

## Sensitivity of Linear Systems

Atkinson §8.4, NMMC<sup>1</sup> §1.2.7, & GvL<sup>2</sup> §2.6

**1** Our first topic is methods for solving nonsingular linear systems of equations. The first question one usually asks about a problem, before analyzing methods to solve it, is whether a solution exists, and whether it's unique. For nonsingular linear systems the answer is yes, so we move on to analyzing how the solution depends on the data, i.e. sensitivity.

We will analyze sensitivity to errors in the inputs: the coefficient matrix and the RHS. These are forward error bounds. Later on we will see that Gaussian Elimination (one particular algorithm for finding the solution) in finite-precision floating point arithmetic finds the exact answer to a perturbed problem, with a bound on the size of the perturbations (a backwards error bound). Those results combined with these will together result in forward error bounds on the solution computed via Gaussian Elimination in finite-precision.

**2** Sensitivity to RHS. (Atkinson §8.4)

$$\mathbf{Ax} = \mathbf{b}$$

where  $\mathbf{A}$  is nonsingular. What happens if you perturb the RHS, eg due to roundoff, and then solve the perturbed system exactly?

$$\mathbf{Ax}_\delta = \mathbf{b} + \delta\mathbf{b}$$

$$\mathbf{x}_\delta = \mathbf{A}^{-1}\mathbf{b} + \mathbf{A}^{-1}(\delta\mathbf{b}) = \mathbf{x} + \mathbf{A}^{-1}(\delta\mathbf{b})$$

$$\|\mathbf{x}_\delta - \mathbf{x}\| = \|\mathbf{A}^{-1}(\delta\mathbf{b})\|.$$

We want the relative error; if the error is  $10^{-30}$  it sounds good unless the norm of the solution is also  $10^{-30}$ . Divide both sides by  $\|\mathbf{x}\|$ :

$$\frac{\|\mathbf{x}_\delta - \mathbf{x}\|}{\|\mathbf{x}\|} = \frac{\|\mathbf{A}\|\|\mathbf{A}^{-1}(\delta\mathbf{b})\|}{\|\mathbf{A}\|\|\mathbf{x}\|} \leq \kappa(\mathbf{A}) \frac{\|\delta\mathbf{b}\|}{\|\mathbf{b}\|}.$$

The condition number is

$$\kappa(\mathbf{A}) = \|\mathbf{A}\|\|\mathbf{A}^{-1}\|.$$

Condition numbers are  $\geq 1$ ; e.g. for the 2-norm the condition number is  $\sigma_{\max}/\sigma_{\min}$  where  $\sigma_i$  are singular values of  $\mathbf{A}$ . An ill-conditioned system is a system with a large condition number. If the system is ill-conditioned, then a small perturbation to the RHS can lead to large changes in the solution. For example,

$$\mathbf{A} = \begin{bmatrix} 1 & 0 \\ 0 & \epsilon \end{bmatrix}, \quad \mathbf{b} = (1, 0)^T, \quad \delta\mathbf{b} = (0, \delta)^T$$

The unperturbed solution is  $(1, 0)^T$ . If  $\delta$  is small then the perturbation is small, but the solution to the perturbed problem is  $(1, \delta/\epsilon)^T$  so if  $\epsilon \ll \delta$  then the error is large.

**3** Perturbing both: Formal asymptotic analysis from GvL §2.6.2.  $\mathbf{A}$  is nonsingular. Consider

$$(\mathbf{A} + \epsilon\mathbf{F})\mathbf{x}(\epsilon) = \mathbf{b} + \epsilon\mathbf{f}.$$

Now differentiate ('implicit differentiation')

$$\mathbf{A}\dot{\mathbf{x}}(\epsilon) + \mathbf{F}\mathbf{x}(\epsilon) + \epsilon\mathbf{F}\dot{\mathbf{x}}(\epsilon) = \mathbf{f}.$$

Now evaluate at  $\epsilon = 0$ , and define the notation  $\mathbf{x}(0) = \mathbf{x}$  = the solution to the unperturbed problem

$$\mathbf{A}\dot{\mathbf{x}} + \mathbf{F}\mathbf{x} = \mathbf{f} \quad \Rightarrow \quad \dot{\mathbf{x}} = \mathbf{A}^{-1}(\mathbf{f} - \mathbf{F}\mathbf{x}).$$

$\mathbf{x}(\epsilon)$  has a Taylor series expansion of the form

$$\mathbf{x}(\epsilon) = \mathbf{x} + \epsilon\dot{\mathbf{x}}(0) + o(\epsilon).$$

---

<sup>1</sup>Numerical Methods in Matrix Computations by Bjork

<sup>2</sup>Matrix Computations by Golub and van Loan

(The ‘little-o’ notation  $y = o(\epsilon)$  means that  $\lim_{\epsilon \rightarrow 0} y/\epsilon = 0$ . The ‘big-o’ notation  $y = \mathcal{O}(\epsilon)$  means that  $\lim_{\epsilon \rightarrow 0} y/\epsilon$  exists and is neither infinite nor zero.)

