear system.

The situation is complicated. For example, let $A=U \Sigma V^{T}$, where $U=\frac{1}{\sqrt{2}}\left[\begin{array}{cc}1 & 1 \\ 1 & -1\end{array}\right]$, $\Sigma=\left[\begin{array}{cc}1 & \\ & 10^{-15}\end{array}\right], V=I$, and let $b=A\left[\begin{array}{l}1 \\ 1\end{array}\right]$. That is, we are solving a linear system whose solution is $x=\left[\begin{array}{l}1 \\ 1\end{array}\right]$.

If we solve this in MATLAB using $\mathrm{x}=\mathrm{A} \backslash \mathrm{b}$, the output is $\left[\begin{array}{c}1.0000 \\ 0.94206\end{array}\right]$. Quite different from the exact solution! Did something go wrong? Did MATLAB or the algorithm fail? The answer is NO, MATLAB and the algorithm (LU) performed just fine. This is a ramification of ill-conditioning, not instability. Make sure that after covering this section, you will be able to explain what happened in the above example.

### 7.1 Floating-point arithmetic

The IEEE (double precision) floating point arithmetic is by far the most commonly used model of computation adopted in computation, and we will assume its use here. We will not get into its details as it becomes too computer scientific, rather than mathematical. Just to give a very sketchy introduction

- Computers store number in base 2 with finite/fixed memory (bits)
- Numbers not exactly representable with finite bits in base 2 , including irrational numbers, are stored inexactly (rounded), e.g. $1 / 3 \approx 0.333 \ldots$. The unit with which rounding takes place is the machine precision, often denoted by $\epsilon$ (or $u$ for unit roundoff). In the most standard setting of IEEE double-precision arithmetic, $u \approx 10^{-16}$.
- Whenever calculations (addition, subtraction, multiplication, division) are performed, the result is rounded to the nearest floating-point number (rounding error); this is where numerical errors really creep in
- Thus the accuracy of the final error is nontrivial; in pathological cases, it is rubbish!

To get an idea of how things can go wrong and how serious the final error can be, here are two examples with MATLAB:

- $\left((\operatorname{sqrt}(2))^{2}-2\right) * 1 \mathrm{e} 15=0.4441$ (should be $\left.0 ..\right)$
- $\sum_{n=1}^{\infty} \frac{1}{n} \approx 30$ (should be $\infty$. .)

For matrices, there are much more nontrivial and surprising phenomena than these. An important (but not main) part of numerical analysis/NLA is to study the effect of rounding errors. This topic can easily span a whole course. By far the best reference on this topic is Higham's book [17].

In this section we denote by $f l(X)$ a computed version of $X$ (fl stands for floating point). For basic operations such as addition and multiplication, one has $f l(x+y)=x+y+c \epsilon$ where $|c| \leq \max (|x|,|y|)$ and $f l(x y)=x y+c \epsilon$ where $|c| \leq \max (|x y|)$.

### 7.2 Conditioning and stability

It is important to solidify the definition of stability (which is a property of an algorithm) and conditioning (which concerns the sensitivity of a problem and has nothing to do with the algorithm used to solve it).

- Conditioning is the sensitivity of a problem (e.g. of finding $y=f(x)$ given $x$ ) to perturbation in inputs, i.e., how large $\kappa:=\sup _{\delta x}\|f(x+\delta x)-f(x)\| /\|\delta x\|$ is in the limit $\delta x \rightarrow 0$. (Very informally, one can think of conditioning as the largest directional derivative).
(this is the absolute condition number; equally important is the relative condition number $\left.\kappa_{r}:=\sup _{\delta x} \frac{\|f(x+\delta x)-f(x)\|}{\|f(x)\|} / \frac{\|\delta x\|}{\|x\|}\right)$
- (Backward) Stability is a property of an algorithm, which describes if the computed solution $\hat{y}$ is a 'good' solution, in that it is an exact solution of a nearby input, that is, $\hat{y}=f(x+\Delta x)$ for a small $\Delta x$ : if $\|\Delta x\|$ can be shown to be small, $\hat{y}$ is a backward stable solution. If an algorithm is guarantee to output a backward stable solution, that algorithm is called backward stable. Throughout this section, $\Delta$ denotes a quantity that is small relative to the quantity to follow: $\|\Delta X\| /\|X\|=O(\epsilon)$, where $\epsilon$ is the machine precision.

To repeat, conditioning is intrinsic in the problem. Stability is a property of an algorithm. Thus we will never say "this problem is backward stable" or "this algorithm is illconditioned". We can say "this problem is ill/well-conditioned", or "this algorithm is/isn't
(backward) stable". If a problem is ill-conditioned $\kappa \gg 1$, and the computed solution is no very accurate, then one should blame the problem, not the algorithm. In such cases, a backward stable solution (see below) is usually still considered a good solution.

Notation/convention in this section: $\hat{x}$ denotes a computed approximation to $x$ (e.g. of $\left.x=A^{-1} b\right)$. $\epsilon$ denotes a small term $O(u)$, on the order of unit roundoff/working precision; so we write e.g. $u, 10 u,(m+n) u, m n u$ all as $\epsilon$. (In other words, here we assume $m, n \ll u^{-1}$.)

- Consequently (in this lecture/discussion) the norm choice does not matter for the discussion.


### 7.3 Numerical stability; backward stability

Let us dwell more on (backward) stability, because it is really at the heart of the discussion on numerical stability. The word backward is key here and it probably differs from the natural notion of stability that you might first think of.

For a computational task $Y=f(X)$ (given input $X$, compute $Y$ ) and computed approximant $\hat{Y}$,

- Ideally, error $\|Y-\hat{Y}\| /\|Y\|=\epsilon$ : but this is seldom true, and often impossible! ( $u$ : unit roundoff, $\approx 10^{-16}$ in standard double precision)
- Good alg has Backward stability $\hat{Y}=f(X+\Delta X), \frac{\|\Delta X\|}{\|X\|}=\epsilon$ "exact solution of a slightly wrong input".
- Justification: The input (matrix) is usually inexact anyway, as storing it on a computer as a floating-point object already incurs rounding errors! Consequently, $f(X+\Delta X)$ is just as good at $f(X)$ at approximating $f\left(X_{*}\right)$ where $\|\Delta X\|=O\left(\left\|X-X_{*}\right\|\right)$.
We shall 'settle with' such solution, though it may not mean $\hat{Y}-Y$ is small.
- Forward stability ${ }^{17}\|Y-\hat{Y}\| /\|Y\|=O(\kappa(f) u)$ "error is as small as backward stable alg".
- Another important notion: mixed forward-backward stability: "The computed output is a slightly perturbed solution for a slightly perturbed problem".

[^15]
### 7.4 Matrix condition number

The best way to illustrate conditioning is to look at the conditioning of linear systems. In fact it leads to the following definition, which arises so frequently in NLA that it merits its own name: the condition number of a matrix.

Definition 7.1 The matrix condition number is defined by

$$
\kappa_{2}(A)=\frac{\sigma_{\max }(A)}{\sigma_{\min }(A)}(\geq 1)
$$

That is, $\kappa_{2}(A)=\frac{\sigma_{1}(A)}{\sigma_{n}(A)}$ for $A \in \mathbb{R}^{m \times n}, \quad m \geq n$.
Let's see how this arises:
Theorem 7.1 Consider a backward stable solution for $A x=b$, s.t. $(A+\Delta A) \hat{x}=b$ with $\|\Delta A\| \leq \epsilon\|A\|$ and $\kappa_{2}(A) \ll \epsilon^{-1}\left(\right.$ so $\left.\left\|A^{-1} \Delta A\right\| \ll 1\right)$. Then we have

$$
\frac{\|\hat{x}-x\|}{\|x\|} \leq \epsilon \kappa_{2}(A)+O\left(\epsilon^{2}\right) .
$$

Proof: By Neumann series

$$
(A+\Delta A)^{-1}=\left(A\left(I+A^{-1} \Delta A\right)\right)^{-1}=\left(I-A^{-1} \Delta A+O\left(\left\|A^{-1} \Delta A\right\|^{2}\right)\right) A^{-1}
$$

So $\hat{x}=(A+\Delta A)^{-1} b=A^{-1} b-A^{-1} \Delta A A^{-1} b+O\left(\left\|A^{-1} \Delta A\right\|^{2}\right)=x-A^{-1} \Delta A x+O\left(\left\|A^{-1} \Delta A\right\|^{2}\right)$, Hence

$$
\|x-\hat{x}\| \lesssim\left\|A^{-1} \Delta A x\right\| \leq\left\|A^{-1}\right\|\|\Delta A\|\|x\| \leq \epsilon\|A\|\left\|A^{-1}\right\|\|x\|=\epsilon \kappa_{2}(A)\|x\|
$$

In other words, even with a backward stable solution, one would only have $O\left(\kappa_{2}(A) \epsilon\right)$ relative accuracy in the solution. If $\kappa_{2}(A) \epsilon>1$, the solution may be rubbish! But the NLA view is that's not the fault of the algorithm, the blame is on the problem being so ill-conditioned.

### 7.4.1 Backward stable+well conditioned=accurate solution

We've seen that backward stability does not necessarily imply the solution is accurate. There is a happier side of this argument and useful rule of thumb that can be used to estimate the accuracy of a computed solution using backward stability and conditioning.

Suppose

- $Y=f(X)$ is computed backward stably i.e., $\hat{Y}=f(X+\Delta X),\|\Delta X\|=\epsilon$.
- Conditioning $\|f(X)-f(X+\Delta X)\| \lesssim \kappa\|\Delta X\|$.

Then (this is the absolute version, relative version possible)

$$
\|\hat{Y}-Y\| \lesssim \kappa \epsilon .
$$

'proof':

$$
\|\hat{Y}-Y\|=\|f(X+\Delta X)-f(X)\| \lesssim \kappa\|\Delta X\|\|f(X)\|=\kappa \epsilon
$$

Here is how to interpret the result: If the problem is well-conditioned $\kappa=O(1)$, this immediately implies good accuracy of the solution! But otherwise the solution might have poor accuracy - but it is still the exact solution of a nearby problem. This is often as good as one can possibly hope for.

The reason this is only a rule of thumb and not exactly rigorous is that conditioning only examines the asymptotic behavior, where the perturbation is infinitesimally small. Nonetheless it often gives an indicative estimate for the error and sometimes we can get rigorous bounds if we know more about the problem. Important examples include the following:

- Well-conditioned linear system $A x=b, \kappa_{2}(A) \approx 1$.
- Eigenvalues of symmetric matrices (via Weyl's bound $\lambda_{i}(A+E) \in \lambda_{i}(A)+\left[-\|E\|_{2},\|E\|_{2}\right]$ ).
- Singular values of any matrix $\sigma_{i}(A+E) \in \sigma_{i}(A)+\left[-\|E\|_{2},\|E\|_{2}\right]$.

Indeed, these problems are well-conditioned, so can be solved with extremely high accuracy, essentially to working precision $O(u)$ (times a small factor, usually bounded by something like $\sqrt{n}$ ).

Note: eigvecs/singvecs can be highly ill-conditioned even for the simplest problems. Again, think of the identity. (Those curious are invited to look up the Davis-Kahan $\sin \theta$ theorem [7].)

### 7.5 Stability of triangular systems

### 7.5.1 Backward stability of triangular systems

While we will not be able to discuss in detail which NLA algorithm is backward stable etc, we will pick a few important examples.

One fact is worth knowing (a proof is omitted; see [33] or [17]): triangular linear systems can be solved in a backward stable manner. This fact is important as these arise naturally in the solution of linear systems: $A x=b$ is solved via $L y=b, U x=y$ (triangular systems), as we've seen in Section 5 .

Let $R$ denote a matrix that is (upper or lower) triangular. The computed solution $\hat{x}$ for a (upper/lower) triangular linear system $R x=b$ solved via back/forward substitution is backward stable, i.e., it satisfies

$$
(R+\Delta R) \hat{x}=b, \quad\|\Delta R\|=O(\epsilon\|R\|) .
$$

Proof: Trefethen-Bau or Higham (nonexaminable but interesting). Notes:

- The backward error can be bounded componentwise.
- Using the previous rule-of-thumb, this means $\|\hat{x}-x\| /\|x\| \lesssim \epsilon \kappa_{2}(R)$.
- (unavoidably) poor worst-case (and attainable) bound when ill-conditioned
- often better with triangular systems; the behavior of triangular linear systems keep surprising experts!


### 7.5.2 (In)stability of $\boldsymbol{A x}=\boldsymbol{b}$ via LU with pivots

We have discussed how to solve a linear system using the LU decomposition. An obvious question given the context is: is it backward stable? This question has a fascinating answer, and transpires to touch on one of the biggest open problems in the field.

Fact (proof nonexaminable): Computed $\hat{L} \hat{U}$ satisfies $\frac{\|\hat{L} \hat{U}-A\|}{\|L\|\|U\|}=\epsilon$.
(note: not $\frac{\|\hat{L} \hat{U}-A\|}{\|A\|}=\epsilon$ )

- By stability of triangular systems $(L+\Delta L)(U+\Delta U) \hat{x}=b$. Now if $\|L\|\|U\|=O(\|A\|)$, then it follows that $\Rightarrow \hat{x}$ backward stable solution (exercise).

Question: Does $L U=A+\Delta A$ or $L U=P A+\Delta A$ with $\|\Delta A\|=\epsilon\|A\|$ (i.e., $\|L\|\|U\|=$ $O(\|A\|))$ hold?

Without pivot $(P=I):\|L\|\|U\| \gg\|A\|$ unboundedly (e.g. [ $\left.\left[\begin{array}{ll}\epsilon & 1 \\ 1 & 1\end{array}\right]\right)$ unstable.
With pivots:

- Worst-case: $\|L\|\|U\| \gg\|A\|$ grows exponentially with $n$, unstable.
- growth governed by that of $\|L\|\|U\| /\|A\| \Rightarrow\|U\| /\|A\|$.
- In practice (average case): perfectly stable.
- Hence this is how $A x=b$ is solved, despite alternatives with guaranteed stability exist (but slower; e.g. via SVD, or QR (next)).

Resolution/explanation: among biggest open problems in numerical linear algebra!

### 7.5.3 Backward stability of Cholesky for $A \succ 0$

We've seen that symmetric positive definite matrices can be solved with half the effort using the Cholesky factorisation. As a bonus, in this case it turns out that stability can be guaranteed.

The key fact is that the Cholesky factorisation $A=R^{T} R$ for $A \succ 0$

- succeeds without pivot (the active matrix is always positive definite).
- $R$ never contains entries $>\sqrt{\|A\|}$.

It follows that computing the Cholesky factorisation (assuming $A \succ 0$ ) is backward stable! Hence positive definite linear system $A x=b$ can always be solved in a backward stable manner via Cholesky.

### 7.6 Matrix multiplication is not backward stable

Here is perhaps a shock-matrix matrix multiplication, one of the most basic operations, is in general not backward stable.

Let's start with the basics.

- Vector-vector multiplication is backward stable: $f l\left(y^{T} x\right)=(y+\Delta y)(x+\Delta x)$; in fact $f l\left(y^{T} x\right)=(y+\Delta y) x$. (direct proof is possible)
- It immediately follows that matrix-vector multiplication is also backward stable: $f l(A x)=$ $(A+\Delta A) x$.
- But it is not true to say matrix-matrix multiplication is backward stable; which would require $f l(A B)$ to be equal to $(A+\Delta A)(B+\Delta B)$. This may not be satisfied!

In general, what we can prove is a bound for the forward error $\|f l(A B)-A B\| \leq$ $\epsilon\|A\|\|B\|$, so $\|f l(A B)-A B\| /\|A B\| \leq \epsilon \min \left(\kappa_{2}(A), \kappa_{2}(B)\right)$ (proof: problem sheet).

This is great news when $A$ or $B$ is orthogonal (or more generally square and wellconditioned): say if $A=Q$ is orthogonal, then we have

$$
\|f l(Q B)-Q B\| \leq \epsilon\|B\|
$$

so it follows that $f l(Q B)=Q B+\epsilon\|B\|$, hence defining $\Delta B=Q^{T} \epsilon\|B\|$ we have $f l(Q B)=$ $Q(B+\Delta B)$, that is, orthogonal multiplication is backward stable (this argument proves this for left multiplication; orthogonal right-multiplication is entirely analogous).

One of the reasons backward stability of matrix-matrix multiplication fails to hold is that there are not enough parameters in the backward error to account for the computational error incurred in the matrix multiplication. Each matrix-vector product is backward stable; but we cannot concentrate the backward errors into a single term without potentially increasing the backward error's norm.

On the other hand, we have seen that a linear system can be solved in a backward stable fashion (by QR if necessary; in practice LU with pivoting suffices). This means the inverse can be applied to a vector in a backward stable fashion: The computed $\hat{x}$ satisfies $(A+\Delta A)^{-1} \hat{x}=b$.

One might wonder, can we not solve $n$ linear systems in order to get a backward stable inverse? However, if one solves $n$ linear systems with $b=e_{1}, e_{2}, \ldots, e_{n}$, the solutions will satisfy $\left(A+\Delta_{i} A\right)^{-1} \hat{x}_{i}=e_{i}$, where $\Delta_{i} A$ depends on $i$. There is no uniform $\Delta A$ that is $O(\epsilon)$ such that $(A+\Delta A)^{-1}\left[\hat{x}_{1}, \ldots, \hat{x}_{n}\right]=I$.

Orthogonality matters for stability A happy and important exception is with orthogonal matrices $Q$ (or more generally with well-conditioned matrices):

$$
\frac{\|f l(Q A)-Q A\|}{\|Q A\|} \leq \epsilon, \quad \frac{\|f l(A Q)-A Q\|}{\|A Q\|} \leq \epsilon
$$

They are also backward stable:

$$
\begin{aligned}
& f l(Q A)=Q A+\epsilon \quad \Leftrightarrow \quad f l(Q A)=Q(A+\Delta A) \\
& f l(A Q)=A Q+\epsilon \quad \Leftrightarrow \quad f l(A Q)=(A+\Delta A) Q
\end{aligned}
$$

Hence algorithms involving ill-conditioned matrices are unstable (e.g. eigenvalue decomposition of non-normal matrices, Jordan form, etc), whereas those based on orthogonal matrices are stable. These include

- Householder QR factorisation, QR-based linear system (next subsection).
- QR algorithm for $A x=\lambda x$.
- Golub-Kahan algorithm for $A=U \Sigma V^{T}$.
- QZ algorithm for $A x=\lambda B x$.

Section 8 onwards treats our second big topic, eigenvalue problems. This includes discussing algorithms shown above in boldface.