So we have the following

$$\frac{\|\mathbf{x}(\epsilon) - \mathbf{x}(0)\|}{\|\mathbf{x}(0)\|} = \epsilon \frac{\|\mathbf{A}^{-1}(\mathbf{f} - \mathbf{F}\mathbf{x})\|}{\|\mathbf{x}\|} + o(\epsilon).$$

We’ll use the bound

$$\frac{\|\mathbf{A}^{-1}(\mathbf{f} - \mathbf{F}\mathbf{x})\|}{\|\mathbf{x}\|} \leq \|\mathbf{A}^{-1}\| \left( \frac{\|\mathbf{f}\|}{\|\mathbf{x}\|} + \|\mathbf{F}\| \right).$$

Note that

$$\mathbf{A}\mathbf{x} = \mathbf{b} \quad \Rightarrow \quad \|\mathbf{A}\|\|\mathbf{x}\| \geq \|\mathbf{b}\| \quad \Rightarrow \quad \frac{1}{\|\mathbf{x}\|} \leq \frac{\|\mathbf{A}\|}{\|\mathbf{b}\|}.$$

Using this above we obtain

$$\frac{\|\mathbf{x}(\epsilon) - \mathbf{x}\|}{\|\mathbf{x}\|} \leq \|\mathbf{A}^{-1}\|\|\mathbf{A}\| \left( \frac{\|\epsilon\mathbf{f}\|}{\|\mathbf{b}\|} + \frac{\|\epsilon\mathbf{F}\|}{\|\mathbf{A}\|} \right) + o(\epsilon).$$

This says that the relative error is proportional to the condition number times the relative amplitude of the perturbations in the RHS and in the coefficient matrix. The above derivation (though correct) is not completely rigorous because we never justified that the derivatives exist (they do for small enough  $\epsilon$ ).

The above result is asymptotic, meaning that the error will be approximately equal to the formula on the right hand side *for small enough*  $\epsilon$ . Unfortunately we don’t know how small  $\epsilon$  needs to be in order for the approximation to be accurate. We want some information about the error that will be correct for a known range of  $\epsilon$ .

**4** Theorem 8.4 in Atkinson is not very good. We will instead prove a bound similar to the above rigorously. To prove a rigorous bound we need a Lemma (see also Atkinson Theorems 7.10 & 7.11). Note that  $\|\mathbf{F}\| < 1$  (any operator norm) implies that all eigenvalues are less than 1. If all evals of  $\mathbf{F}$  are less than 1 then  $\mathbf{I} - \mathbf{F}$  is nonsingular (no evals are 0).

Assuming  $\|\mathbf{F}\| < 1$ , note the following

$$\mathbf{I} - \mathbf{F}^{N+1} = (\mathbf{I} + \mathbf{F} + \dots + \mathbf{F}^N) - \mathbf{F}(\mathbf{I} + \mathbf{F} + \dots + \mathbf{F}^N) = \left( \sum_0^N \mathbf{F}^k \right) (\mathbf{I} - \mathbf{F}).$$

Multiply from the right by  $(\mathbf{I} - \mathbf{F})^{-1}$

$$(\mathbf{I} - \mathbf{F}^{N+1})(\mathbf{I} - \mathbf{F})^{-1} = \sum_0^N \mathbf{F}^k$$

Take the limit as  $N \rightarrow \infty$  and recall that  $\|\mathbf{F}\| < 1$  implies that  $\mathbf{F}^N \rightarrow \mathbf{0}$

$$(\mathbf{I} - \mathbf{F})^{-1} = \sum_0^{\infty} \mathbf{F}^k$$

Note similarity to Taylor series for  $1/(1-x)$ .

Return to fixed  $N$ :

$$\|(\mathbf{I} - \mathbf{F}^{N+1})(\mathbf{I} - \mathbf{F})^{-1}\| \leq \sum_0^N \|\mathbf{F}\|^k$$

Now take limits, use geometric series

$$\|(\mathbf{I} - \mathbf{F})^{-1}\| \leq \frac{1}{1 - \|\mathbf{F}\|}.$$

Finally, consider

$$(\mathbf{I} - \mathbf{F})^{-1} - \mathbf{I} = \sum_0^{\infty} \mathbf{F}^k - \mathbf{I} = \sum_1^{\infty} \mathbf{F}^k = \mathbf{F} \sum_0^{\infty} \mathbf{F}^k = \mathbf{F}(\mathbf{I} - \mathbf{F})^{-1}.$$

This implies that

$$\|(\mathbf{I} - \mathbf{F})^{-1} - \mathbf{I}\| = \|\mathbf{F}(\mathbf{I} - \mathbf{F})^{-1}\| \leq \frac{\|\mathbf{F}\|}{1 - \|\mathbf{F}\|}.$$

**5** Now return to a rigorous proof. Consider  $\mathbf{A} + \mathbf{E} = (\mathbf{I} - \mathbf{F})\mathbf{A}$  where  $\mathbf{F} = -\mathbf{E}\mathbf{A}^{-1}$ . Assume that  $\|\mathbf{E}\mathbf{A}^{-1}\| < 1$  so that  $\mathbf{A} + \mathbf{E}$  is nonsingular.