### 7.7 Stability of Householder QR

Householder QR has excellent numerical stability, basically because it's based on orthogonal transformations. With Householder QR, the computed $\hat{Q}, \hat{R}$ satisfy

$$
\left\|\hat{Q}^{T} \hat{Q}-I\right\|=O(\epsilon), \quad\|A-\hat{Q} \hat{R}\|=O(\epsilon\|A\|)
$$

and (of course) $R$ is upper triangular.
Rough proof: Essentially the key idea is that multiplying by an orthogonal matrix is very stable and Householder QR is based on a sequence of multiplications by orthogonal matrices. To give a bit more detail,

- Each reflector satisfies $f l\left(H_{i} A\right)=H_{i} A+\epsilon_{i}\|A\|$.
- Hence $(\hat{R}=) f l\left(H_{n} \cdots H_{1} A\right)=H_{n} \cdots H_{1} A+\epsilon\|A\|$.
- $f l\left(H_{n} \cdots H_{1}\right)=: \hat{Q}^{T}=H_{n} \cdots H_{1}+\epsilon$.
- Thus $\hat{Q} \hat{R}=A+\epsilon\|A\|$.

Notes:

- This doesn't mean $\|\hat{Q}-Q\|,\|\hat{R}-R\|$ are small at all! Indeed $Q, R$ are as ill-conditioned as $A$ [17, Ch. 20].
- Solving $A x=b$ and least-squares problems via the QR factorisation is stable.


### 7.7.1 (In)stability of Gram-Schmidt

(Nonexaminable) A somewhat surprising fact is that the Gram-Schmidt algorithm, when used for computing the QR factorisation, is not backward stable. Namely, orthogonality of the computed $\hat{Q}$ matrix is not guaranteed.

- Gram-Schmidt is subtle:
- plain (classical) version: $\left\|\hat{Q}^{T} \hat{Q}-I\right\| \leq \epsilon\left(\kappa_{2}(A)\right)^{2}$.
- modified Gram-Schmidt (orthogonalise 'one vector at a time'): $\left\|\hat{Q}^{T} \hat{Q}-I\right\| \leq$ $\epsilon \kappa_{2}(A)$.
- Gram-Schmidt twice (G-S again on computed $\hat{Q}$ ) is excellent: $\left\|\hat{Q}^{T} \hat{Q}-I\right\| \leq \epsilon$.


## 8 Eigenvalue problems

First of all, recall that $A x=\lambda x$ has no explicit solution (neither $\lambda$ nor $x$ ); huge difference from $A x=b$ for which $x=A^{-1} b$.

From a mathematical viewpoint this marks an enormous point of departure: something that is is explicitly written vs. something that has no closed-from solution.

From a practical viewpoint the gulf is much smaller, because we have an extremely reliable algorithm for eigenvalue problems; so robust that it is essentially bulletproof provably backward stable.

Before we start describing the QR algorithm let's discuss a few interesting properties of eigenvalues.

- Eigenvalues are the roots of characteristic polynomial (Abel's theorem).
- For any polynomial $p$, there exist (infinitely many) matrices whose eigvals are roots of p.
- Here is a nice and useful fact ${ }^{18}$ : Let $p(x)=x^{n}+a_{n-1} x^{n-1}+\cdots+a_{1} x+a_{0}, a_{i} \in \mathbb{C}$. Then
$p(\lambda)=0 \Leftrightarrow \lambda$ eigenvalue of

$$
C=\left[\begin{array}{ccccc}
-a_{n-1} & -a_{n-2} & \cdots & -a_{1} & -a_{0} \\
1 & 1 & & & \\
& & \ddots & & \\
& & & 1 & 0
\end{array}\right] \in \mathbb{C}^{n \times n} .
$$

[^16]- So no finite-step algorithm exists for $A x=\lambda x$.

Eigenvalue algorithms are necessarily iterative and approximate.

- Same for SVD, as $\sigma_{i}(A)=\sqrt{\lambda_{i}\left(A^{T} A\right)}$.
- But this doesn't mean they're inaccurate!

Usual goal: compute the Schur decomposition $A=U T U^{*}: U$ unitary, $T$ upper triangular.

### 8.1 Schur decomposition

Theorem 8.1 Let $A \in \mathbb{C}^{n \times n}$ (arbitrary square matrix). Then there exists a unitary matrix $U \in \mathbb{C}^{n \times n}$ s.t.

$$
\begin{equation*}
A=U T U^{*}, \tag{9}
\end{equation*}
$$

with $T$ upper triangular. The decomposition (9) is called the Schur decomposition.
Note that

- $\operatorname{eig}(A)=\operatorname{eig}(T)=\operatorname{diag}(T)$.
- $T$ diagonal iff $A$ is normal, i.e., $A^{*} A=A A^{*}$

Proof: Let $A v=\lambda_{1} v$ and find $U_{1}=\left[v_{1}, V_{\perp}\right]$ unitary. Then $A U_{1}=U_{1}\left[\begin{array}{lllll}* & * & * & * & * \\ * & * & * & * \\ * & * & * & * \\ * & * & * & * \\ * & * & * & *\end{array}\right] \Leftrightarrow$
$U_{1}^{*} A U_{1}=\left[\begin{array}{cccc}* & * & * & * \\ * & * & * & * \\ * & * & * & * \\ * & * & * & * \\ * & * & * & *\end{array}\right]$. Repeat on the lower-right $(n-1) \times(n-1)$ part to get $U_{n-1}^{*} U_{n-2}^{*} \ldots U_{1}^{*} A U_{1} U_{2} \ldots U_{n-1}=T$.

The reason we usually take the Schur form to be the goal is that its computation can be done in a backward stable manner. Essentially it boils down to the fact that the decomposition is based on orthogonal transformations.

- For normal matrices $A^{*} A=A A^{*}, T$ must be diagonal and Schur form is automatically diagonalised.
- For nonnormal $A$, if diagonalisation $A=X \Lambda X^{-1}$ really necessary, done via starting with the Schur decomposition and further reducing $T$ by solving Sylvester equations; but this process involves nonorthogonal transformations and is not backward stable (nonexaminable).
- The Schur decomposition is among the few examples in NLA where the difference between $\mathbb{C}$ and $\mathbb{R}$ matters a bit. When working only in $\mathbb{R}$ one cannot always get a triangular $T$; it will have $2 \times 2$ blocks in the diagonal (these blocks have complex eigenvalues). This is essentially because real matrices can have complex eigenvalues. This is still not a major issue as eigenvalues of $2 \times 2$ matrices can be computed easily.
(This marks the end of "the first half" of the course (i.e., up until the basic facts about eigenvalues, but not its computation). This information is relevant to MMSC students.)


### 8.2 The power method for finding the dominant eigenpair $A x=\lambda x$

We now start describing the algorithms for solving eigenvalue problems. The first algorithm that we introduce is surprisingly simple and is based on the idea of keep multiplying the matrix $A$ to an arbitrary vector.

This algorithm by construction is designed to compute only a single eigenvalue and its corresponding eigenvector, namely the dominant one. It is not able, at least as presented, to compute all the eigenvalues.

Despite its limitation and simplicity it is an extremely important algorithm and the underlying idea is in fact a basis for the ultimate QR algorithm that we use in order to compute all eigenvalues of a matrix. So here is the algorithm, called the power method.

```
\(\overline{\text { Algorithm 8.1 The power method for computing the dominant eigenvalue and eigenvector }}\)
of \(A \in \mathbb{R}^{n \times n}\).
    : Let \(x \in \mathbb{R}^{n}:=\) random initial vector (unless specified)
    Repeat: \(x=A x, x=\frac{x}{\|x\|_{2}}\).
    \(\hat{\lambda}=x^{T} A x\) gives an estimate for the eigenvalue.
```

- Convergence analysis: suppose $A$ is diagonalisable (generic assumption). We can write $x_{0}=\sum_{i=1}^{n} c_{i} v_{i}, A v_{i}=\lambda_{i} v_{i}$ with $\left|\lambda_{1}\right|>\left|\lambda_{2}\right|>\cdots$. Then after $k$ iterations,

$$
x=C \sum_{i=1}^{n}\left(\frac{\lambda_{i}}{\lambda_{1}}\right)^{k} c_{i} v_{i} \rightarrow C c_{1} v_{1} \quad \text { as } k \rightarrow \infty \text { for some scalar } C
$$

- Converges geometrically $(\lambda, x) \rightarrow\left(\lambda_{1}, v_{1}\right)$ with linear rate $\frac{\left|\lambda_{2}\right|}{\left|\lambda_{1}\right|}$
- What does this imply about $A^{k}=Q R$ as $k \rightarrow \infty$ ? First column of $Q \rightarrow v_{1}$ under mild assumptions.

Notes:

- Google pagerank \& Markov chain are linked to the power method.
- As we'll see, the power method is basis for refined algorithms (QR algorithm, Krylov methods (Lanczos, Arnoldi,...))


### 8.2.1 Digression (optional): Why compute eigenvalues? Google PageRank

Let us briefly digress and talk about a famous application that at least used to solve an eigenvalue problem in order to achieve a familiar task: Google web search.

Once Google receives a user's inquiry with keywords, it needs to rank the relevant webpages, to output the most important pages first.

Here, the 'importance' of websites is determined by the dominant eigenvector of columnstochastic matrix (i.e., column sums are all 1)

$$
A=\alpha P+\frac{(1-\alpha)}{n}\left[\begin{array}{ccc}
1 & \cdots & 1 \\
\vdots & \ddots & \vdots \\
1 & \cdots & 1
\end{array}\right]
$$

$P$ : scaled adjacency matrix $\left(P_{i j}=1\right.$ if $i, j$

connected by an edge, 0 otherwise, and then
image from wikipedia scaled s.t. column sums are 1$), \alpha \in(0,1)$

To solve this approximately (note that high accuracy isn't crucial here- getting the ordering wrong isn't the end of the world), Google does (at least in the old days) a few steps of power method: with initial guess $x_{0}, k=0,1, \ldots$

1. $x_{k+1}=A x_{k}$
2. $x_{k+1}=x_{k+1} /\left\|x_{k+1}\right\|_{2}, \quad k \leftarrow k+1$, repeat.

- $x_{k} \rightarrow$ PageRank vector $v_{1}: A v_{1}=\lambda_{1} v_{1}$. As $A$ is a nonnegative matrix $A_{i j} \geq 0$, the dominant eigenvector $v_{1}$ can be taken to be nonnegative (by the Perron-Frobenius theorem; nonexaminable).

For more on pagerank etc, see Gleich [13] if interested (nonexaminable).

### 8.2.2 Shifted inverse power method

We saw that the convergence of the power method is governed by $\left|\lambda_{1} / \lambda_{2}\right|$, the ratio between the absolute value of dominant and the next dominant eigenvalue.

If this ratio is close to 1 the power method would converge very slowly. A natural question arises: can one speed up the convergence in such cases?

This leads to the inverse power method, or more generally shifted inverse power method, which can compute eigenpairs with much improved speed provided that the parameters (shifts) are chosen appropriately.

```
Algorithm 8.2 Shifted inverse power method for computing the eigenvalue and eigenvector
of \(A \in \mathbb{R}^{n \times n}\) closest to a prescribed value \(\mu \in \mathbb{C}\).
    1: Let \(x \in \mathbb{R}^{n}:=\) random initial vector, unless specified.
    2: Repeat: \(x:=(A-\mu I)^{-1} x, x=x /\|x\|\).
    \(\hat{\lambda}=x^{T} A x\) gives an estimate for the eigenvalue.
```

Note that the eigenvalues of the matrix $(A-\mu I)^{-1}$ are $(\lambda(A)-\mu)^{-1}$.

- By the same analysis as above, shifted-inverse power method converges with improved linear rate $\frac{\left|\lambda_{\sigma(1)}-\mu\right|}{\left|\lambda_{\sigma(2)}-\mu\right|}$ to the eigenpair closest to $\mu$. Here $\sigma$ denotes a permutation of $\{1, \ldots, n\}$ such that $\left|\lambda_{\sigma(1)}-\mu\right|$ minimises $\left|\lambda_{i}-\mu\right|$ over $i$.
- $\mu$ can change adaptively with the iterations. The choice $\mu:=x^{T} A x$ gives the Rayleigh quotient iteration, with quadratic convergence $\left\|A x^{(k+1)}-\lambda^{(k+1)} x^{(k+1)}\right\|=O\left(\| A x^{(k)}-\right.$ $\lambda^{(k)} x^{(k)} \|^{2}$ ); this is further improved to cubic convergence if $A$ is symmetric.

It is worth emphasising that the improved convergence comes at a cost: one step of shifted inverse power method requires a solution of a linear system, which is clearly more expensive than power method.

## 9 The QR algorithm

We'll now describe an algorithm called the QR algorithm that is used universally for solving eigenvalue problems of moderate size, e.g. by MATLAB's eig. Given $A \in \mathbb{R}^{n \times n}$, the algorithm

- Finds all eigenvalues (approximately but reliably) in $O\left(n^{3}\right)$ flops,
- Is backward stable.

Sister problem: Given $A \in \mathbb{R}^{m \times n}$ compute SVD $A=U \Sigma V^{T}$

- 'ok' algorithm: $\operatorname{eig}\left(A^{T} A\right)$ to find $V$, then normalise $A V$
- there's a better (but still closely related) algorithm: Golub-Kahan bidiagonalisation, discussed later in Section 10.2.


### 9.1 QR algorithm for $A x=\lambda x$

As the name suggests, the QR algorithm is based on the QR factorisation of a matrix. Another key fact is that the eigenvalues of a product of two matrices remain the same when the order of the product is reversed $(\operatorname{eig}(A B)=\operatorname{eig}(B A)$, problem sheet). The QR algorithm essentially is a simple combination of these two ideas, and the vanilla version is deceptively
simple: basically take the QR factorisation, swap the order, take the QR and repeat the process. Namely,

```
Algorithm 9.1 The QR algorithm for finding the Schur decomposition of a square matrix
A.
    : Set \(A_{1}=A\).
    2: Repeat: \(A_{1}=Q_{1} R_{1}, \quad A_{2}=R_{1} Q_{1}, \quad A_{2}=Q_{2} R_{2}, \quad A_{3}=R_{2} Q_{2}, \quad \ldots\)
```

Notes:

- $A_{k}$ are all similar: $A_{k+1}=Q_{k}^{T} A_{k} Q_{k}$
- We shall 'show' that $A \rightarrow$ triangular (diagonal if $A$ normal), under weak assumptions.
- Basically: $Q R$ (factorise) $\rightarrow R Q$ (swap) $\rightarrow Q R \rightarrow R Q \rightarrow \cdots$
- Fundamental work by Francis $(61,62)$ and Kublanovskaya (63)
- This is a truly magical algorithm!
- The algorithm is backward stable, as based on orthogonal transforms: essentially, $R Q=f l\left(Q^{T}(Q R) Q\right)=Q^{T}(Q R+\Delta(Q R)) Q$.
- always converges (with shifts) in practice, but a global proof is currently unavailable(!)
- uses 'shifted inverse power method' (rational functions) without inversions

Again, the QR algorithm performs: $A_{k}=Q_{k} R_{k}, A_{k+1}=R_{k} Q_{k}$, repeat.
Let's look at its properties.
Theorem 9.1 For $k \geq 1$,

$$
A_{k+1}=\left(Q^{(k)}\right)^{T} A Q^{(k)}, \quad A^{k}=\left(Q_{1} \cdots Q_{k}\right)\left(R_{k} \cdots R_{1}\right)=: Q^{(k)} R^{(k)}
$$

Proof : recall $A_{k+1}=Q_{k}^{T} A_{k} Q_{k}$, repeat.

Proof by induction: $k=1$ trivial.
Suppose $A^{k-1}=Q^{(k-1)} R^{(k-1)}$. We have

$$
A_{k}=\left(Q^{(k-1)}\right)^{T} A Q^{(k-1)}=Q_{k} R_{k} .
$$

Then $A Q^{(k-1)}=Q^{(k-1)} Q_{k} R_{k}$, and so

$$
A^{k}=A Q^{(k-1)} R^{(k-1)}=Q^{(k-1)} Q_{k} R_{k} R^{(k-1)}=Q^{(k)} R^{(k)} \square
$$

QR algorithm and power method We can deduce that the QR algorithm is closely related to the power method. Let's try explain the connection. By Theorem 9.1, $Q^{(k)} R^{(k)}$ is the QR factorisation of $A^{k}$ : as we saw in the analysis of the power method, the columns of $A^{k}$ are 'dominated by the leading eigenvector' $x_{1}$, where $A x_{1}=\lambda_{1} x_{1}$.

In particular, consider $A^{k}[1,0, \ldots, 0]=A^{k} e_{n}$ :

- $A^{k} e_{n}=R^{(k)}(1,1) Q^{(k)}(:, 1)$, which is parallel to the first column of $Q^{(k)}$.
- By the analysis of the power method, this implies $Q^{(k)}(:, 1) \rightarrow x_{1}$
- Hence by $A_{k+1}=\left(Q^{(k)}\right)^{A} Q^{(k)}, A_{k}(:, 1) \rightarrow\left[\lambda_{1}, 0, \ldots, 0\right]^{T}$.

This tells us why the QR algorithm would eventually compute the eigenvalues-at least the dominant ones. One can even go further and argue that once the dominant eigenvalue has converged, we can expect the next dominant eigenvalue to start converging too-as due to the nature of the QR algorithm based on orthogonal transformations, we are then working in the orthogonal complement; and so on and so forth until the matrix $A$ becomes upper triangular (and in the normal case this becomes diagonal), completing the solution of the eigenvalue problem. But there is much better news.

QR algorithm and inverse power method We have seen that the QR algorithm is related to the power method. We have also seen that the power method can be improved by a shift-and-invert technique. A natural question arises: can we introduce a similar technique in the QR algorithm? This question has an elegant and remarkable answer: not only is this possible but it is possible without ever inverting a matrix or solving a linear system (isn't that incredible!?).

Let's try and explain this. We start with the same expression for the QR iterates:

$$
A^{k}=\left(Q_{1} \cdots Q_{k}\right)\left(R_{k} \cdots R_{1}\right)=: Q^{(k)} R^{(k)}, \quad A_{k+1}=\left(Q^{(k)}\right)^{T} A Q^{(k)}
$$

Now take inverse: $A^{-k}=\left(R^{(k)}\right)^{-1}\left(Q^{(k)}\right)^{T}$,
Conjugate transpose: $\left(A^{-k}\right)^{T}=Q^{(k)}\left(R^{(k)}\right)^{-*}$
$\Rightarrow \mathrm{QR}$ factorisation of matrix $\left(A^{-k}\right)^{T}$ with eigvals $r\left(\lambda_{i}\right)=\lambda_{i}^{-k}$
$\Rightarrow$ Connection also with (unshifted) inverse power method. Note that no matrix inverse is performed in the algorithm.

- This means the final column of $Q^{(k)}$ converges to minimum left eigenvector $v_{n}$ with rate $\frac{\left|\lambda_{n}\right|}{\left|\lambda_{n-1}\right|}$, hence $A_{k}(n,:) \rightarrow\left[0, \ldots, 0, \lambda_{n}\right]$.
- (Very) fast convergence if $\left|\lambda_{n}\right| \ll\left|\lambda_{n-1}\right|$.
- Can we achieve this situation? Yes by shifts.

QR algorithm with shifts and shifted inverse power method We are now ready to reveal the connection between the shift-and-invert power method and the QR algorithm with shifts.

First, here is the QR algorithm with shifts.

```
Algorithm 9.2 The QR algorithm with shifts for finding the Schur decomposition of a
square matrix \(A\).
    1: Set \(A_{1}=A, k=1\).
    2: \(A_{k}-s_{k} I=Q_{k} R_{k}\) (QR factorisation)
    3: \(A_{k+1}=R_{k} Q_{k}+s_{k} I, \quad k \leftarrow k+1\), repeat.
```

1. $A_{k}-s_{k} I=Q_{k} R_{k}$ (QR factorisation)
2. $A_{k+1}=R_{k} Q_{k}+s_{k} I, \quad k \leftarrow k+1$, repeat.

Roughly, if $s_{k} \approx \lambda_{n}$, then $A_{k+1} \approx\left[\begin{array}{lllll}* & * & * & * & * \\ * & * & * & * & * \\ * & * & * & * & * \\ * & * & * & * & * \\ & & & & \lambda_{n}\end{array}\right]$ by the argument just made.

## Theorem 9.2

$$
\prod_{i=1}^{k}\left(A-s_{i} I\right)=Q^{(k)} R^{(k)}\left(=\left(Q_{1} \cdots Q_{k}\right)\left(R_{k} \cdots R_{1}\right)\right)
$$

Proof: Suppose true for $k-1$. Then the QR alg computes $\left(Q^{(k-1)}\right)^{T}\left(A-s_{k} I\right) Q^{(k-1)}=Q_{k} R_{k}$, so $\left(A-s_{k} I\right) Q^{(k-1)}=Q^{(k-1)} Q_{k} R_{k}$, hence

$$
\prod_{i=1}^{k}\left(A-s_{i} I\right)=\left(A-s_{k} I\right) Q^{(k-1)} R^{(k-1)}=Q^{(k-1)} Q_{k} R_{k} R^{(k-1)}=Q^{(k)} R^{(k)} .
$$

Inverse conjugate transpose: $\prod_{i=1}^{k}\left(A-s_{i} I\right)^{-*}=Q^{(k)}\left(R^{(k)}\right)^{-*}$.

- This means the algorithm (implicitly) finds the QR factorisation of a matrix with eigvals $r\left(\lambda_{j}\right)=\prod_{i=1}^{k} \frac{1}{\lambda_{j}-s_{i}}$.
- The rate is now the ratio of the smallest $\prod_{i=1}^{k}\left(\bar{\lambda}_{j}-s_{i}\right)$ over the runner-up; very fast if $s_{i} \approx \lambda_{j}$
- This reveals the intimate connection between shifted QR and shifted inverse power method, hence rational approximation .
- Ideally, we would like to choose $s_{k} \approx \lambda_{n}$ to accelerate convergence. This is done by choosing $s_{k}$ to be the bottom-right corner entry ${ }^{19}$ of $A_{k}$; which is sensible given that with the QR algorithm, it tends to converge to an eigenvalue rapidly (with a few steps of the unshifted QR , it tends to converge to the smallest eigenvalue).


### 9.2 QR algorithm preprocessing: reduction to Hessenberg form

We've seen the QR iterations drives colored entries to 0 (esp. red ones)

$$
A=\left[\begin{array}{lllll}
* & * & * & * & * \\
* & * & * & * & * \\
* & * & * & * & * \\
* & * & * & * & * \\
* & * & * & * & *
\end{array}\right]
$$

- Hence $A_{n, n} \rightarrow \lambda_{n}$, so choosing $s_{k}=A_{n, n}$ is sensible.
- This reduces $\# \mathrm{QR}$ iterations to $O(n)$ (empirical but reliable estimate).
- But each iteration of the QR algorithm is $O\left(n^{3}\right)$ for QR , overall $O\left(n^{4}\right)$.
- We next discuss a preprocessing technique to reduce the overall cost to $O\left(n^{3}\right)$.

The idea is to initially apply a series of deterministic orthogonal transformations that reduces the matrix to a form that is closer to upper triangular. This is done before starting the QR iterates, hence called a preprocessing step.

More precisely, to improve the cost of QR factorisation, we first reduce the matrix via orthogonal Householder transformations as follows:

$$
A=\left[\begin{array}{ccccc}
* & * & * & * & * \\
* & * & * & * & * \\
* & * & * & * & * \\
* & * & * & * & * \\
* & * & * & * & *
\end{array}\right], \quad H_{1} A=\left[\begin{array}{lllll}
* & * & * & * & * \\
* & * & * & * & * \\
* & * & * & * \\
* & * & * & * \\
* & * & * & *
\end{array}\right], \quad H_{1}=I-2 v_{1} v_{1}^{T}, v_{1}=\left[\begin{array}{c}
0 \\
* \\
* \\
* \\
*
\end{array}\right]
$$

Then $H_{1} A H_{1}=\left[\begin{array}{rrrrr}* & * & * & * & * \\ * & * & * & * & * \\ * & * & * & * \\ & * & * & * & * \\ & * & * & * & *\end{array}\right]$. Repeat with $H_{2}=I-2 v_{2} v_{2}^{T}, v_{2}=[0,0, *, *, *]^{T}, \ldots$ :

$$
H_{2} H_{1} A H_{1} H_{2}=\left[\begin{array}{lllll}
* & * & * & * & * \\
* & * & * & * & * \\
& * & * & * & * \\
& * & * & * \\
& & * & * & *
\end{array}\right], \quad H_{3} H_{2} H_{1} A H_{1} H_{2} H_{3}=\left[\begin{array}{lllll}
* & * & * & * & * \\
* & * & * & * & * \\
& * & * & * & * \\
& & * & * & * \\
& & & * & *
\end{array}\right],
$$

[^17]\[

A=\left[$$
\begin{array}{lllll}
* & * & * & * & * \\
* & * & * & * & * \\
* & * & * & * & * \\
* & * & * & * & * \\
* & * & * & * & *
\end{array}
$$\right] \xrightarrow[\rightarrow]{H_{1}}\left[$$
\begin{array}{lllll}
* & * & * & * & * \\
* & * & * & * & * \\
* & * & * & * \\
* & * & * & * \\
* & * & * & *
\end{array}
$$\right] \xrightarrow[\rightarrow]{H_{2}}\left[$$
\begin{array}{lllll}
* & * & * & * & * \\
* & * & * & * & * \\
& * & * & * & * \\
& * & * & * \\
& * & * & *
\end{array}
$$\right] \xrightarrow[\rightarrow]{H_{3}} \underset{\rightarrow}{H_{n-2}}\left[$$
\begin{array}{ccccc}
H_{n} & * & * & * & * \\
* & * & * & * \\
& * & * & * \\
& & * & *
\end{array}
$$\right] .
\]

We have thus transformed the matrix $A$ into upper Hessenberg form.

- Crucially, the Hessenberg structure is preserved throughout the QR algorithm: if $A_{1}=$ $Q R$ Hessenberg, then so is $A_{2}=R Q$.
- Using Givens rotations, each QR iter is $O\left(n^{2}\right)$, not $O\left(n^{3}\right)$.

We are thus going to first reduce $A$ to Hessenberg form via a sequence of orthogonal transformations (thus preserving eigenvalues), and then apply the QR algorithm with shifts to the Hessenberg matrix. The remaining task is thus to drive the subdiagonal elements $*$ to 0 . $A$ is thus reduced to a triangular form, revealing the Schur decomposition of $A$ (if one traces back all the orthogonal transformations employed).

Once the bottom-right is converged $|*|<\epsilon$, we take

$$
\left[\begin{array}{lllll}
* & * & * & * & * \\
* & * & * & * & * \\
& * & * & * & * \\
& & * & * & * \\
& & & * & *
\end{array}\right] \approx\left[\begin{array}{lllll}
* & * & * & * & * \\
* & * & * & * & * \\
& * & * & * & * \\
& & * & * & * \\
& & & & *
\end{array}\right]
$$

and continue with shifted QR on $(n-1) \times(n-1)$ block, repeat. This process is called deflation.

- Empirically, the shifted QR algorithm needs $2--4$ iterations per eigenvalue. Overall, the cost is $O\left(n^{3}\right), \approx 25 n^{3}$ flops.


### 9.2.1 The (shifted) QR algorithm in action

Let's see how the QR algorithm works with a small example. The plots show the convergence of the subdiagonal entries $\left|A_{i+1, i}\right|$ (note that their convergence signifies the convergence of the QR algorithm as we initially reduce the matrix to Hessenberg form).


In light of the connection to rational functions as discussed above, here we plot the underlying functions (red dots: eigvals). The idea here is that we want the functions to take large values at the target eigenvalue (at the current iterate) in order to accelerate convergence.

### 9.2.2 (Optional) QR algorithm: other improvement techniques

We have completed the description of the main ingredients of the QR algorithm. Nonetheless, there are a few more bells and whistles that go into to a production code. We will not get into too much detail but here is a list of the key players.

- Double-shift strategy for $A \in \mathbb{R}^{n \times n}$
$-(A-s I)(A-\bar{s} I)=Q R$ using only real arithmetic
- Aggressive early deflation
[Braman-Byers-Mathias 2002 [5]]
- Examine lower-right (say $100 \times 100$ ) block instead of $(n, n-1)$ element
- dramatic speedup $(\approx \times 10)$
- Balancing $A \leftarrow D A D^{-1}, D$ : diagonal
- Aims at reducing $\left\|D A D^{-1}\right\|$ : often yields better-conditioned eigenvalues.

Finally, let us mention a cautionary tale:

- For nonsymmetric $A$, global convergence is NOT established(!)
- Of course it always converges in practice.. proving convergence is another big open problem in numerical linear algebra.
- For symmetric $A$, global convergence analysis is available, see [28, Ch. 8].


## 10 QR algorithm continued

### 10.1 QR algorithm for symmetric $A$

so far we have not assumed anything about the matrix besides that it is square. This is for good reason-because the QR algorithm is applicable to solving any eigenvalue problem. However, in many situations the eigenvalue problem is structured. In particular the case where the matrix is symmetric arises very frequently in practice and it comes with significant simplification of the QR algorithm, so it deserves a special mention.

- Most importantly, symmetry immediately implies that the initial reduction to Hessenberg form reduces $A$ to tridiagonal:

$$
A=\left[\begin{array}{lllll}
* & * & * & * & * \\
* & * & * & * & * \\
* & * & * & * & * \\
* & * & * & * & * \\
* & * & * & * & *
\end{array}\right] \xrightarrow[\rightarrow]{Q_{1}}\left[\begin{array}{lllll}
* & * & & & \\
* & * & * & * & * \\
& * & * & * & * \\
& * & * & * & * \\
& * & * & * & *
\end{array}\right] \xrightarrow[\rightarrow]{Q_{2}}\left[\begin{array}{lllll}
* & * & & & \\
* & * & * & & \\
& * & * & * & * \\
& * & * & * \\
& * & * & *
\end{array}\right] \xrightarrow[\rightarrow]{Q_{3}}\left[\begin{array}{lllll}
* & * & & \\
* & * & * & \\
& * & * & * & \\
& & * & * & * \\
& & & * & *
\end{array}\right]
$$

- QR steps for tridiagonal: requires $O(n)$ flops instead of $O\left(n^{2}\right)$ per step.
- Powerful alternatives are available for tridiagonal eigenvalue problem (divide-conquer [Gu-Eisenstat 95], HODLR [Kressner-Susnjara 19],...)
- Cost: $\frac{4}{3} n^{3}$ flops for eigvals, $\approx 10 n^{3}$ for eigvecs (store Givens rotations to compute eigvecs).
- Another approach (nonexaminable): spectral divide-and-conquer (Nakatsukasa-Freund, Higham); which is all about using a carefully chosen rational approximation to find orthogonal matrices $V_{i}$ such that


### 10.2 Computing the SVD: Golub-Kahan's bidiagonalisation algorithm

The key ideas of the QR algorithm turns out to be applicable for the computation of the SVD. Clearly this is a major step-given the importance of the SVD. In particular the SVD algorithm has strong connections to the symmetric QR algorithm. This is perhaps not surprising given the strong connection between the SVD and symmetric eigenvalue problems as we discussed previously.

A noteworthy difference is that in the SVD $A=U \Sigma V^{T}$ the two matrices $U$ and $V$ are allowed to be different. The algorithm respects this and instead of initially reducing the matrix to tridiagonal form, it reduces it to a so-called bidiagonal form.

Here is how it works. Apply Householder reflectors from left and right (different ones) to bidiagonalise

$$
\begin{aligned}
& A \rightarrow B=H_{L, n} \cdots H_{L, 1} A H_{R, 1} H_{R, 2} \cdots H_{R, n-2}
\end{aligned}
$$

- Since the transformations are all orthogonal multiplications, singular values are preserved $\sigma_{i}(A)=\sigma_{i}(B)$.
- Once bidiagonalised, one can complete the SVD as follows:
- Mathematically, do QR alg on $B^{T} B$ (symmetric tridiagonal)
- More elegant: divide-and-conquer [Gu-Eisenstat 1995] or dqds algorithm [FernandoParlett 1994] (nonexaminable)
- Cost: $\approx 4 m n^{2}$ flops for singular values $\Sigma, \approx 20 m n^{2}$ flops to also compute the singular vectors $U, V$.


## 10.3 (Optional but important) QZ algorithm for generalised eigenvalue problems

An increasingly important class of eigenvalue problems is the so-called generalised eigenvalue problems involving two matrices. You have probably not seen them before, but the number of applications that boil down to a generalised eigenvalue problem has been increasing rapidly.

A generalised eigenvalue problem is of the form

$$
A x=\lambda B x, \quad A, B \in \mathbb{C}^{n \times n}
$$

- The matrices $A, B$ are given. The goal is to find the eigenvalues $\lambda$ and their corresponding eigenvectors $x$.
- There are usually (incl. when $B$ is nonsingular) $n$ eigenvalues, which are the roots of $\operatorname{det}(A-\lambda B)$.
- When $B$ is invertible, one can reduce the problem to $B^{-1} A x=\lambda x$.
- Important case: $A, B$ symmetric, $B$ positive definite: in this case $\lambda$ are all real.

QZ algorithm: look for unitary $Q, Z$ s.t. $Q A Z, Q B Z$ both upper triangular.

- Then $\operatorname{diag}(Q A Z) / \operatorname{diag}(Q B Z)$ are the eigenvalues.
- Algorithm: first reduce $A, B$ to Hessenberg-triangular form.
- Then implicitly do QR to $B^{-1} A$ (without inverting $B$ ).
- Cost: $\approx 50 n^{3}$.
- See [14] for details.


## 10.4 (Optional) Tractable eigenvalue problems

Beyond generalised eigenvalue problems there are more exotic generalisations of eigenvalue problems and reliable algorithms have been proposed for solving them.

Thanks to the remarkable developments in NLA, the following problems are 'tractable' in that reliable algorithms exist for solving them, at least when the matrix size $n$ is modest (say in the thousands).

- Standard eigenvalue problems $A x=\lambda x$
- symmetric ( $4 / 3 n^{3}$ flops for eigvals, $+9 n^{3}$ for eigvecs)
- nonsymmetric ( $10 n^{3}$ flops for eigvals, $+15 n^{3}$ for eigvecs)
- SVD $A=U \Sigma V^{T}$ for $A \in \mathbb{R}^{m \times n}:\left(\frac{8}{3} m n^{2}\right.$ flops for singvals, $+20 m n^{2}$ for singvecs)
- Generalised eigenvalue problems $A x=\lambda B x, A, B \in \mathbb{C}^{n \times n}$
- Polynomial eigenvalue problems, e.g. (degree $k=2$ ) $P(\lambda) x=\left(\lambda^{2} A+\lambda B+C\right) x=0$, $A, B, C \in \mathbb{C}^{n \times n}: \approx 20(n k)^{3}$
- Nonlinear problems, e.g. $N(\lambda) x=(A \exp (\lambda)+B) x=0$
- often solved via approximating by polynomial $N(\lambda) \approx P(\lambda)$
- more difficult: $A(x) x=\lambda x$ : eigenvector nonlinearity

Further speedup is often possible when structure present (e.g. sparse, low-rank)

## 11 Iterative methods: introduction

This section marks a point of departure from previous sections. So far we've been discussing direct methods. Namely, we've covered the LU-based linear system solution and the QR algorithm for eigenvalue problems ${ }^{20}$.

Direct methods are

- Incredibly reliable, and backward stable
- Works like magic if $n \lesssim 10000$
- But not if $n$ is larger!

A 'big' matrix problem is one for which direct methods aren't feasible. Historically, as computers become increasingly faster, roughly

- 1950: $n \geq 20$
- 1965: $n \geq 200$
- 1980: $n \geq 2000$
- 1995: $n \geq 20000$
- 2010: $n \geq 100000$
- 2020: $n \geq 500000$ ( $n \geq 50000$ on a standard desktop)
was considered 'too large for direct methods'. While it's clearly good news that our ability to solve problems with direct methods has been improving, the scale of problems we face in data science has been growing at the same (or even faster) pace! For such problems, we need to turn to alternative algorithms: we'll cover iterative and randomised methods. We first discuss iterative methods, with a focus on Krylov subspace methods.

Direct vs. iterative methods Broadly speaking, the idea of iterative methods is to:

- Gradually refine the solution iteratively.
- Each iteration should be (a lot) cheaper than direct methods, usually $O\left(n^{2}\right)$ or less.
- Iterative methods can be (but not always) much faster than direct methods.
- Tends to be (slightly) less robust, nontrivial/problem-dependent analysis.
- Often, after $O\left(n^{3}\right)$ work it still gets the exact solution (ignoring roundoff errors). But one would hope to get a (acceptably) good solution long before that!


Figure 2: Rough comparison of direct vs. iterative methods, image from [Trefethen-Bau [33]]

Each iteration of most iterative methods is based on multiplying $A$ to a vector; clearly cheaper than an $O\left(n^{3}\right)$ direct algorithm. We'll focus on Krylov subspace methods. (Other iterative methods we won't get into details include the Gauss-Seidel, SOR and Chebyshev semi-iterative methods.)

### 11.1 Polynomial approximation: basic idea of Krylov

The big idea behind Krylov subspace methods is to approximate the solution in terms of a polynomial of the matrix times a vector. Namely, in Krylov subspace methods, we look for an (approximate) solution $\hat{x}$ (for $A x=b$ or $A x=\lambda x$ ) of the form (after $k$ th iteration)

$$
\hat{x}=p_{k-1}(A) v,
$$

where $p_{k-1}$ is a polynomial of degree (at most) $k-1$, that is, $p_{k-1}(z)=\sum_{i=0}^{k-1} c_{i} z^{i}$ for some coefficients $c_{i} \in \mathbb{C}$, and $v \in \mathbb{R}^{n}$ is the initial vector (usually $v=b$ for linear systems, for eigenproblems $v$ is usually a random vector). That is, $p_{k-1}(A)=\sum_{i=0}^{k-1} c_{i} A^{i}$.