$$\begin{aligned} (\mathbf{A} + \mathbf{E})\mathbf{y} &= \mathbf{b} + \delta\mathbf{b}, \quad \mathbf{A}\mathbf{x} = \mathbf{b} \\ (\mathbf{I} - \mathbf{F})\mathbf{A}\mathbf{y} &= \mathbf{b} + \delta\mathbf{b} \\ \mathbf{y} &= \mathbf{A}^{-1}(\mathbf{I} - \mathbf{F})^{-1}(\mathbf{b} + \delta\mathbf{b}) \\ \mathbf{y} - \mathbf{x} &= \mathbf{A}^{-1} [(\mathbf{I} - \mathbf{F})^{-1} - \mathbf{I}] \mathbf{b} + (\mathbf{I} - \mathbf{F})^{-1} \delta\mathbf{b} \\ \|\mathbf{y} - \mathbf{x}\| &\leq \|\mathbf{A}^{-1}\| \left[ \frac{\|\mathbf{F}\|}{1 - \|\mathbf{F}\|} \|\mathbf{b}\| + \frac{1}{1 - \|\mathbf{F}\|} \|\delta\mathbf{b}\| \right] \end{aligned}$$

Divide by  $\|\mathbf{x}\|$  and multiply RHS by  $\|\mathbf{A}\|/\|\mathbf{A}\|$

$$\begin{aligned} \frac{\|\mathbf{y} - \mathbf{x}\|}{\|\mathbf{x}\|} &\leq \kappa(\mathbf{A}) \left[ \frac{\|\mathbf{F}\|}{1 - \|\mathbf{F}\|} \frac{\|\mathbf{b}\|}{\|\mathbf{A}\|\|\mathbf{x}\|} + \frac{1}{1 - \|\mathbf{F}\|} \frac{\|\delta\mathbf{b}\|}{\|\mathbf{A}\|\|\mathbf{x}\|} \right] \\ \frac{\|\mathbf{y} - \mathbf{x}\|}{\|\mathbf{x}\|} &\leq \kappa(\mathbf{A}) \left[ \frac{\|\mathbf{F}\|}{1 - \|\mathbf{F}\|} + \frac{1}{1 - \|\mathbf{F}\|} \frac{\|\delta\mathbf{b}\|}{\|\mathbf{b}\|} \right]. \end{aligned}$$

Now let  $\epsilon = \max\{\|\mathbf{F}\|, \|\delta\mathbf{b}\|/\|\mathbf{b}\|\} (< 1)$

$$\frac{\|\mathbf{y} - \mathbf{x}\|}{\|\mathbf{x}\|} \leq \frac{2\epsilon\kappa(\mathbf{A})}{1 - \epsilon}.$$

This relies on the fact that  $x/(1-x)$  and  $1/(1-x)$  are increasing on  $x \in (0, 1)$ .

The above result depends on  $\|\mathbf{F}\|$  while the asymptotic result depended on  $\|\mathbf{E}\|/\|\mathbf{A}\|$ . You can make the connection using  $\mathbf{F}\mathbf{A} = -\mathbf{E} \Rightarrow \|\mathbf{F}\|\|\mathbf{A}\| \geq \|\mathbf{E}\| \Rightarrow \|\mathbf{F}\| \geq \|\mathbf{E}\|/\|\mathbf{A}\|$ .

The above result is not asymptotic; it's an upper bound. The relative error might actually be a lot smaller than this.

In summary, if the system is ill-conditioned then small perturbations in the matrix or RHS can lead to large changes in the solution. A large perturbation in the matrix or RHS can lead to large changes in the solution even when the system is well-conditioned.

**6** Preconditioning. Suppose that  $\mathbf{A}$ ,  $\mathbf{M}$ ,  $\mathbf{M}_l$ , and  $\mathbf{M}_r$  are all invertible matrices. The following systems all have the same solution

- Standard:  $\mathbf{A}\mathbf{x} = \mathbf{b}$
- Left-preconditioned  $\mathbf{M}\mathbf{A}\mathbf{x} = \mathbf{M}\mathbf{b}$
- Right-preconditioned  $\mathbf{A}\mathbf{M}\mathbf{y} = \mathbf{b}$  and  $\mathbf{M}\mathbf{y} = \mathbf{x}$
- Split-preconditioned  $\mathbf{M}_l\mathbf{A}\mathbf{M}_r\mathbf{y} = \mathbf{M}_l\mathbf{b}$  and  $\mathbf{M}_r\mathbf{y} = \mathbf{x}$ .

For the left- and right-preconditioned systems, if  $\mathbf{M} \approx \mathbf{A}^{-1}$  then the preconditioned system has the same solution as the original system but a better condition number. For the split-preconditioned system, if  $\mathbf{A} = \mathbf{L}\mathbf{U}$  and  $\mathbf{M}_l \approx \mathbf{L}^{-1}$  and  $\mathbf{M}_r \approx \mathbf{U}^{-1}$  then the preconditioned system has a better condition number. We won't discuss methods for constructing preconditioners; take the numerical linear algebra course APPM 5620 offered every spring.

---

**Gaussian Elimination & LU Factorization**  
Atkinson §8.1–8.4, NMMC §1.2, & GvL Chapter 3

**1** Gaussian Elimination with Partial Pivoting (GEPP) is an algorithm that is guaranteed to compute the exact solution (in the absence of roundoff errors) to a nonsingular linear system in a finite number of steps. Any algorithm that computes an exact solution in a finite number of steps is a ‘direct method.’

Form the augmented matrix  $[\mathbf{A}|\mathbf{b}]$ . Then:

For  $k = 1 : n - 1$

- Swap row  $k$  with the row that has the largest element in the  $k^{\text{th}}$  column on or below the diagonal.
- Add multiples of row  $k$  to rows  $k + 1 : n$  to set all elements below the  $k^{\text{th}}$  diagonal to 0.