Natural questions:

- Why would this be a good idea?
- Clearly, 'easy' to compute.
- One example: recall power method $\hat{x}=A^{k-1} v=p_{k-1}(A) v$ Krylov finds a "better/optimal" polynomial $p_{k-1}(A)$. When the eigenvalues of $A$ are well-behaved, we'll see that $O(1)$ iterations suffices for convergence.
- We'll see more cases where Krylov is powerful.
- How to turn this idea into an algorithm?
- Find an orthonormal basis: Arnoldi (next), Lanczos.

[^18]
### 11.2 Orthonormal basis for $\mathcal{K}_{k}(A, b)$

Goal: Find approximate solution $\hat{x}=p_{k-1}(A) b$, i.e. in Krylov subspace

$$
\mathcal{K}_{k}(A, b):=\operatorname{span}\left(\left[b, A b, A^{2} b, \ldots, A^{k-1} b\right]\right)
$$

You would want to convince yourself that any vector in the Krylov subspace can be written as a polynomial of $A$ times the vector $b$.

An important and non-trivial step towards finding a good solution is to form an orthonormal basis for the Krylov subspace.

First step: form an orthonormal basis $Q$, s.t. solution $x \in \mathcal{K}_{k}(A, b)$ can be written as $x=Q y$

- Naive idea: Form matrix $\left[b, A b, A^{2} b, \ldots, A^{k-1} b\right]$, then compute its QR factorisation.
- $\left[b, A b, A^{2} b, \ldots, A^{k-1} b\right]$ is usually terribly conditioned! Dominated by leading eigvec - $Q$ is therefore extremely ill-conditioned, and hence inaccurately computed
- A much better solution is to do the Arnoldi iteration
- Multiply $A$ once at a time to the latest orthonormal vector $q_{i}$
- Then orthogonalise $A q_{i}$ against previous $q_{j}$ 's $(j=1, \ldots, i-1)$ (as in GramSchmidt).


### 11.3 Arnoldi iteration

Here is a pseudocode of the Arnold iteration. Essentially, what it does is multiply the matrix $A$, orthogonalise against the previous vectors, and repeat. This is massively (often exponentially wrt conditioning) better than multiplying $A k-1$ times and then orthogonalising.

```
Algorithm 11.1 The Arnoldi iteration for finding an orthonormal basis for Krylov subspace
\(\mathcal{K}_{k}(A, b)\). Given \(A \in \mathbb{R}^{n \times n}\) and \(b \in \mathbb{R}^{n}\), outputs orthonormal \(\left[q_{1}, q_{2}, \ldots,\right]\) and Hessenberg
\(H_{i j}=h_{i j}\).
    Set \(q_{1}=b /\|b\|_{2}\)
    For \(k=1,2, \ldots\),
        set \(v=A q_{k}\)
        for \(j=1,2, \ldots, k\)
            \(h_{j k}=q_{j}^{T} v, v=v-h_{j k} q_{j} \%\) orthogonalise against \(q_{j}\) via (modified) Gram-Schmidt
        end for
        \(h_{k+1, k}=\|v\|_{2}, q_{k+1}=v / h_{k+1, k}\)
    End for
```

Let's verify a basic result about Arnoldi, that the vectors $q_{i}$ computed by Arnoldi spans the Krylov subspace.

Theorem 11.1 Suppose in Algorithm 9.2 that $h_{k+1, k} \neq 0$ for $k=1, \ldots, \ell$. Then for $k=$ $1, \ldots, \ell$,

$$
\begin{equation*}
\operatorname{Span}\left(q_{1}, \ldots, q_{k}\right)=\mathcal{K}_{k}(A, b) \tag{10}
\end{equation*}
$$

Proof: Induction on $\ell$. When $\ell=1$ the result is trivial. Note that $q_{1}=q_{\ell}=p_{\ell-1}(A) b$, where $p_{\ell-1}$ is a polynomial of degree $\ell-1=0$ (constant).

We take the induction hypothesis to be that (10) holds for $k=1, \ldots, \hat{\ell}$, and $q_{\hat{\ell}}$ is of exact degree $\hat{\ell}-1$.

Then the $(\hat{\ell}+1)$ th Arnoldi iteration gives $q_{\hat{\ell}+1}=\frac{1}{h_{\hat{\ell}+1, \hat{\ell}}}\left(A q_{\hat{\ell}}-\sum_{j=1}^{\hat{\ell}} h_{j, \hat{\ell}} q_{j}\right)$, which is of exact degree $\hat{\ell}$.

Theorem 11.1 assumes that $h_{k+1, k} \neq 0$; if this doesn't hold it's not bad news, we are (usually) actually in a lucky situation. For example, this means the solution for $A x=b$ lies in the subspace (so with a reasonable method like GMRES below, we obtain the exact solution), and when solving an eigenvalue problem, an invariant subspace has been found. This situation is called a happy breakdown (and once this happens, we cannot/don't proceed further with Arnoldi).

- Hence after $k$ steps, $A Q_{k}=Q_{k+1} \tilde{H}_{k}=Q_{k} H_{k}+q_{k+1}\left[0, \ldots, 0, h_{k+1, k}\right]$, with $Q_{k}=$ $\left[q_{1}, q_{2}, \ldots, q_{k}\right], Q_{k+1}=\left[Q_{k}, q_{k+1}\right], \operatorname{span}\left(Q_{k}\right)=\operatorname{span}\left(\left[b, A b, \ldots, A^{k-1} b\right]\right)$

$$
A Q_{k}=Q_{k+1} \tilde{H}_{k}, \quad \tilde{H}_{k}=\underbrace{\left[\begin{array}{cccc}
h_{1,1} & h_{1,2} & \cdots & h_{1, k} \\
h_{2,1} & h_{2,2} & \cdots & h_{2, k} \\
& \ddots & & \vdots \\
& & h_{k, k-1} & h_{k, k} \\
& & h_{k+1, k}
\end{array}\right]}_{\mathbb{R}^{(k+1) \times k}}, \quad Q_{k+1}^{T} Q_{k+1}=I_{k+1}
$$

- Cost is $k A$-multiplications $+O\left(k^{2}\right)$ inner products ( $O\left(n k^{2}\right)$ flops)


## 12 Arnoldi and GMRES for $\boldsymbol{A x}=\boldsymbol{b}$

This is an exciting section as we will describe the GMRES (Generalised Minimal RESidual) algorithm. This algorithm has been so successful that in the 90 s the paper that introduced it was the most cited paper in all of applied mathematics.

GMRES attempts to find an approximate solution in the Krylov subspace that is in some sense the best possible, in that the residual is minimised.

Idea (very simple!): minimise residual in Krylov subspace:
[Saad-Schulz 86 [29]]

$$
\begin{equation*}
x=\operatorname{argmin}_{x \in \mathcal{K}_{k}(A, b)}\|A x-b\|_{2} . \tag{11}
\end{equation*}
$$

In order to solve this, The algorithm cleverly takes advantage of the structure afforded by Arnoldi.

Algorithm: Given $A Q_{k}=Q_{k+1} \tilde{H}_{k}$ and writing $x=Q_{k} y$, rewrite as

$$
\begin{aligned}
\min _{y}\left\|A Q_{k} y-b\right\|_{2} & =\min _{y}\left\|Q_{k+1} \tilde{H}_{k} y-b\right\|_{2} \\
& =\min _{y}\left\|\left[\begin{array}{c}
\tilde{H}_{k} \\
0
\end{array}\right] y-\left[\begin{array}{c}
Q_{k+1}^{T} \\
Q_{k+1, \perp}^{T}
\end{array}\right] b\right\|_{2} \\
& =\min _{y}\left\|\left[\begin{array}{c}
\tilde{H}_{k} \\
0
\end{array}\right] y-\right\| b\left\|_{2} e_{1}\right\|_{2}, \quad e_{1}=[1,0, \ldots, 0]^{T} \in \mathbb{R}^{n}
\end{aligned}
$$

( where $\left[Q_{k+1}, Q_{k+1, \perp}\right]$ is orthogonal; we're using the same trick as in least-squares)

- Minimised when $\left\|\tilde{H}_{k} y-Q_{k+1}^{T} b\right\|_{2}$ is minimised; this is a Hessenberg least-squares problem.
- Solve via the QR-based approach described in Section 6.5. Here it's even simpler as the matrix is Hessenberg: $k$ Givens rotations yields the QR factorisation, then a triangular solve will complete the solution, so $O\left(k^{2}\right)$ work in addition to Arnoldi; recall Section 6.4.

Notice how the Arnoldi decomposition reduced the problem (11) effortlessly to a very convenient one $\left\|\tilde{H}_{k} y-Q_{k+1}^{T} b\right\|_{2}$; this wouldn't be the case if we "merely" had a QR factorisation of the Krylov matrix $\left[b, A b, A^{2} b, \ldots, A^{k-1} b\right]$. We'll see this phenomenon many times in this section and next.

### 12.1 GMRES convergence: polynomial approximation

We now study the convergence of GMRES. As with any Krylov subspace method, the analysis is based on the theory of polynomial approximation.

Theorem 12.1 (GMRES convergence) Assume that $A$ is diagonalisable, $A=X \Lambda X^{-1}$. Then the $k$ th GMRES iterate $x_{k}$ satisfies

$$
\left\|A x_{k}-b\right\|_{2} \leq \kappa_{2}(X) \min _{p \in \mathcal{P}_{k}, p(0)=1} \max _{z \in \lambda(A)} \mid p(z)\| \| b \|_{2} .
$$

Proof: Recall that $x_{k} \in \mathcal{K}_{k}(A, b) \Rightarrow x_{k}=p_{k-1}(A) b$, where $p_{k-1}$ is a polynomial of degree at most $k-1$. Hence GMRES solution is

$$
\begin{aligned}
\min _{x_{k} \in \mathcal{K}_{k}(A, b)}\left\|A x_{k}-b\right\|_{2} & =\min _{p_{k-1} \in \mathcal{P}_{k-1}}\left\|A p_{k-1}(A) b-b\right\|_{2} \\
& =\min _{\tilde{p} \in \mathcal{P}_{k}, \tilde{\mathcal{p}}(0)=0}\|(\tilde{p}(A)-I) b\|_{2} \\
& =\min _{p \in \mathcal{P}_{k}, p(0)=1}\|p(A) b\|_{2}
\end{aligned}
$$

If $A$ is diagonalisable, $A=X \Lambda X^{-1}$,

$$
\begin{aligned}
\|p(A)\|_{2} & =\left\|X p(\Lambda) X^{-1}\right\|_{2} \leq\|X\|_{2}\left\|X^{-1}\right\|_{2}\|p(\Lambda)\|_{2} \\
& =\kappa_{2}(X) \max _{z \in \lambda(A)}|p(z)|
\end{aligned}
$$

(recall that $\lambda(A)$ is the set of eigenvalues of $A$ )
Noting that $\|p(A) b\|_{2} \leq\|p(A)\|_{2}\|b\|_{2}$, here is the interpretation of this analysis: We would like to find a polynomial s.t. $p(0)=1$ and $|p|$ is small at the eigenvalues of $A$, that is, $\left|p\left(\lambda_{i}\right)\right|$ is small for all $i$.

The question becomes, what type of eigenvalues are 'nice' for this to hold?
GMRES example $G$ : Gaussian random matrix $\left(G_{i j} \sim N(0,1)\right.$, i.i.d.) $G / \sqrt{n}, n=1000$ : eigenvalues in unit disk.




In the left case GMRES converges rapidly, whereas it does not in the right case. The difference can be explained by noting that in the left case, there is a polynomial $p(z)=$ $2^{-k}(z-2)^{k}$ that is very small in the yellow disk while $p(0)=1$, for which $\max _{z \in \lambda(A)}|p(z)| \lesssim$ $2^{-k}$. By contrast, in the right example we don't have such an obvious choice; and indeed there isn't a polynomial for which $\max _{z \in \lambda(A)}|p(z)|$ decays rapidly.

Now suppose the eigenvalues are real (and positive); this would be the case obviously if $A=A^{T}$.


In this case the naive convergence bound $1 / 2^{k}$ obtained via the polynomial $p(z)=2^{-k}(z-$ $2)^{k}$ isn't very sharp. The reason is that a more specific polynomial can "concentrate its energy" on the real interval [1,3]; we'll return to this shortly in Section 13.3 in the analysis of the CG algorithm.

Initial vector. Sometimes a good initial guess $x_{0}$ for $x_{*}$ is available. In this case we take the initial residual $r_{0}=A x_{0}-b$, and work in the affine space $x_{k}=x_{0}+\mathcal{K}_{k}\left(A, r_{0}\right)$. All the analysis above can be modified readily to allow for this situation with essentially the same conclusion (clearly if one has a good $x_{0}$ by all means that should be used).

### 12.2 When does GMRES converge fast?

Recall GMRES solution satisfies (assuming $A$ is diagonalisable+nonsingular)

$$
\min _{x \in \mathcal{K}_{k}(A, b)}\|A x-b\|_{2}=\min _{p \in \mathcal{P}_{k}, p(0)=1}\|p(A) b\|_{2} \leq \kappa_{2}(X) \max _{z \in \lambda(A)}|p(z)|\|b\|_{2} .
$$

As the example above illustrates, $\max _{z \in \lambda(A)}|p(z)|$ is small when

- $\lambda(A)$ are clustered away from 0
- a good $p$ can be found quite easily
- e.g. example above

Another interesting case is:

- When $\lambda(A)$ takes $k(\ll n)$ distinct values
- Then convergence in $k$ GMRES iterations; to see this, note that a polynomial $p$ can be found such that $p\left(\lambda_{i}\right)=0$ at the $k$ eigenvalues, that is, $p(z)=C \prod_{i=1}^{k}\left(z-\lambda_{i}\right)$, and choose $C$ such that $p(0)=1$.


### 12.3 Preconditioning for GMRES

We've seen that GMRES is great if the eigenvalues are clustered away from 0 . If this is not true, GMRES can really require close to (or equal to!) the full $n$ iterations for convergence. This is undesirable - at this point we've spent more computation than a direct method! We need a workaround.

The idea of preconditioning is to instead of solving

$$
A x=b
$$

find a "preconditioner" $M \in \mathbb{R}^{n \times n}$ and solve

$$
M A x=M b
$$

Of course, $M$ cannot be an arbitrary matrix. It has to be chosen very carefully. Desiderata of $M$ are

- $M$ is simple enough s.t. applying $M$ to a vector is easy (note that each GMRES iteration requires $M A$-multiplication to a vector, say $M A x$, which is done as $M(A x))$, and one of

1. $M A$ has clustered eigenvalues away from 0 .
2. $M A$ has a small number of distinct nonzero eigenvalues.
3. $M A$ is well-conditioned $\kappa_{2}(M A)=O(1)$; then solve the normal equation $(M A)^{T} M A x=$ $(M A)^{T} M b$.

## Preconditioners: examples

- ILU (Incomplete LU) preconditioner: $A \approx L U, M=(L U)^{-1}=U^{-1} L^{-1}, L, U$ 'as sparse as $A^{\prime} \Rightarrow M A \approx I$ (hopefully; 'cluster away from 0 ').
- For $\tilde{A}=\left[\begin{array}{cc}A & B \\ C & 0\end{array}\right]$, set $M=\left[\begin{array}{ll}A^{-1} & \\ & \left(C A^{-1} B\right)^{-1}\end{array}\right]$. Then if $M$ nonsingular, $M \tilde{A}$ has eigvals $\in\left\{1, \frac{1}{2}(1 \pm \sqrt{5})\right\} \Rightarrow 3$-step convergence. [Murphy-Golub-Wathen 2000 [23]]
- Multigrid-based, operator preconditioning, ...
- A "perfect" preconditioner is $M=A^{-1}$; as then preconditioned GMRES will converge in one step. Obviously this $M$ isn't easy to apply (if it was then we're done!). Preconditioning, therefore, can be regarded as an act of efficiently approximating the inverse.

Finding effective preconditioners is a never-ending research topic.
Prof. Andy Wathen is our Oxford expert!

### 12.4 Restarted GMRES

Another practical GMRES technique is restarting. For $k$ iterations, GMRES costs $k$ matrix multiplications $+O\left(n k^{2}\right)$ for orthogonalisation $\rightarrow$ Arnoldi eventually becomes expensive.

Practical solution: restart by solving 'iterative refinement':

1. Stop GMRES after $k_{\text {max }}$ (prescribed) steps to get approx. solution $\hat{x}_{1}$.
2. Solve $A \tilde{x}=b-A \hat{x}_{1}$ via GMRES. (This is a linear system with a different right-hand side).
3. Obtain solution $\hat{x}_{1}+\tilde{x}$.

Sometimes multiple restarts are needed.

## 13 Symmetric case: Lanczos and Conjugate Gradient method for $A x=b, A \succ 0$

As is often the case, when $A=A^{T}$ the situation simplifies significantly. First, the Arnoldi decomposition becomes what is called Lanczos.

### 13.1 Lanczos iteration and Lanczos decomposition

When $A$ is symmetric, the Arnoldi decomposition simplifies to

$$
A Q_{k}=Q_{k} T_{k}+q_{k+1}\left[0, \ldots, 0, t_{k+1, k}\right],
$$

where $T_{k}$ is symmetric tridiagonal (proof: just note $H_{k}=Q_{k}^{T} A Q_{k}$ in Arnoldi). This yields the Lanczos decomposition

$$
A Q_{k}=Q_{k+1}\lfloor\tilde{T}_{k}, \quad \tilde{T}_{k}=\underbrace{\left[\begin{array}{cccc}
t_{1,1} & t_{1,2} & & \\
t_{2,1} & t_{2,2} & \ddots & \\
& \ddots & & t_{k-1, k} \\
& & t_{k, k-1} & t_{k, k} \\
& & t_{k+1, k}
\end{array}\right]}_{\mathbb{R}^{(k+1) \times k}}, \quad Q_{k+1}^{T} Q_{k+1}=I_{k+1}
$$

- The vectors $q_{k}$ form a 3-term recurrence $t_{k+1, k} q_{k+1}=\left(A-t_{k, k}\right) q_{k}-t_{k-1, k} q_{k-1}$. Orthogonalisation is necessary only against last two vectors $q_{k}, q_{k-1}$
- Significant speedup over Arnoldi; cost is $k A$-multiplication plus $O(k)$ inner products ( $O(n k)$ ).
- In floating-point arithmetic, sometimes the computed $Q_{k}$ loses orthogonality and reorthogonalisation may be necessary (nonexaminable, see e.g. Demmel [8])


### 13.2 CG algorithm for $A x=b, A \succ 0$

Here we introduce the conjugate gradient (CG) method. CG is not only important in practice but has historical significance in that it was the very first Krylov algorithm to be introduced (and initially made a big hype, but then took a while to be recognised as a competitive method).

First recall that when $A$ is symmetric, Lanczos gives $Q_{k}, T_{k}$ such that $A Q_{k}=Q_{k} T_{k}+$ $q_{k+1}\left[0, \ldots, 0, t_{k+1, k}\right], T_{k}$ : tridiagonal.