The augmented matrix will now have the form  $[\mathbf{U}|\mathbf{c}]$ . Set  $x_n = c_n/u_{n,n}$ . For  $k = n - 1 : -1 : 1$

- Set  $x_k = (c_k - \sum_{i=k+1}^n u_{k,i}x_i)/u_{k,k}$ .

The first loop is the forward solve, the second loop is the back solve. In the forward solve, if there’s a zero on the diagonal you have to pivot (pivot means swap rows), and you can pivot in any way that produces a nonzero on the diagonal. Now in finite-precision arithmetic the probability of getting a hard zero is small anyways, so that’s not really why we pivot. Instead, the strategy that you use for pivoting has an impact on the magnitude of the roundoff error; we pivot to minimize roundoff errors. I have specified a particular pivoting strategy: pivoting the largest element to the diagonal. We will later quote an error bound that depends on this particular strategy. An even better strategy from the perspective of roundoff errors is to pivot the rows and columns so that the largest element in the trailing submatrix gets moved to the diagonal. This is called ‘complete’ pivoting and leads to a better roundoff error bound, but is also more costly.

‘Gauss-Jordan’ elimination is a variant where the second loop is like the first: you eliminate elements *above* the diagonal (without pivoting). After the second loop the augmented matrix has the form  $[\mathbf{I}|\mathbf{x}]$ . Gauss-Jordan is more expensive than Gaussian Elimination.

**2** Gaussian Elimination & LU Factorization. If you solve  $\mathbf{Ax} = \mathbf{b}$  multiple times with different right hand sides, the ‘forward solve’ part of the loop is the same every time. You will always arrive at an augmented matrix of the form  $[\mathbf{U}|\mathbf{c}]$ ; the upper triangular matrix  $\mathbf{U}$  will always be the same, only the vector  $\mathbf{c}$  will depend on the right hand side  $\mathbf{b}$ . So it’s wasteful to re-do the whole forward loop each time. Instead, you should save some record of the operations (pivoting, multiplying and adding rows) that occur during the forward loop. If you have to solve the system again with a different  $\mathbf{b}$ , you just re-apply the same sequence of operations to the new RHS vector  $\mathbf{b}$  to get the new  $\mathbf{c}$ , then proceed to the back solve.

It turns out that storing the steps in the forward loop of Gaussian Elimination is like computing a matrix factorization (called a permuted LU factorization) of the form

$$\mathbf{PA} = \mathbf{LU}$$

where  $\mathbf{P}$  is a permutation matrix,  $\mathbf{L}$  is lower-triangular with 1 on the diagonal (‘special lower triangular’) and  $\mathbf{U}$  is upper-triangular. Permuting rows is equivalent to multiplying from the left by a permutation matrix  $\mathbf{P}_i$ , and eliminating everything below the diagonal is equivalent to multiplying from the left by a special lower-triangular elimination matrix  $\mathbf{L}_i$ , so the whole elimination process can be written as

$$\mathbf{L}_{N-1} (\mathbf{P}_{N-1} \mathbf{L}_{N-2} \mathbf{P}_{N-1}) (\mathbf{P}_{N-1} \mathbf{P}_{N-2} \mathbf{L}_{N-3} \mathbf{P}_{N-2} \mathbf{P}_{N-1}) \cdots (\mathbf{P}_{N-1} \cdots \mathbf{P}_2 \mathbf{L}_1 \mathbf{P}_2 \cdots \mathbf{P}_{N-1}) \mathbf{P}_{N-1} \cdots \mathbf{P}_1 \mathbf{A} = \mathbf{U}.$$

The elimination matrix  $\mathbf{L}_i$  is an identity except for nonzeros below the diagonal of the  $i^{\text{th}}$  column. The matrices in parentheses above are symmetric permutations of  $\mathbf{L}_i$  where the permutations operate on rows

& columns with indices *greater than i*. As a result, the symmetric permutations have exactly the same lower-triangular structure of nonzeros as  $\mathbf{L}_i$ , though the entries below the diagonal are different (they are permuted). We can thus write

$$\mathbf{L}_{N-1} \tilde{\mathbf{L}}_{N-2} \cdots \tilde{\mathbf{L}}_1 \mathbf{P}_{N-1} \cdots \mathbf{P}_1 \mathbf{A} = \mathbf{U}.$$

where all the  $\tilde{\mathbf{L}}_i$  are (special) lower-triangular. The inverse of a (special) lower-triangular matrix is also (special) lower triangular (in fact the inverse of an elimination matrix just changes the sign of the off-diagonals), and a product of permutation matrices is itself a permutation matrix, so we can write

$$\mathbf{PA} = \mathbf{LU}$$

where

$$\mathbf{P} = \mathbf{P}_{N-1} \cdots \mathbf{P}_1$$

and

$$\mathbf{L} = \tilde{\mathbf{L}}_1^{-1} \cdots \tilde{\mathbf{L}}_{N-2}^{-1} \mathbf{L}_{N-1}^{-1}.$$

Thus GEPP is equivalent to a matrix factorization.

The matrix  $\mathbf{U}$  is computed during the forward loop of GEPP. The  $\mathbf{P}$  and  $\mathbf{L}$  matrices can be computed during the GEPP algorithm with no extra cost (since they basically just keep a record of the operations performed by the GEPP algorithm). Generally we don't store the matrix  $\mathbf{P}$  since it only has  $n$  nonzero elements, all of which are 1. Instead we store a permutation vector. The matrix  $\mathbf{L}$  can be computed online easily even with pivoting; see, e.g., Algorithm 1.2.4 in NMMC. Do not use the above analysis to compute  $\mathbf{L}$ !