The idea of CG is as follows: when $A \succ 0 \mathrm{PD}$, solve $Q_{k}^{T}\left(A Q_{k} y-b\right)=T_{k} y-Q_{k}^{T} b=0$, and $x=Q_{k} y$.
This is known as "Galerkin orthogonality": it imposes that the residual $A x-b$ is orthogonal to $Q_{k}$.

- $T_{k} y=Q_{k}^{T} b$ is a tridiagonal linear system, so it requires only $O(k)$ operations to solve.
- Three-term recurrence reduces cost to $O(k) A$-multiplications, making the orthogonalisation cost almost negligible.
- The CG algorithm serendipitously minimises $A$-norm of error in the Krylov subspace $x_{k}=\operatorname{argmin}_{x \in Q_{k}}\left\|x-x_{*}\right\|_{A}\left(A x_{*}=b\right):$ writing $x_{k}=Q_{k} y$, we have

$$
\begin{aligned}
\left(x_{k}-x_{*}\right)^{T} A\left(x_{k}-x_{*}\right) & =\left(Q_{k} y-x_{*}\right)^{T} A\left(Q_{k} y-x_{*}\right) \\
& =y^{T}\left(Q_{k}^{T} A Q_{k}\right) y-2 b^{T} Q_{k} y+b^{T} x_{*},
\end{aligned}
$$

and minimiser wrt $y$ (either by completing the squares, or via convexity of the function wrt $y)$ is $y=\left(Q_{k}^{T} A Q_{k}\right)^{-1} Q_{k}^{T} b$, so $Q_{k}^{T}\left(A Q_{k} y-b\right)=0$.

- Note $\|x\|_{A}=\sqrt{x^{T} A x}$ defines a norm (exercise)
- More generally, for inner-product norm $\|z\|_{M}=\sqrt{\langle z, z\rangle_{M}}, \min _{x=Q y}\left\|x_{*}-x\right\|_{M}$ attained when $\left\langle q_{i}, x_{*}-x\right\rangle_{M}=0, \forall q_{i}$ (cf. Part A Numerical Analysis).

We've described the CG algorithm conceptually. To derive the practical algorithm some clever manipulations are necessary. We won't go over them in detail but here is the outcome:

Set $x_{0}=0, r_{0}=-b, p_{0}=r_{0}$ and do for $k=1,2,3, \ldots$

$$
\begin{aligned}
& \alpha_{k}=\left\langle r_{k}, r_{k}\right\rangle /\left\langle p_{k}, A p_{k}\right\rangle \\
& x_{k+1}=x_{k}+\alpha_{k} p_{k} \\
& r_{k+1}=r_{k}-\alpha_{k} A p_{k} \\
& \beta_{k}=\left\langle r_{k+1}, r_{k+1}\right\rangle /\left\langle r_{k}, r_{k}\right\rangle \\
& p_{k+1}=r_{k+1}+\beta_{k} p_{k}
\end{aligned}
$$

where $r_{k}=b-A x_{k}$ (residual) and $p_{k}$ (search direction). $x_{k}$ is the CG solution after $k$ iterations.

One can show among others (exercise/sheet)

- $\mathcal{K}_{k}(A, b)=\operatorname{span}\left(r_{0}, r_{1}, \ldots, r_{k-1}\right)=\operatorname{span}\left(x_{1}, x_{2}, \ldots, x_{k}\right)\left(\right.$ also equal to $\left.\operatorname{span}\left(p_{0}, p_{1}, \ldots, p_{k-1}\right)\right)$
- $r_{j}^{T} r_{k}=0, j=0,1,2, \ldots, k-1$

Thus $x_{k}$ is $k$ th CG solution, satisfying Galerkin orthogonality $Q_{k}^{T}\left(A x_{k}-b\right)=0$ : residual is orthogonal to the (Krylov) subspace.

### 13.3 CG convergence

Let's examine the convergence of the CG iterates.
Theorem 13.1 Let $A \succ 0$ be an $n \times n$ positive definite matrix and $b \in \mathbb{R}^{n}$. Let $e_{k}:=x_{*}-x_{k}$ be the error after the $k$ th $C G$ iteration ( $x_{*}$ is the exact solution $A x_{*}=b$ ). Then

$$
\frac{\left\|e_{k}\right\|_{A}}{\left\|e_{0}\right\|_{A}} \leq 2\left(\frac{\sqrt{\kappa_{2}(A)}-1}{\sqrt{\kappa_{2}(A)}+1}\right)^{k}
$$

Proof: We have $e_{0}=x_{*}\left(x_{0}=0\right)$, and

$$
\begin{aligned}
\frac{\left\|e_{k}\right\|_{A}}{\left\|e_{0}\right\|_{A}} & =\min _{x \in \mathcal{K}_{k}(A, b)}\left\|x_{k}-x_{*}\right\|_{A} /\left\|x_{*}\right\|_{A} \\
& =\min _{p_{k-1} \in \mathcal{P}_{k-1}}\left\|p_{k-1}(A) b-A^{-1} b\right\|_{A} /\left\|e_{0}\right\|_{A} \\
& =\min _{p_{k-1} \in \mathcal{P}_{k-1}}\left\|\left(p_{k-1}(A) A-I\right) e_{0}\right\|_{A} /\left\|e_{0}\right\|_{A} \\
& =\min _{p \in \mathcal{P}_{k}, p(0)=1}\left\|p(A) e_{0}\right\|_{A} /\left\|e_{0}\right\|_{A} \\
& =\min _{p \in \mathcal{P}_{k}, p(0)=1}\left\|V\left[\begin{array}{rrr}
p\left(\lambda_{1}\right) & & \\
& \ddots & \\
& & p\left(\lambda_{n}\right)
\end{array}\right] V^{T} e_{0}\right\|_{A} /\left\|e_{0}\right\|_{A} .
\end{aligned}
$$

Now (blue) ${ }^{2}=\sum_{i} \lambda_{i} p\left(\lambda_{i}\right)^{2}\left(V^{T} e_{0}\right)_{i}^{2} \leq \max _{j} p\left(\lambda_{j}\right)^{2} \sum_{i} \lambda_{i}\left(V^{T} e_{0}\right)_{i}^{2}=\max _{j} p\left(\lambda_{j}\right)^{2}\left\|e_{0}\right\|_{A}^{2}$.
We've shown

$$
\frac{\left\|e_{k}\right\|_{A}}{\left\|e_{0}\right\|_{A}} \leq \min _{p \in \mathcal{P}_{k}, p(0)=1} \max _{j}\left|p\left(\lambda_{j}\right)\right| \leq \min _{p \in \mathcal{P}_{k}, p(0)=1} \max _{x \in\left[\lambda_{\min }(A), \lambda_{\max }(A)\right]}|p(x)|
$$

To complete the proof, in the next subsection we will show that

$$
\begin{equation*}
\min _{p \in \mathcal{P}_{k}, p(0)=1} \max _{x \in\left[\lambda_{\min }(A), \lambda_{\max }(A)\right]}|p(x)| \leq 2\left(\frac{\sqrt{\kappa_{2}(A)}-1}{\sqrt{\kappa_{2}(A)}+1}\right)^{k} \tag{12}
\end{equation*}
$$

- Note that $\kappa_{2}(A)=\frac{\sigma_{\max }(A)}{\sigma_{\min }(A)}=\frac{\lambda_{\max }(A)}{\lambda_{\min }(A)}\left(=: \frac{b}{a}\right)$.
- The above bound is obtained using Chebyshev polynomials on $\left[\lambda_{\min }(A), \lambda_{\max }(A)\right]$. This is a class of polynomials that arise in a varieties of contexts in computational maths. Let's next look at their properties.


### 13.3.1 Chebyshev polynomials

For $z=\exp (i \theta), x=\frac{1}{2}\left(z+z^{-1}\right)=\cos \theta \in[-1,1], \theta=\operatorname{acos}(x), T_{k}(x)=\frac{1}{2}\left(z^{k}+z^{-k}\right)=\cos (k \theta)$. $T_{k}(x)$ is a polynomial in $x$ :

$$
\frac{1}{2}\left(z+z^{-1}\right)\left(z^{k}+z^{-k}\right)=\frac{1}{2}\left(z^{k+1}+z^{-(k+1)}\right)+\frac{1}{2}\left(z^{k-1}+z^{-(k-1)}\right) \Leftrightarrow \underbrace{2 x T_{k}(x)=T_{k+1}(x)+T_{k-1}(x)}_{3 \text {-term recurrence }}
$$

One can argue why Chebyshev polynomials grow rapidly outside $[-1,1]$ as follows: $T_{k}(x)=2^{k} x^{k}+$ (poly degree $\leq k-1$ ), and so $\left|T_{k}(x)\right|=O\left(2^{k} x^{k}\right)$. We'll derive a more precise estimate (and actual lower bound) below.


These polynomials grow very fast outside the interval (here the 'standard' $[-1,1]$ ). For example, plots on $[-2,1]$ look like



Here's a nice plot of several Chebyshev polynomials:


### 13.3.2 Properties of Chebyshev polynomials

For $z=\exp (i \theta), x=\frac{1}{2}\left(z+z^{-1}\right)=\cos \theta \in[-1,1], \theta=\operatorname{acos}(x), T_{k}(x)=\frac{1}{2}\left(z^{k}+z^{-k}\right)=\cos (k \theta)$.

- Inside $[-1,1],\left|T_{k}(x)\right| \leq 1$
- Outside $[-1,1],\left|T_{k}(x)\right| \gg 1$ grows rapidly with $|x|$ and $k$ (fastest growth among $p \in \mathcal{P}_{k}$ with $|p(x)| \leq 1$ on $x \in[-1,1])$
- Symmetry $\left|T_{k}(x)\right|=\left|T_{k}(-x)\right|$ for all $x \in \mathbb{R}$

Shift+scale s.t. $p(x)=c_{k} T_{k}\left(\frac{2 x-b-a}{b-a}\right)$ where $c_{k}=1 / T_{k}\left(\frac{-(b+a)}{b-a}\right)$ so $p(0)=1$. Then

- $|p(x)| \leq 1 /\left|T_{k}\left(\frac{b+a}{b-a}\right)\right|$ on $x \in[a, b]$
- $T_{k}(z)=\frac{1}{2}\left(z^{k}+z^{-k}\right)$ with $\frac{1}{2}\left(z+z^{-1}\right)=\frac{b+a}{b-a} \Rightarrow z=\frac{\sqrt{b / a}+1}{\sqrt{b / a}-1}=\frac{\sqrt{\kappa_{2}(A)}+1}{\sqrt{\kappa_{2}(A)}-1}$, so

$$
|p(x)| \leq 1 / T_{k}\left(\frac{b+a}{b-a}\right) \leq 2\left(\frac{\sqrt{\kappa}-1}{\sqrt{\kappa}+1}\right)^{k}
$$

This establishes (12), completing the proof of Theorem 13.1.
For much more about $T_{k}$, see C6.3 Approximation of Functions

### 13.4 MINRES: symmetric (indefinite) version of GMRES (nonexaminable)

The content and analysis on MINRES is nonexaminable, although it is really a straightforward mix of GMRES (minimise residual) and CG (simplifications from symmetry). When the matrix is symmetric but not positive definite, GMRES can still be simplified although some care is needed as the Matrix $A$ ceases to define a norm unlike for positive definite matrices.

Symmetric analogue of GMRES: MINRES (minimum-residual method) for $A=A^{T}$ (but not necessarily $A \succ 0$ )

$$
x=\operatorname{argmin}_{x \in \mathcal{K}_{k}(A, b)}\|A x-b\|_{2} .
$$

Algorithm: Given $A Q_{k}=Q_{k+1} \tilde{T}_{k}$ and writing $x=Q_{k} y$, rewrite as

$$
\begin{aligned}
\min _{y}\left\|A Q_{k} y-b\right\|_{2} & =\min _{y}\left\|Q_{k+1} \tilde{T}_{k} y-b\right\|_{2} \\
& =\min _{y}\left\|\left[\begin{array}{c}
\tilde{T}_{k} \\
0
\end{array}\right] y-\left[\begin{array}{c}
Q_{k}^{T} \\
Q_{k, \perp}^{T}
\end{array}\right] b\right\|_{2} \\
& =\min _{y}\left\|\left[\begin{array}{c}
\tilde{T}_{k} \\
0
\end{array}\right] y-\right\| b\left\|_{2} e_{1}\right\|_{2}, \quad e_{1}=[1,0, \ldots, 0]^{T} \in \mathbb{R}^{n}
\end{aligned}
$$

( where $\left[Q_{k}, Q_{k, \perp}\right]$ orthogonal; same trick as in least-squares)

- Minimised when $\left\|\tilde{T}_{k} y-\tilde{Q}_{k}^{T} b\right\| \rightarrow$ min; tridiagonal least-squares problem
- Solve via QR ( $k$ Givens rotations)+ tridiagonal solve, $O(k)$ in addition to Lanczos


### 13.4.1 MINRES convergence

As in GMRES, we can examine the MINRES residual in terms of minimising the values of a polynomial at the eigenvalues of $A$.

$$
\begin{aligned}
\min _{x \in \mathcal{K}_{k}(A, b)}\|A x-b\|_{2} & =\min _{p_{k-1} \in \mathcal{P}_{k-1}}\left\|A p_{k-1}(A) b-b\right\|_{2}=\min _{\tilde{p} \in \mathcal{P}_{k}, \tilde{\mathcal{p}}(0)=0}\|(\tilde{p}(A)-I) b\|_{2} \\
& =\min _{p \in \mathcal{P}_{k}, p(0)=1}\|p(A) b\|_{2} .
\end{aligned}
$$

Since $A=A^{T}, A$ is diagonalisable $A=Q \Lambda Q^{T}$ with $Q$ orthogonal, so

$$
\begin{aligned}
\|p(A)\|_{2} & =\left\|Q p(\Lambda) Q^{T}\right\|_{2} \leq\|Q\|_{2}\left\|Q^{T}\right\|_{2}\|p(\Lambda)\|_{2} \\
& =\max _{z \in \lambda(A)}|p(z)| .
\end{aligned}
$$

Interpretation: (again) find polynomial s.t. $p(0)=1$ and $\left|p\left(\lambda_{i}\right)\right|$ small

$$
\frac{\|A x-b\|_{2}}{\|b\|_{2}} \leq \min _{p \in \mathcal{P}_{k}, p(0)=1} \max \left|p\left(\lambda_{i}\right)\right| .
$$

One can prove (nonexaminable)

$$
\min _{p \in \mathcal{P}_{k}, p(0)=1} \max \left|p\left(\lambda_{i}\right)\right| \leq 2\left(\frac{\kappa_{2}(A)-1}{\kappa_{2}(A)+1}\right)^{k / 2}
$$

- obtained by Chebyshev + Möbius change of variables [Greenbaum's book 97]
- minimisation needed on positive and negative sides, hence slower convergence when $A$ is indefinite

CG and MINRES, optimal polynomials Here are some optimal polynomials for CG and MINRES. Note that $\kappa_{2}(A)=10$ in both cases; observe how much faster CG is than MINRES.





### 13.5 Preconditioned CG/MINRES

The preceding analysis suggests that the linear system

$$
A x=b, \quad A=A^{T}(\succ 0)
$$

may not converge rapidly with CG or MINRES if the eigenvalues are not distributed in a favorable manner (i.e., clustered away from 0).

In this case, as with GMRES, a workaround is to find a good preconditioner $M$ such that " $M^{T} M \approx A^{-1}$ " and solve

$$
M^{T} A M y=M^{T} b, \quad M y=x
$$

As before, desiderata of $M$ :

- $M^{T} A M$ is easy to multiply to a vector.
- $M^{T} A M$ has clustered eigenvalues away from 0 .

Note that reducing $\kappa_{2}\left(M^{T} A M\right)$ directly implies rapid convergence.

- It is possible to implement preconditioned CG with just $M^{T} M$ (no need to find $M$ ).


### 13.6 The Lanczos algorithm for symmetric eigenproblem (nonexaminable)

So far we've discussed Krylov subspace methods for linear systems $A x=b$; we next consider its use for solving eigenvalue problems $A x=\lambda x$. The idea is broadly the same: build the Krylov subspace using Arnoldi or Lanczos, and use it to find a solution (i.e., eigenvector) in the subspace ${ }^{21}$.

We are now ready to describe one of the most successful algorithms for large-scale symmetric eigenvalue problems: the Lanczos algorithm. In simple words, it finds an eigenvalue and eigenvector in a Krylov subspace by a projection method, called the Rayleigh-Ritz process.

```
\(\overline{\text { Algorithm 13.1 Rayleigh-Ritz: given symmetric } A \text { and orthonormal } Q \text {, find approximate }}\)
eigenpairs
    1: Compute \(Q^{T} A Q\).
    2: Eigenvalue decomposition \(Q^{T} A Q=V \hat{\Lambda} V^{T}\).
    3: Approximate eigenvalues \(\operatorname{diag}(\hat{\Lambda})\) (Ritz values) and eigenvectors \(Q V\) (Ritz vectors).
```

This is a projection method (similar alg is available for SVD).
Now we can describe the Lanczos algorithm as follows:

[^19]Algorithm 13.2 Lanczos algorithm: $A \in \mathbb{R}^{n \times n}$ find (extremal) eigenpairs.
1: Perform Lanczos iterations to obtain $A Q_{k}=Q_{k} T_{k}+q_{k+1}\left[0, \ldots, 0, h_{k+1, k}\right]$.
2: Compute the eigenvalue decomposition of $T_{k}=V_{k} \hat{\Lambda} V_{k}^{T}$ (Rayleigh-Ritz with subspace $Q_{k}$ ).
3: $\operatorname{diag}(\hat{\Lambda})$ are the approximate eigenvalues (Ritz values), and the columns of $Q_{k} V_{k}$ are the approximate eigenvectors (Ritz vectors).

- In this case $Q=Q_{k}$, so simply $Q_{k}^{T} A Q_{k}=T_{k}$ (tridiagonal eigenproblem)
- Very good convergence is usually observed to extremal eigenpairs. To see this:
- Recall from Courant-Fischer $\lambda_{\max }(A)=\max _{x} \frac{x^{T} A x}{x^{T} x}$
- Hence $\lambda_{\max }(A) \geq \underbrace{\max _{x \in \mathcal{K}_{k}(A, b)} \frac{x^{T} A x}{x^{T} x}}_{\text {Lanczos output }} \geq \underbrace{\frac{v^{T} A v}{v^{T} v}, \quad v=A^{k-1} b}_{\text {power method }}$
- Same for $\lambda_{\text {min }}$, similar for e.g. $\lambda_{2}$

This is admittedly a very crude estimate for the convergence of Lanczos. A thorough analysis turns out to be very complicated; if interested see for example [26].

Experiments with Lanczos Symmetric $A \in \mathbb{R}^{n \times n}, n=100$, Lanczos/power method with random initial vector $b$


Convergence to dominant eigenvalue
Convergence of Ritz values (approximate eigenvalues)
The same principles of projecting the matrix onto a Krylov subspace applies to nonsymmetric eigenvalue problems. Essentially it boils down to finding the eigenvalues of the upper

Hessenberg matrix $H$ arising in the Arnoldi iteration, rather than the tridiagonal matrix as in Lanczos.