**3 GEPP: Computational Cost.** Operation count for row reduction (not applied to RHS). We will ignore the cost of row swaps. In step  $k$  you eliminate  $n - k$  rows. At each row you compute the 'multiplier' (one division), then you perform  $n - k$  multiplications and additions. So the cost at step  $k$  is  $2(n - k)^2 + (n - k)$  floating point operations. So the total cost is

$$\sum_{k=1}^{n-1} 2(n - k)^2 + (n - k) \sim \mathcal{O}(n^3).$$

The precise coefficient of the leading order term is  $2n^3/3$ ; generally we only keep the term that grows fastest with  $n$  so that we can gauge cost for large matrices.

The back solve costs  $n$  divisions plus  $n(n - 1)/2$  multiplications and the same number of additions, so overall it is  $\mathcal{O}(n^2)$ . We say that the overall cost of GEPP/permuted LU is order  $n^3$ . Once you've computed the LU factorization of a matrix it only takes  $\mathcal{O}(n^2)$  operations to use it to solve a system.

**4 Cholesky.** If the matrix  $\mathbf{A}$  is symmetric and positive definite, then (i) no pivoting is needed, and (ii) the diagonal elements of  $\mathbf{U}$  are positive. (Applying GE to  $\mathbf{A}$  is the best way to check if a symmetric matrix is positive definite; not by computing the eigenvalues.) Moreover, symmetry implies

$$\mathbf{A} = \mathbf{LDL}^T$$

where  $\mathbf{D}$  is a diagonal matrix whose diagonal entries are those of  $\mathbf{U}$ . So we can define  $\mathbf{G} = \mathbf{LD}^{1/2}$ , and we have  $\mathbf{A} = \mathbf{GG}^T$  where  $\mathbf{G}$  is the Cholesky factor of  $\mathbf{A}$ .

**5 Tridiagonal matrices.** If  $\mathbf{A}$  is tridiagonal then  $\mathbf{L}$  will be lower bidiagonal and  $\mathbf{U}$  will have nonzeros at most two elements above the diagonal (graphical proof). If no pivoting is required (e.g. because  $\mathbf{A}$  is SPD), then  $\mathbf{U}$  will be upper-bidiagonal. In either case the cost of GEPP is only  $\mathcal{O}(n)$ .

---

## Gaussian Elimination & LU Factorization: Roundoff Errors

Atkinson Theorem 8.5, NMMC §1.4.3, & GvL §3.3

Atkinson Theorem 8.5 is not as accurate as the results presented here, but is still OK.

### 1 Backward errors in the computed solution using GE with either complete or partial pivoting.

(cf NMMC Theorem 1.4.6) Let  $\bar{\mathbf{x}}$  be the computed solution to  $\mathbf{A}\mathbf{x} = \mathbf{b}$ . Then there is a matrix  $\mathbf{E}$  such that

$$(\mathbf{A} + \mathbf{E})\bar{\mathbf{x}} = \mathbf{b}$$

where

$$\|\mathbf{E}\|_\infty \leq 1.5 \frac{n^2(n+1)\mathbf{u}}{1-n\mathbf{u}} \rho_n \|\mathbf{A}\|_\infty$$

and  $\rho_n$  is the ‘growth factor.’

Technically this is not a rigorous bound because it ignores terms on the order of  $(n\mathbf{u})^2$ , but when  $n\mathbf{u}$  is small it is accurate. You could insert this into our sensitivity analysis to get a forward error bound on the computed solution; set

$$\epsilon = \|\mathbf{E}\|_\infty / \|\mathbf{A}\|_\infty \leq 1.5 \frac{n^2(n+1)\mathbf{u}}{1-n\mathbf{u}} \rho_n$$

and insert in

$$\frac{\|\bar{\mathbf{x}} - \mathbf{x}\|}{\|\mathbf{x}\|} \leq \frac{2\epsilon\kappa(\mathbf{A})}{1-\epsilon}.$$

### 2 Some results on the growth factor $\rho_n$ (See NMMC)

- With partial pivoting the growth factor is sharply bounded by  $\rho_n \leq 2^{n-1}$ , but this upper bound is usually not attained. Partial pivoting is usually deemed ‘good enough’ for most applications.
- With complete pivoting, the upper bound is  $\rho_n < 1.8n^{1/2+\ln(n)/4}$  (Atkinson forgets the 1/2 in the exponent). Complete pivoting is extremely stable to roundoff errors.
- For a ‘diagonally-dominant’ matrix (definition deferred) and for tridiagonal matrices the growth factor is bounded by  $\rho_n \leq 2$ .

### 3 Small residuals. The residual in the computed solution $\bar{\mathbf{r}} = \mathbf{b} - \mathbf{A}\bar{\mathbf{x}}$ satisfies

$$\frac{\|\bar{\mathbf{r}}\|_\infty}{\|\mathbf{A}\|_\infty \|\bar{\mathbf{x}}\|_\infty} \leq 1.5 \frac{n^2(n+1)\mathbf{u}}{1-n\mathbf{u}} \rho_n.$$

This is independent of the condition number of  $\mathbf{A}$ . So Gaussian elimination can reliably produce solutions that have relatively small residuals, even when the matrix is ill-conditioned. (Again, this is not a rigorous bound because it ignores terms on the order of  $(n\mathbf{u})^2$ .)

4 It’s worth noting that if you’re worried about roundoff errors, the easiest thing to do is to compute the solution twice and examine the differences. For example, compute once with partial pivoting and once with complete pivoting and compare. Or compute once in single precision and once in double precision and compare (or double and quad). If roundoff is limiting your accuracy, there are methods for obtaining improved accuracy; all hope is not lost, just do some reading.

---

## Stationary Iterations: Richardson, Jacobi, Seidel, SOR

Atkinson §8.6, NMMC §4.1, GvL §11.2

1 General idea: Let  $\mathbf{A} = \mathbf{M} - \mathbf{N}$  where it's easy to solve  $\mathbf{M}\mathbf{y} = \mathbf{c}$ , i.e. it's easy to invert  $\mathbf{M}$  (though we never compute  $\mathbf{M}^{-1}$ !). Then define the following iteration

$$\mathbf{x}_{k+1} = \mathbf{M}^{-1}\mathbf{N}\mathbf{x}_k + \mathbf{M}^{-1}\mathbf{b}.$$

This is a 'splitting' for obvious reasons, and it's 'stationary' because it does the same operation at every step. We'll cover some classical splittings and their associated iterations. Generally we want to know whether  $\mathbf{x}_k$  will converge to the solution of  $\mathbf{A}\mathbf{x} = \mathbf{b}$  as  $k \rightarrow \infty$  for every starting vector  $\mathbf{x}_0$ .

First let's suppose that the iteration does converge to something, and call it  $\mathbf{x}_\infty$ . We'll verify that this limit does solve  $\mathbf{A}\mathbf{x} = \mathbf{b}$ :

$$\begin{aligned}\mathbf{x}_\infty &= \mathbf{M}^{-1}\mathbf{N}\mathbf{x}_\infty + \mathbf{M}^{-1}\mathbf{b} \Rightarrow \\ (\mathbf{I} - \mathbf{M}^{-1}\mathbf{N})\mathbf{x}_\infty &= \mathbf{M}^{-1}\mathbf{b} \\ (\mathbf{M} - \mathbf{N})\mathbf{x}_\infty &= \mathbf{A}\mathbf{x}_\infty = \mathbf{b}.\end{aligned}$$

Next we'll ask under what conditions the splitting will converge. It's easiest to show that the error goes to 0, rather than that  $\mathbf{x}_k \rightarrow \mathbf{x}_\infty$ . The error is  $\mathbf{e}_k = \mathbf{x}_\infty - \mathbf{x}_k$ , and it solves (use  $\mathbf{M}\mathbf{x}_\infty = \mathbf{N}\mathbf{x}_\infty + \mathbf{b}$ )

$$\mathbf{e}_{k+1} = \mathbf{B}\mathbf{e}_k, \quad \mathbf{B} = \mathbf{M}^{-1}\mathbf{N}.$$

Now we want to know under what conditions this will converge to 0 for any initial error. A necessary and sufficient condition is:

The iteration is convergent (i.e. converges for any initial condition) iff the largest eigenvalue of  $\mathbf{B}$  has amplitude less than 1. (I.e. the spectral radius  $\rho(\mathbf{B}) < 1$ .)

See Atkinson Theorem 7.9. This is basically linear algebra background, so we won't prove it.

It's not always easy to check this condition. The following condition is sufficient but not necessary:

If  $\|\mathbf{B}\| < 1$  for some matrix norm induced by a vector norm, then the iteration is convergent.

The proof of this is simple and is omitted.

How fast is the iteration converging?

If  $\rho(\mathbf{B}) < 1$  then

$$\lim_{k \rightarrow \infty} \left( \frac{\|\mathbf{e}_k\|}{\|\mathbf{e}_0\|} \right)^{1/k} = \rho(\mathbf{B})$$

for any vector norm.

This means that for large enough  $k$ ,

$$\|\mathbf{e}_k\| \lesssim \|\mathbf{e}_0\| \rho(\mathbf{B})^k.$$

The proof of the above is based on a nontrivial fact about the spectral radius, found in NMMC Lemma 4.1.2. If the iteration matrix  $\mathbf{B}$  is non-normal, actual errors can *grow* at low  $k$ , even when the method is convergent. They can grow by many orders of magnitude before eventually settling into asymptotic decay, so be careful.

2 The standard splittings first decompose  $\mathbf{A}$  into diagonal, upper, and lower parts:

$$\mathbf{A} = \mathbf{D} - \mathbf{E} - \mathbf{F} = \mathbf{D}(\mathbf{I} - \mathbf{L} - \mathbf{U}).$$

- Gauss-Jacobi, aka Jacobi. Based on the idea that it's easy to invert a diagonal matrix. The splitting is  $\mathbf{M} = \mathbf{D}$ , and the iteration is

$$\mathbf{x}_{k+1} = \mathbf{D}^{-1}(\mathbf{E} + \mathbf{F})\mathbf{x}_k + \mathbf{D}^{-1}\mathbf{b}.$$

- Gauss-Seidel. Based on the idea that a triangular matrix is easy to invert. The splitting is  $\mathbf{M} = \mathbf{D} - \mathbf{E}$ , and the iteration is

$$\mathbf{x}_{k+1} = (\mathbf{D} - \mathbf{E})^{-1}\mathbf{F}\mathbf{x}_k + (\mathbf{D} - \mathbf{E})^{-1}\mathbf{b}.$$