### 13.7 Arnoldi for nonsymmetric eigenvalue problems (nonexaminable)

Arnoldi for eigenvalue problems: the algorithm, just like Lanczos algorithm, combines Arnoldi iteration+Rayleigh-Ritz.

1. Compute $Q^{T} A Q$.
2. Eigenvalue decomposition $Q^{T} A Q=X \hat{\Lambda} X^{-1}$.
3. Approximate eigenvalues $\operatorname{diag}(\hat{\Lambda})$. (Ritz values) and eigenvectors $Q X$ (Ritz vectors).

As in Lanczos, $Q=Q_{k}=\mathcal{K}_{k}(A, b)$, so simply $Q_{k}^{T} A Q_{k}=H_{k}$ (Hessenberg eigenproblem, note that this is ideal for the QR algorithm as the preprocessing step can be skipped).

Which eigenvalues are found by Arnoldi? We give a rather qualitative answer:

- First note that Krylov subspace is invariant under shift: $\mathcal{K}_{k}(A, b)=\mathcal{K}_{k}(A-s I, b)$.
- Thus any eigenvector that power method applied to $A-s I$ converges to should be contained in $\mathcal{K}_{k}(A, b)$.
- To find other (e.g. interior) eigvals, one can use shift-invert Arnoldi: $Q=\mathcal{K}_{k}((A-$ $\left.s I)^{-1}, b\right)$.


## 14 Randomised algorithms in NLA

In this final part of this lecture series we will talk about randomised algorithms. This takes us to the forefront of research in the field: A major idea in numerical linear algebra since 2005 or so has been the use of randomisation where in a matrix sketch is used in order to extract information about the matrix, for example in order to construct a low-rank approximation.

So far, all algorithms have been deterministic (always same output).

- Direct methods (LU for $A x=b, \mathrm{QR}$ alg for $A x=\lambda x$ or $A=U \Sigma V^{T}$ ) are
- Incredibly reliable, backward stable.
- Works like magic if $n \lesssim 10000$.
- But not beyond; cubic complexity $O\left(n^{3}\right)$ or $O\left(m n^{2}\right)$.
- Iterative methods (GMRES, CG, Arnoldi, Lanczos) are
- Very fast when it works (nice spectrum etc).
- Otherwise, not so much; need for preconditioning.
- Randomised algorithms
- Output differs at every run.
- Ideally succeed with enormous probability, e.g. $1-\exp (-c n)$.
- Often by far the fastest\&only feasible approach.
- Not for all problems - active field of research.

Why do we need randomisation? Randomised algorithms are somehow attached to a negative connotation that the algorithm is not reliable and always produces different results. Well, this is a valid concern. We hope to remove such concerns in what follows. The reason randomisation is necessary is the sheer size of datasets that we face today, with the rise of data science. As you will see, one core idea of randomisation is that by allowing for a very small error (which can easily be in the order of machine precision) instead of getting an exact solution, one can often dramatically speed up the algorithm.

We'll cover two NLA topics where randomisation has been very successful: low-rank approximation (randomised SVD), and overdetermined least-squares problems

### 14.1 Gaussian matrices

In randomised algorithms we typically introduce a random matrix. For the analysis, Gaussian matrices $G \in \mathbb{R}^{m \times n}$ are the most convenient (not always for computation). These are matrices whose entries are drawn independently from the standard normal (Gaussian) distribution $G_{i j} \sim N(0,1)$. We cannot do justice to random matrix theory (there is a course in Hilary term!), but here is a summary of the properties of Gaussian matrices that we'll be using, namely orthogonal invariance and the Marchenko-Pastur rule, which informally states "rectangular random matrices are well-conditioned".

### 14.1.1 Orthogonal invariance

A useful fact about Gaussian random matrices $G$ is that its distribution is invariant under orthogonal transformations. That is, if $G$ is Gaussian, so is $Q G$ and $G Q$, where $Q$ is any orthogonal matrix independent of $G$. To see this (nonexaminable): note that a sum of Gaussian (scalar) random variables is Gaussian, with 0 mean. Recall (non-examinable, but worth knowing if you don't) that the distribution of (multivariate) Gaussian random variables is determined completely by the mean and the (co-)variance. Now let $g_{i}$ denote the $i$ th column of $G$. Then the mean is $\mathbb{E}\left[\left(Q g_{i}\right)\right]=Q \mathbb{E}\left[g_{i}\right]=0$. Hence the covariance matrix is $\mathbb{E}\left[\left(Q g_{i}\right)^{T}\left(Q g_{i}\right)\right]=Q \mathbb{E}\left[g_{i}^{T} g_{i}\right] Q^{T}=Q Q^{T}=I$, so each $Q g_{i}$ is multivariate Gaussian with the same distribution $N(0,1)$ as $g_{i}$. It follows that $Q g_{i}$ takes iid $N(0,1)$ entries. Independence of $Q g_{i}, Q g_{j}$ for $i \neq j$ is immediate.

Another (perhaps quicker) way to see this is to note that the joint pdf of $g_{i}=\left[g_{11}, \ldots, g_{n 1}\right]^{T}$ is $\frac{1}{(2 \pi)^{n / 2}} \exp \left(-\frac{1}{2}\left(g_{11}^{2}+\cdots+g_{n 1}^{2}\right)\right)$, and that of $Q g_{i}=\left[\tilde{g}_{11}, \ldots, \tilde{g}_{n 1}\right]^{T}$ is (by a change of variables, noting $\operatorname{det} Q=1)$ is also $\frac{1}{(2 \pi)^{n / 2}} \exp \left(-\frac{1}{2}\left(\tilde{g}_{11}^{2}+\cdots+\tilde{g}_{n 1}^{2}\right)\right)$.

### 14.1.2 Marchenko-Pastur: Rectangular random matrices are well conditioned

Another very useful fact is that when the matrix is rectangular $m>n$ (or $m<n$ ), the singular values of a Gaussian $G$ where $G_{i j} \sim N(0,1)$ are known to lie in the interval $[\sqrt{m}-$ $\sqrt{n}, \sqrt{m}+\sqrt{n}]$. This is a consequence of the Marchenko-Pastur (M-P) rule.

M-P is a classical result in random matrix theory, which we will not be able to prove here (and hence is clearly nonexaminable). We refer those interested to the part C course on Random Matrix Theory. However, understanding the statement and the ability to use this fact is indeed examinable.

The key message is easy to state: a rectangular random matrix is well conditioned with extremely high probability. This fact is enormously important and useful in a variety of contexts in computational mathematics.

Here is a more precise statement:
Theorem 14.1 (Marchenko-Pastur) The singular values of random matrix $X \in \mathbb{R}^{m \times n}$ ( $m \geq n$ ) with iid $X_{i j}$ (mean 0, variance 1) follow the Marchenko-Pastur (M-P) distribution (proof nonexaminable), with density $\sim \frac{1}{x} \sqrt{((\sqrt{m}+\sqrt{n})-x)(x-(\sqrt{m}-\sqrt{n}))}$, and support $[\sqrt{m}-\sqrt{n}, \sqrt{m}+\sqrt{n}]$.

Proof: omitted, and nonexaminable. Strictly speaking, the theorem concerns the limit $m, n \rightarrow \infty$; but the result holds in the nonasymptotic limit with enormous probability. The 'failure' probability that a singular value lies outside the the interval $[\sqrt{m}-\sqrt{n}, \sqrt{m}+\sqrt{n}]$ by $t$ or more is bounded by $\exp \left(-c t^{2}\right)$ for a constant $c>0$; see [6, Thm II.13] if interested. Here and below, this is what we mean when we say informally that an event holds with high/enormous probability.

Note that the statement does not require $X$ to be Gaussian; it holds for a wide variety of random matrices with independent entries (we'll still mostly stick with Gaussian matrices, as the orthogonal invariance is very attractive for the analysis).

Here is an illustration of the theorem.


Histogram of singular values of random Gaussian matrices with varying aspect ratio $m / n$.
$\sigma_{\max }(X) \approx \sqrt{m}+\sqrt{n}, \sigma_{\min }(X) \approx \sqrt{m}-\sqrt{n}$, hence $\kappa_{2}(X) \approx \frac{1+\sqrt{m / n}}{1-\sqrt{m / n}}=O(1)$.
As stated above, M-P is a key fact in many breakthroughs in computational maths! Examples include

- Randomised SVD, Blendenpik (randomised least-squares)
- (nonexaminable:) Compressed sensing (RIP) [Donoho 06, Candes-Tao 06], Matrix concentration inequalities [Tropp 11], Function approximation by least-squares [Cohen-DavenportLeviatan 13]
- (nonexaminable:) You might have heard of the Johnson-Lindenstrauss (JL) Lemma. A host (thousands) of papers have appeared that use JL to prove interesting results. It turns out that many of these can be equally well be proven using M-P.

In the remainder, we will use these facts to develop and understand randomised algorithms for (i) least-squares problems, and (ii) low-rank approximation.

### 14.2 Randomised least-squares

Let's discuss ${ }^{22}$ randomised algorithms for least-squares problems that are highly overdetermined, i.e., $A$ is tall-skinny. Throughout this section, we continue to assume $\operatorname{rank}(A)=n$.


Let $x_{*}$ denote the solution $x_{*}=\operatorname{argmin}_{x}\|A x-b\|_{2}$. First recall that with a traditional method, one either solves the normal equation $x_{*}=\left(A^{T} A\right)^{-1} A^{T} b$ or better yet the thin QR $A=Q R$, then $x_{*}=R^{-1}\left(Q^{T} b\right)$. Both require $O\left(m n^{2}\right)$ cost, and the latter is stable.

Here we'll look into randomised algorithms, which can have a lower complexity, which means the algorithm can run much faster than traditional methods when $m, n$ are large.

## 14.3 "Fast" algorithm: row subset selection

To get a sense of why randomisation is needed, let's think of a dirty 'fast' algorithm for highly overdetermined (i.e. $m \gg n, A$ is very tall-skinny) least-squares problems, based on row subset selection.

- Let's first suppose that $A x_{*}=b$, i.e., $\left\|A x_{*}-b\right\|_{2}=0$; such overdetermined systems are called consistent. Then we see that $x_{*}$ can be found by solving a smaller problem: $x_{*}$ is equal to $\hat{x}:=\operatorname{argmin}_{x}\left\|A_{1} x-b_{1}\right\|_{2}\left(\right.$ or $\left.A_{1} \hat{x}=b_{1}\right)$, where $A_{1} \in \mathbb{R}^{s \times n}(s \geq n)$ is any submatrix of $A$ obtained by taking a subset of $s$ rows, and $b_{1} \in \mathbb{R}^{s}$ is the corresponding entries of $b$. Since we still have $A_{1} x_{*}=b_{1}$, we find the solution $x_{*}$ as long as this system determines $x_{*}$ uniquely, that is, when $\operatorname{rank}\left(A_{1}\right)=n$. This suggests that one can try take any $s$ of the rows, for example simply the first $s=n$ rows. If they are linearly independent, we get the solution in $O\left(s n^{2}\right)$ operations, without even looking at the whole matrix!

[^20]- The above process can easily fail when $A x_{*} \neq b$. Consider, for example, the case where $A=\left[\begin{array}{c}A_{1} \\ A_{2} \\ \vdots \\ A_{k}\end{array}\right], b=\left[\begin{array}{c}b_{1} \\ b_{2} \\ \vdots \\ b_{k}\end{array}\right]$ where $A_{1}=\epsilon I_{n}(\epsilon \ll 1)$, and $A_{i}=I_{n}$ for $i \geq 2$, and $b_{i}=b_{j}$ if $i, j \geq 2$. Then the least-squares problem has solution $x \approx b_{2}$ (convince yourself). However, if the top $s=n$ rows are chosen as the row subset, by solving $\min _{x}\left\|A_{1} x-b_{1}\right\|_{2}$ one gets $\hat{x}=\frac{1}{\epsilon} b_{1}$; this can be arbitrarily bad!

One would like to avoid such pathological cases, irrespective of how 'badly behaved' $A$ and $b$ are. This is where randomisation comes to the rescue.

### 14.4 Sketch-and-solve for least-squares problems

In the last problem above where row subset selection failed, one can see that we would have been much better off had we chosen other rows, say the bottom ones. It turns out that there is always a subset selection that gives a good solution. The trouble is, we don't know which! There are ${ }_{m} C_{s}$ possibilities, and we don't want to try them all.

The goal of randomisation is to mix up the rows so that any subset selection will be good (with enormous probability). The algorithm simply generates a random sketch matrix $G \in \mathbb{R}^{s \times m}$ (which we'll take to be Gaussian for the analysis) with $s=c n(c>1)$, say $c=4$,
 problem

$$
\begin{equation*}
\min _{x}\|G(A x-b)\|_{2} \tag{14}
\end{equation*}
$$

To see a connection with the subset selection algorithm, one can think of first left-multiplying a square Gaussian matrix $\hat{G}:=\left[\begin{array}{c}G \\ G_{2}\end{array}\right]$, consider $\min _{x}\|\hat{G}(A x-b)\|_{2}$, and think of selecting the first $s$ rows, which gives (14).

This is an instance of the fundamental idea behind randomised algorithms, called sketching: multiply a random matrix to reduce the dimension. Here the matrix $G$ is called the sketching matrix, and it is important that the sketch size $s$ is at least as large as the intrinsic dimension of the problem, which here is $n$. That is, we require $s=c n$ where $c>1$, say $s=2 n$.

The algorithm is commonly referred to as a "sketch and solve" approach, as it forms the sketch $G A$ (and $G b$ ), then solves the lower-dimensional problem $\operatorname{argmin}_{x}\|G A x-G b\|_{2}$.

The upshot of sketching is that the solution $\hat{x}=\operatorname{argmin}_{x}\|G(A x-b)\|_{2}$ is a good solution for any problem (regardless of $A, b$ ), with enormous probability.
Theorem 14.2 (Sketch-and-solve for least-squares) Let $A \in \mathbb{R}^{m \times n}$ with $m \gg n$ and $\operatorname{rank}(A)=n$. Let $x_{*}=\operatorname{argmin}_{x}\|A x-b\|_{2}$ and $\hat{x}=\operatorname{argmin}_{x}\|G(A x-b)\|_{2}$, where $G \in \mathbb{R}^{s \times m}$
is a Gaussian matrix with $s>c n, c>1$. Then with high probability

$$
\begin{equation*}
\|A \hat{x}-b\|_{2} \leq \frac{\sqrt{s}+\sqrt{n+1}}{\sqrt{s}-\sqrt{n+1}}\left\|A x_{*}-b\right\|_{2} \tag{15}
\end{equation*}
$$

We can verify this claim using the properties of Gaussian matrices in Section 14.1 as follows.

## Proof:

- Let $[A, b]=Q R$ be the QR factorisation. We have $G[A, b]=(G Q) R$. We can write $\|A v-b\|_{2}=\|Q w\|_{2}$ for some $w \in \mathbb{R}^{n+1}$, and $\|G(A v-b)\|_{2}=\|(G Q) w\|_{2}$.
- The key point is that the $G Q$ matrix is Gaussian, by the orthogonal invariance of Gaussian matrices (Section 14.1.1). Furthermore, $G Q$ is rectangular $s \times(n+1)$, so by M-P, we have $\sigma_{i}(G Q) \in[\sqrt{s}-\sqrt{n+1}, \sqrt{s}+\sqrt{n+1}]$. Roughly, $\|(G Q) w\|_{2} \approx\|Q w\|_{2}$ for all $w$.
- It follows that, using the subordinate property $\|X Y v\|_{2} \leq\|X\|_{2}\|Y v\|_{2}$

$$
\|G(A v-b)\|_{2}=\left\|G[A, b]\left[\begin{array}{c}
v \\
-1
\end{array}\right]\right\|_{2} \leq(\sqrt{s}+\sqrt{n+1})\left\|R\left[\begin{array}{c}
v \\
-1
\end{array}\right]\right\|_{2}=(\sqrt{s}+\sqrt{n+1})\|A v-b\|_{2}
$$

and similarly

$$
\|G(A v-b)\|_{2} \geq(\sqrt{s}-\sqrt{n+1})\|A v-b\|_{2}
$$

These hold for any vector $v$. Since by the definition of $\hat{x}$ we have $\|G(A \hat{x}-b)\|_{2} \leq$ $\|G(A x-b)\|_{2}$, it follows that

$$
\|A \hat{x}-b\|_{2} \leq \frac{1}{\sqrt{s}-\sqrt{n+1}}\|G(A x-b)\|_{2} \leq \frac{\sqrt{s}+\sqrt{n+1}}{\sqrt{s}-\sqrt{n+1}}\|A x-b\|_{2} .
$$

The result shows that by sketch-and-solve, one obtains a solution within $\tau=\frac{\sqrt{s}+\sqrt{n+1}}{\sqrt{s}-\sqrt{n+1}}$ of optimal in terms of minimising the residual. For example, by taking $s=4(n+1)$, we have $\tau=3$. If the residual of the original problem can be made small, say $\left\|A x_{*}-b\right\|_{2}=10^{-10}$, then $\|A \hat{x}-b\|_{2} \leq 3 \times 10^{-10}$, giving an excellent least-squares fit. If $A$ is well-conditioned, this also implies the solution $x$ is close to the exact solution.

Complexity and faster sketches (Nonexaminable) You might have noticed that the sketch-and-solve algorithm isn't faster! The simple-looking computation of $G A$ already requires $O(s m n)$ flops, which is at least $O\left(m n^{2}\right)$, the same as the QR-based algorithm discussed in Section 6, which gives the exact solution. What's the point!?

The answer is that we use a structured (but still random) matrix $G$ such that $G A$ can be computed much more efficiently. For example, $G=S F D$ is a popular choice, where $F$ is an FFT matrix, $D$ is random diagonal (e.g. $\pm 1$ with equal probability) and $S$ is a random
subsampling matrix. Such matrices can be multiplied to $A$ in $O(m n \log s)$ operations (nonexaminable), so the overall cost of solving (13) becomes $O\left(m n \log s+s n^{2}\right)$. In addition, while the theory above assumed $G$ is Gaussian, in practice very similar behavior is often observed with such structured sketch [22].

### 14.5 Sketch-to-precondition: Blendenpik

There is even better news: the sketch $S A$ can be used to find an excellent preconditioner for an iterative algorithm. This way one can obtain essentially the best solution with residual equal to $\left\|A x_{*}-b\right\|_{2}$. The algorithm Blendenpik [1], which takes a "sketch-to-precondition" approach, proceeds as follows. The first step is the same, to sketch the matrix. But the sketch is then used to find a preconditioner, via the QR factorisation.
(QR factorisation), then solve $\min _{y}\left\|\left(A \hat{R}^{-1}\right) y-b\right\|_{2}$ 's normal equation via Krylov (CG) method. The vector $\hat{x}=\hat{R}^{-1} y$ is then the final solution.

- The cost is $O\left(m n \log m+n^{3}\right)$ using fast FFT-type transforms ${ }^{23}$ for $G$.
- Crucially, $A \hat{R}^{-1}$ is well-conditioned. Why? Marchenko-Pastur (next)

The name Blendenpik comes from the fact that the matrix $G$ is 'blending' the rows, and the matrix $G A$ can be seen as picking a submatrix of $\hat{G} A$ as discussed after (14).

### 14.5.1 Explaining $\kappa_{2}\left(A \hat{R}^{-1}\right)=O(1)$ via Marchenko-Pastur

Let us prove that $\kappa_{2}\left(A \hat{R}^{-1}\right)=O(1)$ with high probability. A key result, once again, is M-P.

Theorem 14.3 Let $G \in \mathbb{R}^{s \times m}(s=c n, c>1)$ be a Gaussian matrix and let $\quad G \quad A=$ $\hat{Q} \hat{R}$ (QR factorisation). Then $A \hat{R}^{-1}$ is well-conditioned with high probability.

Let's prove this for $G \in \mathbb{R}^{4 n \times m}$ Gaussian, and for concreteness assuming $s=4 n$ :

[^21]Proof: Let $A=Q R$. Then $G A=(G Q) R=: \tilde{G} R$

- $\tilde{G}$ is $4 n \times n$ rectangular Gaussian, hence well-conditioned.
- So by M-P, $\kappa_{2}\left(\tilde{R}^{-1}\right)=O(1)$ where $\tilde{G}=\tilde{Q} \tilde{R}$ is the QR factorisation.
- Thus $G A=\tilde{G} R=(\tilde{Q} \tilde{R}) R=\tilde{Q}(\tilde{R} R)=\tilde{Q} \hat{R}$, so $\hat{R}^{-1}=R^{-1} \tilde{R}^{-1}$.
- Hence $A \hat{R}^{-1}=(Q R)\left(R^{-1} \hat{R}^{-1}\right)=Q \tilde{R}^{-1}$, so $\kappa_{2}\left(A \hat{R}^{-1}\right)=\kappa_{2}\left(\tilde{R}^{-1}\right)=O(1)$.


### 14.5.2 Blendenpik: solving $\min _{x}\|A x-b\|_{2}$ using $\hat{R}$

We have seen that $\kappa_{2}\left(A \hat{R}^{-1}\right)=: \kappa_{2}(B)=O(1)$; defining $\hat{R} x=y, \min _{x}\|A x-b\|_{2}=$ $\min _{y}\left\|\left(A \hat{R}^{-1}\right) y-b\right\|_{2}=\min _{y}\|B y-b\|_{2}$.

- $B$ is well-conditioned $\Rightarrow$ in normal equation

$$
\begin{equation*}
B^{T} B y=B^{T} b \tag{16}
\end{equation*}
$$

$B^{T} B$ is also well-conditioned $\kappa_{2}\left(B^{T} B\right)=O(1)$; so positive definite and well-conditioned.

- Thus we can solve (16) via CG (or LSQR [27], a more stable variant in this context; nonexaminable)
- The convergence theory for CG implies it would achieve exponential convergence, so convergence in $O(1)$ iterations! (or $O\left(\log \frac{1}{\epsilon}\right)$ iterations for $\epsilon$ accuracy)
- each iteration requires $w \leftarrow B w$ and $w \leftarrow B^{T} w$. It is important to note that these can be computed without computing $A \hat{R}^{-1}$ explicitly. For example, to $w \leftarrow B w$ we first do $w \leftarrow \hat{R}^{-1} w(n \times n$ triangular solve $)$, then $w \leftarrow A w(m \times n$ matrix-vector multiplication). This takes $O(m n)$ cost overall.


### 14.5.3 Blendenpik experiments

Let's illustrate our findings. Since Blendenpik finds a preconditioner $\hat{R}^{-1}$ such that $A \hat{R}^{-1}$ is well-conditioned (regardless of $\kappa_{2}(A)$ ), we expect the convergence of CG to be independent of $\kappa_{2}(A)$. This is indeed what we see here, in Figure 3.

In practice, Blendenpik is observed to achieve $\approx \times 5$ speedup over classical (HouseholderQR based) method when $m \gg n$.


Figure 3: Solving $\min _{x}\|A x-b\|_{2}$ via CG for $A^{T} A x=A^{T} b$ vs. Blendenpik $\left(A R^{-1}\right)^{T}\left(A R^{-1}\right) x=\left(A R^{-1}\right)^{T} b, m=10000, n=100$. The matrix $A$ is generated to have a specified condition number (10 or $10^{4}$ ), then we set a unit norm vector $b=A x+e$ where $\|e\|_{2}=10^{-6}$. Solution with sketch-and-solve (Section 14.4) is also shown, which gets a (suboptimal but good) solution without iterations. The Blendenpik solutions have the same quality as a classical, QR-based method.

## 15 Randomised algorithms for low-rank approximation

The final topic we treat is low-rank approximation, where the goal is to devise fast algorithms for finding an approximation to the truncated SVD. This is arguably the very topic that made randomised algorithms get appreciated by researchers and practicioners as a compelling approach to solve important NLA problems. This course starts with the SVD, and ends with it too.

### 15.1 Randomised SVD by Halko-Martinsson-Tropp

We start with what has been arguably the most successful usage of randomisation in NLA, low-rank approximation. Probably the best reference is the paper by Halko, Martinsson and Tropp (HMT) [16]. This paper has been enormously successful, with almost 4000 Google Scholar citations as of 2022. See also the recent survey by the same authors [22].

The algorithm itself is astonishingly simple:

```
Algorithm 15.1 Randomised SVD (HMT): given \(A \in \mathbb{R}^{m \times n}\) and rank \(r\), find a rank- \(r\)
approximation \(\hat{A} \approx A\).
    Form a random matrix \(G \in \mathbb{R}^{n \times r}\), usually \(r \ll n\).
    2: Compute \(A G\).
    3: Compute the QR factorisation \(A G=Q R\).
    4: \(\left.A \approx Q \quad Q^{T} A \quad\left(Q U_{0}\right) \Sigma_{0} V_{0}^{T}\right)\) is a rank- \(r\) approximation.
```

Here, $G$ is a random matrix taking independent and identically distributed (iid) entries. A convenient choice (for the theory, not necessarily for computation) is a Gaussian matrix, with iid entries $G_{i j} \sim N(0,1)$.

To gain insight, suppose that $\operatorname{rank}(A)=r$, that is, $A=U_{r} \Sigma_{r} V_{r}^{T}$ where $U \in \mathbb{R}^{m \times r}$. Then $A G=U_{r}\left(\Sigma_{r} V_{r}^{T} G\right)$, where $\left(\Sigma_{r} V_{r}^{T} G\right)$ is nonsingular with probability 1 (nonexaminable but I hope this is convincing), and its conditioning improves if $G$ is fatter. Then $Q$ spans exactly $\operatorname{span}\left(U_{r}\right)$, and hence $Q Q^{T} A=U_{r} U_{r}^{T} A=A$. Usually $\operatorname{rank}(A)>r$, but surprisingly, the HMT algorithm still behaves robustly.

Here are some properties of the HMT algorithm:

- $O(m n r)$ cost for dense $A$. This is much lower than the $O\left(m n^{2}\right)$ for the SVD!
- Near-optimal approximation guarantee:

Theorem 15.1 For any $\hat{r}<r$,

$$
\mathbb{E}\|A-\hat{A}\|_{F} \leq \sqrt{1+\frac{r}{r-\hat{r}-1}}\left\|A-A_{\hat{r}}\right\|_{F}
$$

where $A_{\hat{r}}$ is the rank $\hat{r}$-truncated $S V D$ (expectation w.r.t. random matrix $G$ ).
This is a remarkable result; make sure to pause and think about what it says! The approximant $\hat{A}$ has error $\|A-\hat{A}\|_{F}$ that is within a factor $\left(1+\frac{r}{r-\hat{r}-1}\right)$ of the optimal truncated SVD, for a slightly lower rank $\hat{r}<r$ (think, say, $\hat{r}=\lfloor 0.9 r\rfloor$ ).

Goal: understand this, or at least why $\mathbb{E}\|A-\hat{A}\|=O(1)\left\|A-A_{\hat{r}}\right\|$.
Pseudoinverse To understand why the HMT algorithm works, we need to introduce the pseudoinverse of rectangular matrices.

Given $M \in \mathbb{R}^{m \times n}$ with economical SVD $M=U_{r} \Sigma_{r} V_{r}^{T}\left(U_{r} \in \mathbb{R}^{m \times r}, \Sigma_{r} \in \mathbb{R}^{r \times r}, V_{r} \in \mathbb{R}^{n \times r}\right.$ where $r=\operatorname{rank}(M)$ so that $\Sigma_{r} \succ 0$ ), the pseudoinverse $M^{\dagger}$ is

$$
M^{\dagger}=V_{r} \Sigma_{r}^{-1} U_{r}^{T} \in \mathbb{R}^{n \times m}
$$

- $M^{\dagger}$ satisfies $M M^{\dagger} M=M, M^{\dagger} M M^{\dagger}=M^{\dagger}, M M^{\dagger}=\left(M M^{\dagger}\right)^{T}, M^{\dagger} M=\left(M^{\dagger} M\right)^{T}$ (these are often taken to be the definition of the pseudoinverse - the above definition is much simpler IMO).
- $M^{\dagger}=M^{-1}$ if $M$ nonsingular.
- $M^{\dagger} M=I_{n}$ if $m \geq n\left(M M^{\dagger}=I_{m}\right.$ if $\left.m \leq n\right)$ and $M$ is full rank.
- Given a full-rank underdetermined system $\square x=\square$ where $A \in \mathbb{R}^{m \times n}$ with $m<n$ and $r=m$, the general solution is $x=A^{\dagger} b+V_{r, \perp} z$ for arbitrary $z$, and minimum-norm soln is $x=A^{\dagger} b$.

Below we're going to use the pseudoinverse for rectangular matrices of full (column or row) rank.

### 15.2 HMT approximant: analysis (down from 70 pages!)

We are now in a position to explain why the HMT algorithm works. The original HMT paper is over 70 pages with long analysis. Here we attempt to condense the arguments to the essence (and with a different proof). Recall that our low-rank approximation is $\hat{A}=Q Q^{T} A$, where $A G=Q R$. Goal: $\|A-\hat{A}\|=\left\|\left(I_{m}-Q Q^{T}\right) A\right\|=O\left(\left\|A-A_{\hat{r}}\right\|\right)$.

1. $Q Q^{T} A G=A G\left(Q Q^{T}\right.$ is orthogonal projector onto $\left.\operatorname{span}(A G)\right)$. Hence $\left(I_{m}-\right.$ $\left.Q Q^{T}\right) A G=0$, so $A-\hat{A}=\left(I_{m}-Q Q^{T}\right) A\left(I_{n}-G M^{T}\right)$ for any $M \in \mathbb{R}^{n \times r}$.
The idea then is to choose $M$ cleverly so that the expression $\left(I_{m}-Q Q^{T}\right) A\left(I_{n}-G M^{T}\right)$ can be shown to have small norm.
2. Set $M^{T}=\left(V_{1}^{T} G\right)^{\dagger} V_{1}^{T}$ where $V_{1}=\left[v_{1}, \ldots, v_{\hat{r}}\right] \in \mathbb{R}^{n \times \hat{r}}$ is the top right singular vectors of $A(\hat{r} \leq r)$. Recall that $\hat{r}$ is any integer bounded by $r$.
3. $V_{1} V_{1}^{T}\left(I-G M^{T}\right)=V_{1} V_{1}^{T}\left(I-G\left(V_{1}^{T} G\right)^{\dagger} V_{1}^{T}\right)=0$ if $V_{1}^{T} G$ full row-rank (this is a generic assumption that holds with probability 1 if $G$ is Gaussian; note that then $V_{1}^{T} G$ is also Gaussian), so $A-\hat{A}=\left(I_{m}-Q Q^{T}\right) A\left(I-V_{1} V_{1}^{T}\right)\left(I_{n}-G M^{T}\right)$.
4. Taking norms yields $\|A-\hat{A}\|_{2}=\left\|\left(I_{m}-Q Q^{T}\right) A\left(I-V_{1} V_{1}^{T}\right)\left(I_{n}-G M^{T}\right)\right\|_{2}=\|\left(I_{m}-\right.$ $\left.Q Q^{T}\right) U_{2} \Sigma_{2} V_{2}^{T}\left(I_{n}-G M^{T}\right) \|_{2}$ where $\left[V_{1}, V_{2}\right]$ is orthogonal (and $A=\left[U, U_{2}\right]\left[\begin{array}{ll}\Sigma & \\ & \Sigma_{2}\end{array}\right]\left[V_{1}, V_{2}\right]^{T}$ is the SVD), so

$$
\|A-\hat{A}\|_{2} \leq\left\|\Sigma_{2}\right\|_{2}\left\|\left(I_{n}-G M^{T}\right)\right\|_{2}=\underbrace{\left\|\Sigma_{2}\right\|_{2}}_{\text {optimal rank- } \hat{r}}\left\|G M^{T}\right\|_{2}
$$

It remains to prove $\left\|G M^{T}\right\|_{2}=O(1)$. To see why this should hold with high probability, we again invoke the Marchenko-Pastur rule.
$\left\|G M^{T}\right\|_{2}=O(1) \quad$ Recall that we've shown for $M^{T}=\left(V_{1}^{T} G\right)^{\dagger} V_{1}^{T}$ where $G \in \mathbb{R}^{n \times r}$ is random, that

$$
\|A-\hat{A}\|_{2} \leq\left\|\Sigma_{2}\right\|_{2}\left\|\left(I_{n}-G M^{T}\right)\right\|_{2}=\underbrace{\left\|\Sigma_{2}\right\|_{2}}_{\text {optimal rank- } \hat{r}}\left\|G M^{T}\right\|_{2} .
$$

Now $\left\|G M^{T}\right\|_{2}=\left\|G\left(V_{1}^{T} G\right)^{\dagger} V_{1}^{T}\right\|_{2}=\left\|G\left(V_{1}^{T} G\right)^{\dagger}\right\|_{2} \leq\|G\|_{2}\left\|\left(V_{1}^{T} G\right)^{\dagger}\right\|_{2}$.
Now let's analyse the (standard) case where $G$ is random Gaussian $G_{i j} \sim \mathcal{N}(0,1)$. Then

- By the orthogonal invariance of Gaussian matrices (Sec. 14.1.1) $V_{1}^{T} G$ is another Gaussian matrix hence $\left\|\left(V_{1}^{T} G\right)^{\dagger}\right\|=1 / \sigma_{\min }\left(V_{1}^{T} G\right) \lesssim 1 /(\sqrt{r}-\sqrt{\hat{r}})$ by M-P.
- $\|G\|_{2} \lesssim \sqrt{n}+\sqrt{r}$ by M-P.

Together we get $\left\|G M^{T}\right\|_{2} \lesssim \frac{\sqrt{n}+\sqrt{r}}{\sqrt{r}-\sqrt{\hat{r}}}=" O(1)$ ".
Remark:

- When $G$ is a non-Gaussian random matrix, the performance is similar, but is harder to analyze. A popular choice is again the so-called SRFT matrices, which use the FFT (fast Fourier transform) and can be applied to $A$ with $O(m n \log m)$ cost rather than $O(m n r)$.


### 15.3 Precise analysis for HMT (nonexaminable)

A slightly more elaborate analysis again using random matrix theory will give us a very sharp bound on the expected value of the error $E_{\mathrm{HMT}}=: A-\hat{A}$. (this again is non-examinable).

To understand below, we need the notion of projectors. A square matrix $P \in \mathbb{R}^{n \times n}$ is called a projector if $P^{2}=P$.

- $P$ is always diagonalisable and all eigenvalues are 1 or 0 . To see this (nonexaminable): consider the Jordan decomposition of $P$, and see that the eigenvalues 1,0 cannot have a Jordan block of size $>1$.
- $\|P\|_{2} \geq 1$ and $\|P\|_{2}=1$ iff $P=P^{T}$; in this case $P$ is called an orthogonal projector, and $P$ can be written as $P=Q Q^{T}$ where $Q$ is orthonormal.
- One can easily show that $I-P$ is another projector as $(I-P)^{2}=I-P$, and unless $P=0$ or $P=I$, we have $\|I-P\|_{2}=\|P\|_{2}$ :
Schur form $Q P Q^{*}=\left[\begin{array}{cc}I & B \\ 0 & 0\end{array}\right], Q(I-P) Q^{*}=\left[\begin{array}{cc}0 & -B \\ 0 & I\end{array}\right]$; See Szyld $2006[32]$ for many more about projectors.

Theorem 15.2 [Reproduces HMT 2011 Thm.10.5] If $G$ is Gaussian, for any $\hat{r}<r, \mathbb{E}\left\|E_{\mathrm{HMT}}\right\|_{F} \leq$ $\sqrt{\mathbb{E}\left\|E_{\mathrm{HMT}}\right\|_{F}^{2}}=\sqrt{1+\frac{r}{r-\hat{r}-1}}\left\|A-A_{\hat{r}}\right\|_{F}$.

Proof: First ineq: Cauchy-Schwarz. Defining $G\left(V_{1}^{T} G\right)^{\dagger} V_{1}^{T}=: \mathcal{P}_{G, V_{1}}$ (this is a projector as $\left.\mathcal{P}_{G, V_{1}}^{2}=\mathcal{P}_{G, V_{1}}\right),\left\|E_{\mathrm{HMT}}\right\|_{F}^{2}$ is

$$
\begin{aligned}
& \left\|A\left(I-V_{1} V_{1}^{T}\right)\left(I-\mathcal{P}_{G, V_{1}}\right)\right\|_{F}^{2}=\left\|A\left(I-V_{1} V_{1}^{T}\right)\right\|_{F}^{2}+\left\|A\left(I-V_{1} V_{1}^{T}\right) \mathcal{P}_{G, V_{1}}\right\|_{F}^{2} \\
& =\left\|\Sigma_{2}\right\|_{F}^{2}+\left\|\Sigma_{2} \mathcal{P}_{G, V_{1}}\right\|_{F}^{2}=\left\|\Sigma_{2}\right\|_{F}^{2}+\left\|\Sigma_{2}\left(V_{2}^{T} G\right)\left(V_{1}^{T} G\right)^{\dagger} V_{1}^{T}\right\|_{F}^{2}
\end{aligned}
$$

Now if $G$ is Gaussian then $V_{2}^{T} G \in \mathbb{R}^{(n-\hat{r}) \times r}$ and $V_{1}^{T} G \in \mathbb{R}^{\hat{r} \times r}$ are independent Gaussian. Hence by [HMT Prop. 10.1] $\mathbb{E}\left\|\Sigma_{2}\left(V_{2}^{T} G\right)\left(V_{1}^{T} G\right)^{\dagger}\right\|_{F}^{2}=\frac{r}{r-\hat{r}-1}\left\|\Sigma_{2}\right\|_{F}^{2}$, so

$$
\mathbb{E}\left\|E_{\mathrm{HMT}}\right\|_{F}^{2}=\left(1+\frac{r}{r-\hat{r}-1}\right)\left\|\Sigma_{2}\right\|_{F}^{2}
$$

Note how remarkable the theorem is - the 'lazily' computed approximant is nearly optimal up to a factor $\sqrt{1+\frac{r}{r-\hat{r}-1}}$ for a near rank $\hat{r}$ (one can take, e.g. $\hat{r}=0.9 r$ ).