- Successive Over-Relaxation (SOR). Part of the diagonal goes into  $\mathbf{M}$  and part goes into  $\mathbf{N}$ :

$$\mathbf{x}_{k+1} = (\mathbf{D} - \omega\mathbf{E})^{-1}[\omega\mathbf{F} + (1 - \omega)\mathbf{D}]\mathbf{x}_k + \omega(\mathbf{D} - \omega\mathbf{E})^{-1}\mathbf{b}.$$

Both GS and SOR have 'symmetric' variants. A single step of symmetric GS consists of a step with  $\mathbf{M} = \mathbf{D} - \mathbf{E}$ , followed by a step with  $\mathbf{M} = \mathbf{D} - \mathbf{F}$ .

### Convergence Theorems

(Definition) The square matrix  $\mathbf{A}$  is diagonally dominant by rows when the absolute value of the diagonal element in each row is larger than the sum of the absolute values of the remaining terms

$$|a_{ii}| \geq \sum_{j \neq i} |a_{ij}|.$$

"Strict" diagonal dominance replaces the  $\geq$  by  $>$ . Diagonal dominance by columns is similar.

The Gershgorin Disk theorem (Atkinson Theorem 9.1) guarantees that strictly diagonally dominant matrices are nonsingular.

**Convergence Theorem #1** If  $\mathbf{A}$  is strictly diagonally dominant then the Gauss Jacobi and Seidel iterations are convergent.

Proof for Jacobi: The Jacobi iteration matrix is  $\mathbf{B} = \mathbf{D}^{-1}(\mathbf{E} + \mathbf{F})$ . If  $\mathbf{A}$  is row-dominant then the  $\infty$  norm is strictly less than 1. (If  $\mathbf{A}$  is column-dominant then the 1 norm is strictly less than 1.)

Proof for Seidel: Note  $\mathbf{B} = (\mathbf{D} - \mathbf{E})^{-1}\mathbf{F}$ . Let  $\mathbf{e}_k$  be the standard basis vector such that  $\|\mathbf{B}^T\|_1 = \|\mathbf{B}\|_\infty = \|\mathbf{B}^T\mathbf{e}_k\|_1$ . Note that

$$\begin{aligned} (\mathbf{I} - \mathbf{D}^{-1}\mathbf{E})\mathbf{B} &= \mathbf{D}^{-1}(\mathbf{D} - \mathbf{E})(\mathbf{D} - \mathbf{E})^{-1}\mathbf{F} = \mathbf{D}^{-1}\mathbf{F} \Rightarrow \\ \mathbf{B} &= \mathbf{D}^{-1}\mathbf{E}\mathbf{B} + \mathbf{D}^{-1}\mathbf{F} \end{aligned}$$

Take transpose, then multiply by  $\mathbf{e}_k$  and take the 1-norm

$$\begin{aligned} \|\mathbf{B}\|_\infty &\leq \|\mathbf{B}\|_\infty \|(\mathbf{D}^{-1}\mathbf{E})^T\mathbf{e}_k\|_1 + \|(\mathbf{D}^{-1}\mathbf{F})^T\mathbf{e}_k\|_1 \\ \|\mathbf{B}\|_\infty &\leq \frac{\|(\mathbf{D}^{-1}\mathbf{F})^T\mathbf{e}_k\|_1}{1 - \|(\mathbf{D}^{-1}\mathbf{E})^T\mathbf{e}_k\|_1} \end{aligned}$$

To justify the division without reversing the inequality we need that  $\|(\mathbf{D}^{-1}\mathbf{E})^T\mathbf{e}_k\|_1 < 1$ . The fact that  $\mathbf{A}$  is strictly-diagonally-dominant by rows guarantees this (since the vector is a column sum of a transpose). Now we're trying to prove that  $\|\mathbf{B}\|_\infty < 1$ ; this is transparent if

$$\|(\mathbf{D}^{-1}\mathbf{E})^T\mathbf{e}_k\|_1 + \|(\mathbf{D}^{-1}\mathbf{F})^T\mathbf{e}_k\|_1 < 1.$$

Again this is just a full row sum of the Jacobi iteration matrix broken into two absolute-value pieces (absolute value of the sum of elements left of the diagonal plus absolute value of the sum of those right of the diagonal); strict diagonal-dominance then guarantees the desired result. Column-wise DD is similar.

It is possible to weaken the conditions to *irreducible diagonal dominance*.



**Convergence Theorem #2** If  $\mathbf{A}$  is symmetric positive definite then the Gauss Seidel iteration is convergent.