### 15.4 Generalised Nyström (nonexaminable)

Let us briefly mention an algorithm for low-rank approximation that is even faster than HMT, especially when $r \gg 1$.

Let $X \in \mathbb{R}^{n \times r}$ be a random matrix, say Gaussian as before; and set another random matrix $Y \in \mathbb{R}^{n \times(r+\ell)}$, and
[Nakatsukasa arXiv 2020 [24]]

$$
\hat{A}=\left(A X\left(Y^{T} A X\right)^{\dagger} Y^{T}\right) A=\mathcal{P}_{A X, Y} A
$$

Then $\hat{A}$ is another rank-r approximation to $A$, and $A-\hat{A}=\left(I-\mathcal{P}_{A X, Y}\right) A=\left(I-\mathcal{P}_{A X, Y}\right) A(I-$ $X M^{T}$ ); choose $M$ s.t. $X M^{T}=X\left(V_{1}^{T} X\right)^{\dagger} V_{1}^{T}=\mathcal{P}_{X, V_{1}}$. Then $\mathcal{P}_{A X, Y}, \mathcal{P}_{X, V_{1}}$ are (nonorthogonal) projectors, and

$$
\begin{aligned}
\|A-\hat{A}\| & =\left\|\left(I-\mathcal{P}_{A X, Y}\right) A\left(I-\mathcal{P}_{X, V_{1}}\right)\right\| \\
& \leq\left\|\left(I-\mathcal{P}_{A X, Y}\right) A\left(I-V_{1} V_{1}^{T}\right)\left(I-\mathcal{P}_{X, V_{1}}\right)\right\| \\
& \leq\left\|A\left(I-V_{1} V_{1}^{T}\right)\left(I-\mathcal{P}_{X, V_{1}}\right)\right\|+\left\|\mathcal{P}_{A X, Y} A\left(I-V_{1} V_{1}^{T}\right)\left(I-\mathcal{P}_{X, V_{1}}\right)\right\| .
\end{aligned}
$$

- Note that the $\left\|A\left(I-V_{1} V_{1}^{T}\right)\left(I-\mathcal{P}_{X, V_{1}}\right)\right\|$ term is the exact same as in the HMT error.
- Extra term $\left\|\mathcal{P}_{A X, Y}\right\|_{2}=O(1)$ as before if $c>1$ in $Y \in \mathbb{R}^{m \times c r}$ (again, by MarchenkoPastur).
- Overall, about $\left(1+\left\|\mathcal{P}_{A X, Y}\right\|_{2}\right) \approx\left(1+\frac{\sqrt{n}+\sqrt{r+\ell}}{\sqrt{r+\ell}-\sqrt{r}}\right)$ times bigger expected error than HMT (this bound can be improved by a more careful analysis, similar to Thm. 15.2), still near-optimal and much faster $O\left(m n \log n+r^{3}\right)$.

Here's some experiments with HMT and GN (generalised Nyström ).
Let $A=U \Sigma V_{1}^{T}$ be a dense $30,000 \times 30,000$ matrix with geometrically decaying singular values $\sigma_{i}$, and choose $U, V_{1}$ to be random (orthogonal factors of a random Gaussian matrix).

We use HMT and GN to find low-rank approximations to $A$ varying the rank $r$. We also compare with MATLAB's SVD, which gives the optimal truncated SVD (up to numerical errors, which are $O\left(10^{-15}\right)$ so invisible). Here the sketch matrices $G, X(, Y)$ are SRFT matrices (as they run faster than Gaussian matrices).



We see that randomised algorithms can outperform the standard SVD significantly in efficiency, while (unless the truncated SVD is really the goal) the slight suboptimality in accuracy is almost always not an issue.

### 15.5 MATLAB code

Implementing the HMT and GN algorithms (at least with Gaussian matrices, which don't always give optimal speed performance) is very easy!

Here's some sample MATLAB code. We can set up a matrix with exponentially decaying singular values (the matrix is constructed by computing $A=U \Sigma V^{T}$, where $U, V$ are orthonormal and $\Sigma$ is a geometric sequence from $10^{-100}$ to 1 ).

```
n = 1000; % size
A = gallery('randsvd',n,1e100); % geometrically decaying singvals
r = 200; % rank
```

Then to do HMT as follows:

```
X = randn(n,r);
AX = A*X;
[Q,R] = qr(AX,0); % QR fact.
At = Q*(Q'*A);
```

which with high probability gives me an excellent approximation

```
norm(At-A,'fro')/norm(A,'fro')
ans = 1.2832e-15
```

And for Generalised Nyström :

```
X = randn(n,r); Y = randn(n,1.5*r);
AX = A*X; YA = Y'*A; YAX = YA*X;
[Q,R] = qr(YAX,0); % stable pseudo-inverse via the QR factorisation
At = (AX/R)*(Q'*YA);
norm(At-A,'fro')/norm(A,'fro')
ans = 2.8138e-15
```

Both algorithms give an excellent low-rank approximation to $A$.

### 15.6 Randomised algorithm for $A x=b, A x=\lambda x$ ?

We have seen that randomization can be very powerful for low-rank approximation and least-squares problems. What about the two core problems in NLA, linear systems $A x=b$ and eigenvalue problems $A x=\lambda x$ ? A recent preprint [25] proposes the use of sketching (as described above) for GMRES (for $A x=b$ ) and Rayleigh-Ritz (for eigenproblems) combined with a generation of a basis for a Krylov subspace that is not orthonormal. The resulting algorithm can be surprisingly efficient, but does not always work. This is certainly not the end of the story.

Randomised algorithms promise to be employed in more and more applications and problems. We will almost surely see more breakthroughs in randomised NLA. Who be the inventors? Would the young and fresh brains like to take a shot?

## 16 Conclusion and discussion

We have tried to give a good overview of the field of numerical linear algebra in this course. However a number of topics have been omitted completely due to lack of time. Here is an incomplete list of other topics.

### 16.1 Important (N)LA topics not treated

These are clearly nonexaminable but definitely worth knowing if you want to get seriously into the field. These will be discussed in lecture if and only if time permits, and in any case only superficially.

- tensors
[Kolda-Bader 2009 [21]]
- FFT (values $\leftrightarrow$ coefficients map for polynomials)
- sparse direct solvers
- multigrid
[e.g. Golub and Van Loan 2012 [14]]
[Duff, Erisman, Reid 2017 [10]]
[e.g. Elman-Silvester-Wathen 2014 [11]]
- fast (Strassen-type) matrix multiplication etc
[Strassen 1969 [31]+many follow-ups]
- functions of matrices
[Higham 2008 [18]]
- generalised, polynomial/nonlinear eigenvalue problems
[Guttel-Tisseur 2017 [15]]
- perturbation theory (Davis-Kahan [7] etc)
[Stewart-Sun 1990 [30]]
- compressed sensing (this deals with $A x=b$ where $A$ is very 'fat') [Foucart-Rauhut 2013 [12]]
- model order reduction
[Benner-Gugercin-Willcox 2015 [3]]
- communication-avoiding algorithms

$$
\text { [e.g. Ballard-Demmel-Holtz-Schwartz } 2011 \text { [2]] }
$$

### 16.2 Course summary

This information is provided for MSc students who take two separate exams based on the 1 st and 2 nd halfs.

1st half

- SVD and its properties (Courant-Fischer etc), applications (low-rank)
- Direct methods (LU) for linear systems and least-squares problems (QR)
- Stability of algorithms

2nd half

- Direct method (QR algorithm) for eigenvalue problems, SVD
- Krylov subspace methods for linear systems (GMRES, CG) and eigenvalue problems (Arnoldi, Lanczos)
- Randomised algorithms for SVD and least-squares


### 16.3 Related courses you can take

Courses with significant intersection with NLA include

- C6.3 Approximation of Functions: Chebyshev polynomials/approximation theory
- C7.7 Random Matrix Theory: for theoretical underpinnings of Randomised NLA
- C6.4 Finite Element Method for PDEs: NLA arising in solutions of PDEs
- C6.2 Continuous Optimisation: NLA in optimisation problems
and many more: differential equations, data science, optimisation, machine learning,... NLA is everywhere in computational mathematics.

Thank you for your interest in NLA!

## References

[1] H. Avron, P. Maymounkov, and S. Toledo. Blendenpik: Supercharging LAPACK's least-squares solver. SIAM J. Sci. Comp., 32(3):1217-1236, 2010.
[2] G. Ballard, J. Demmel, O. Holtz, and O. Schwartz. Minimizing communication in numerical linear algebra. SIAM J. Matrix Anal. Appl., 32(3):866-901, 2011.
[3] P. Benner, S. Gugercin, and K. Willcox. A survey of projection-based model reduction methods for parametric dynamical systems. SIAM Rev., 57(4):483-531, 2015.
[4] R. Bhatia. Positive Definite Matrices. Princeton University Press, 2009.
[5] K. Braman, R. Byers, and R. Mathias. The multishift QR algorithm. Part II: Aggressive early deflation. SIAM J. Matrix Anal. Appl., 23:948-973, 2002.
[6] K. R. Davidson and S. J. Szarek. Local operator theory, random matrices and banach spaces. Handbook of the geometry of Banach spaces, 1(317-366):131, 2001.
[7] C. Davis and W. M. Kahan. The rotation of eigenvectors by a perturbation. III. SIAM J. Numer. Anal., 7(1):1-46, 1970.
[8] J. Demmel. Applied Numerical Linear Algebra. SIAM, Philadelphia, USA, 1997.
[9] J. Dongarra and F. Sullivan. Guest editors' introduction: The top 10 algorithms. IEEE Computer Architecture Letters, 2(01):22-23, 2000.
[10] I. S. Duff, A. M. Erisman, and J. K. Reid. Direct Methods for Sparse Matrices. Oxford University Press, 2017.
[11] H. C. Elman, D. J. Silvester, and A. J. Wathen. Finite elements and fast iterative solvers: with applications in incompressible fluid dynamics. Oxford University Press, USA, 2014.
[12] S. Foucart and H. Rauhut. A Mathematical Introduction to Compressive Sensing. Springer, 2013.
[13] D. F. Gleich. Pagerank beyond the web. SIAM Rev., 57(3):321-363, 2015.
[14] G. H. Golub and C. F. Van Loan. Matrix Computations. The Johns Hopkins University Press, 4th edition, 2012.
[15] S. Güttel and F. Tisseur. The nonlinear eigenvalue problem. Acta Numer., 26:1-94, 2017.
[16] N. Halko, P.-G. Martinsson, and J. A. Tropp. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. SIAM Rev., 53(2):217-288, 2011.
[17] N. J. Higham. Accuracy and Stability of Numerical Algorithms. SIAM, Philadelphia, PA, USA, second edition, 2002.
[18] N. J. Higham. Functions of Matrices: Theory and Computation. SIAM, Philadelphia, PA, USA, 2008.
[19] R. A. Horn and C. R. Johnson. Topics in Matrix Analysis. Cambridge University Press, 1991.
[20] R. A. Horn and C. R. Johnson. Matrix Analysis. Cambridge University Press, second edition, 2012.
[21] T. G. Kolda and B. W. Bader. Tensor decompositions and applications. SIAM Rev., 51(3):455-500, 2009.
[22] P.-G. Martinsson and J. A. Tropp. Randomized numerical linear algebra: Foundations and algorithms. Acta Numer., pages 403-572, 2020.
[23] M. F. Murphy, G. H. Golub, and A. J. Wathen. A note on preconditioning for indefinite linear systems. SIAM J. Sci. Comp., 21(6):1969-1972, 2000.
[24] Y. Nakatsukasa. Fast and stable randomized low-rank matrix approximation. arXiv:2009.11392.
[25] Y. Nakatsukasa and J. A. Tropp. Fast \& accurate randomized algorithms for linear systems and eigenvalue problems. arXiv 2111.00113.
[26] C. C. Paige. Error analysis of the Lanczos algorithm for tridiagonalizing a symmetric matrix. IMA J. Appl. Math., 18(3):341-349, 1976.
[27] C. C. Paige and M. A. Saunders. LSQR: An algorithm for sparse linear equations and sparse least squares. ACM Trans. Math. Soft., 8(1):43-71, 1982.
[28] B. N. Parlett. The Symmetric Eigenvalue Problem. SIAM, Philadelphia, 1998.
[29] Y. Saad and M. H. Schultz. GMRES - A generalized minimal residual algorithm for solving nonsymmetric linear systems. SIAM J. Sci. Stat. Comp., 7(3):856-869, 1986.
[30] G. W. Stewart and J.-G. Sun. Matrix Perturbation Theory (Computer Science and Scientific Computing). Academic Press, 1990.
[31] V. Strassen. Gaussian elimination is not optimal. Numer. Math., 13:354-356, 1969.
[32] D. B. Szyld. The many proofs of an identity on the norm of oblique projections. Numerical Algorithms, 42(3-4):309-323, 2006.
[33] L. N. Trefethen and D. Bau. Numerical Linear Algebra. SIAM, Philadelphia, 1997.
[34] M. Udell and A. Townsend. Why are big data matrices approximately low rank? SIAM Journal on Mathematics of Data Science, 1(1):144-160, 2019.
[35] D. P. Woodruff. Sketching as a tool for numerical linear algebra. Foundations and Trends $(B)$ in Theoretical Computer Science, 10(1-2):1-157, 2014.


[^0]:    *Last update: May 3, 2023. (minor typos fixed) Please report any corrections or comments on these lecture notes to nakatsukasa@maths.ox.ac.uk

[^1]:    ${ }^{1}$ There are exciting recent developments involving eigenvalue problems for rectangular matrices, but these are outside the scope of this course.

[^2]:    ${ }^{2}$ A pertinent quote is Richard Feynman's "Linear systems are important because we can solve them". Because we can solve them, we do all sorts of tricks to reduce difficult problems to linear systems!

[^3]:    ${ }^{3}$ While there are a small number of cases where the distinction between real and complex matrices matters, in the majority of cases it does not, and the argument carries over to complex matrices by replacing. ${ }^{T}$ with .*. Therefore for the most part, we lose no generality in assuming the matrix is real (which slightly simplifies our mindset). Whenever necessary, we will highlight the subtleties that arise resulting from the difference between real and complex. (For the curious, these are the Schur form/decomposition, $L D L^{T}$ factorisation and eigenvalue decomposition for (real) matrices with complex eigenvalues.)

[^4]:    ${ }^{4}$ Just in case, here's a proof: for any scalar $c,\|x-c y\|^{2}=\|x\|^{2}-2 c x^{T} y+c^{2}\|y\|^{2}$. This is minimised wrt $c$ at $c=\frac{x^{T} y}{\|y\|^{2}}$ with minimiser $\|x\|^{2}-\frac{\left(x^{T} y\right)^{2}}{\|y\|^{2}}$. Since this must be $\geq 0$, the CS inequality follows.

[^5]:    ${ }^{5}$ If this argument isn't convincing to you now, probably the easiest way to see this is via the SVD; so stay tuned and we'll resolve this in footnote 9, once you've seen the SVD!

[^6]:    ${ }^{6}$ I like to think this is the shortest proof out there.

[^7]:    ${ }^{7}$ I presume you've seen the big-Oh notation before $-f(n)=O\left(n^{k}\right)$ means there exists a constant $C$ s.t. $f(n) / n^{k} \leq C$ for all sufficiently large $n$. In NLA it is a convenient way to roughly measure the operation count.
    ${ }^{8}$ If $r \geq \min (m, n)$ then we can simply take $A_{r}=A$.

[^8]:    ${ }^{9}$ Let us now resolve the cliffhanger in footnote 5 . The claim is that any "fat" $m \times n(m<n)$ matrix $M$ has a right null vector $y \neq 0$ such that $M y=0$. To prove this, use the full SVD $M=U \Sigma V^{T}$ to see that $M v_{n}=0$.

[^9]:    ${ }^{10}$ It is somewhat surprising that images are approximable by low-rank matrices. See https://www. youtube.com/watch?v=9BYsNpTCZGg for a nice explanation.

[^10]:    ${ }^{11}$ exact analogues hold for the $i$ th smallest eigenvalue and singular values.
    ${ }^{12}$ Hermann Weyl was one of the prominent mathematicians of the 20 th centry.

[^11]:    ${ }^{13}$ While we won't discuss complete pivoting further, you might be interested to know that when LU with complete pivoting is applied to a low-rank matrix (with rapidly decaying singular values), one tends to find a good low-rank approximation, almost as good as truncated SVD.

[^12]:    ${ }^{14}$ Positive definite matrices are so important a class of matrices; also full of interesting mathematical properties, enough for a nice book to be written about it [4].

[^13]:    ${ }^{15}$ This algorithm does the following: $A^{T} A=R^{T} R$ (Cholesky), then $Q=A R^{-1}$. As stated it's a very fast but unstable algorithm.

[^14]:    ${ }^{16}$ Just to be clear, if one uses a norm that is not unitarily invariant (e.g. 1-norm), the truncated SVD may cease to be a solution for the low-rank approximation problem.

[^15]:    ${ }^{17}$ The definition here follows Higham's book [17]. The phrase is sometimes used to mean small error; However, as hopefully you will be convinced after this section, it is very often impossible to get a solution of full accuracy if the original problem was ill-conditioned. The notion of backward stability that we employ here (and in much of numerical analysis) is therefore much more realistic.

[^16]:    ${ }^{18} \mathrm{~A}$ great application of this is univariate optimisation: to minimise a polynomial $p(x)$, one can find the critical points by solving for $p^{\prime}(x)=0$, via the companion eigenvalues.

[^17]:    ${ }^{19}$ In production code a so-called Wilkinson shift is often used, which computes the eigenvalue of the $2 \times 2$ bottom-right submatrix of $A_{k}$, and takes $s_{k}$ to be the one closer to the bottom-right entry of $A_{k}$.

[^18]:    ${ }^{20}$ Note that the QR algorithm is iterative so it is not exactly correct to call it direct; however the nature of its consistency and robustness together with the fact that the cost is more less predictable has rightfully earned it the classification as a direct method.

[^19]:    ${ }^{21}$ This section was moved to the end of this section in November 2022, hoping it gives a more natural flow; the content is the same.

[^20]:    ${ }^{22}$ This section was brought forward in October 2022; the content is the same (just a bit spruced up).

[^21]:    ${ }^{23}$ The FFT (fast Fourier transform) is one of the important topics that we can't treat properly-for now just think of it as a matrix-vector multiplication that can be performed in $O(n \log n)$ flops rather than $O\left(n^{2}\right)$.)