Proof from GvL, Theorem 11.2.3. The proof examines the spectral radius, not the norm (naturally, since we only know something about the eigenvalues of  $\mathbf{A}$ ). The iteration matrix for a symmetric  $\mathbf{A}$  is

$$\mathbf{B}_{GS} = (\mathbf{D} - \mathbf{E})^{-1}\mathbf{E}^T.$$

First notice that this matrix has the same eigenvalues as

$$\mathbf{G} = \mathbf{D}^{1/2}\mathbf{B}\mathbf{D}^{-1/2} = (\mathbf{I} - \mathbf{L})^{-1}\mathbf{L}^T$$

where

$$\mathbf{L} = \mathbf{D}^{-1/2}\mathbf{E}\mathbf{D}^{-1/2}.$$

Suppose that  $\mathbf{v}$  is a unit-2-norm eigenvector of  $\mathbf{G}$

$$(\mathbf{I} - \mathbf{L})^{-1}\mathbf{L}^T\mathbf{v} = \lambda\mathbf{v}.$$

Move inverse to RHS then take standard inner product:

$$\mathbf{v}^*\mathbf{L}^T\mathbf{v} = \lambda(1 - \mathbf{v}^*\mathbf{L}\mathbf{v}).$$

Note that (if  $\mathbf{A}$  is real) then  $\mathbf{v}^*\mathbf{L}^T\mathbf{v} = \overline{\mathbf{v}^*\mathbf{L}\mathbf{v}}$ . Let  $\mathbf{v}^*\mathbf{L}\mathbf{v} = a + ib$ , so that

$$a - ib = \lambda(1 - a - ib) \Rightarrow |\lambda|^2 = \frac{a^2 + b^2}{1 - 2a + a^2 + b^2}.$$

Now consider that since  $\mathbf{A}$  is SPD, we have

$$0 < \mathbf{v}^*\mathbf{D}^{-1/2}\mathbf{A}\mathbf{D}^{-1/2}\mathbf{v} = 1 - \mathbf{v}^*\mathbf{L}\mathbf{v} - \mathbf{v}^*\mathbf{L}^T\mathbf{v} = 1 - 2a.$$

This implies that  $|\lambda|^2 < 1$ , as desired.

**SOR** First show that a necessary condition is  $0 < \omega < 2$ . The SOR iteration matrix is

$$\mathbf{B} = (\mathbf{D} - \omega\mathbf{E})^{-1}[\omega\mathbf{F} + (1 - \omega)\mathbf{D}] = (\mathbf{I} - \mathbf{L})^{-1}[(1 - \omega)\mathbf{I} + \omega\mathbf{U}].$$

The determinant of  $\mathbf{B}$  is

$$\det[\mathbf{B}] = \det[(\mathbf{I} - \mathbf{L})^{-1}]\det[(1 - \omega)\mathbf{I} + \omega\mathbf{U}] = 1 \times (1 - \omega)^n.$$

The determinant is the product of the eigenvalues. Let  $\lambda_i$  be the eigenvalues of  $\mathbf{B}$ , then

$$\rho(\mathbf{B})^n = |\lambda_{\max}|^n \geq \prod_{i=1}^n |\lambda_i|^n = |1 - \omega|^n.$$

In order to have  $\rho(\mathbf{B}) < 1$  we have to have  $0 < \omega < 2$ .

**Convergence Theorem #3** If  $\mathbf{A}$  is symmetric positive definite and  $0 < \omega < 2$  then SOR is convergent.

This is Theorem 4.10 from Saad “Iterative methods for sparse linear systems” given there without proof or reference. SOR was a hot item ~40 years ago, and has a very well-developed theory. It’s not widely used any more as a stand-alone method.

---

## Steepest Descent & Conjugate Gradients

**1** We will next look at Conjugate Gradient algorithm, which is a bridge between iterative & direct methods. It is ‘iterative’ in the sense that it produces a new, better approximation at each step, but it is also ‘direct’ in the sense that after  $n$  steps in exact arithmetic it produces the true solution. To understand CG, we start with steepest descent, which is a true iterative method. Both CG and steepest descent apply to systems with  $\mathbf{A}$  symmetric positive-definite.

Consider

$$f(\mathbf{x}) = \frac{1}{2} \mathbf{x}^T \mathbf{A} \mathbf{x} - \mathbf{x}^T \mathbf{b}$$

where  $\mathbf{A}$  is SPD. The gradient is

$$\nabla f = \mathbf{A} \mathbf{x} - \mathbf{b}.$$

The only critical point is the solution to  $\mathbf{A} \mathbf{x} = \mathbf{b}$ . The Hessian is  $\mathbf{A}$ ; since it’s SPD, the critical point is a minimizer. This function is quadratic, so the solution to  $\mathbf{A} \mathbf{x} = \mathbf{b}$  is the unique global minimizer of  $f$ . Solving  $\mathbf{A} \mathbf{x} = \mathbf{b}$  is therefore equivalent to finding the minimizer of  $f$ .

Steepest descent is an iterative method for finding the minimum of  $f$ . It takes the form

$$\mathbf{x}_{k+1} = \mathbf{x}_k - \alpha_k \mathbf{g}_k$$

where  $\mathbf{g}_k$  is the gradient of  $f$  at  $\mathbf{x}_k$ , and  $\alpha_k$  is some scalar chosen so that  $f(\mathbf{x}_{k+1}) < f(\mathbf{x}_k)$ . The gradient of  $f$  at  $\mathbf{x}_k$  is

$$\mathbf{g}_k = \nabla f|_{\mathbf{x}_k} = \mathbf{A} \mathbf{x}_k - \mathbf{b} = -\mathbf{r}_k$$

where  $\mathbf{r}_k$  is the residual at step  $k$ . (Recall that the direction of steepest increase of a function is the gradient, and the direction of steepest decrease/descent is minus the gradient.)

How can we choose  $\alpha_k$ ? Note that

$$\begin{aligned} f(\mathbf{x}_{k+1}) &= \frac{1}{2} (\mathbf{x}_k - \alpha_k \mathbf{g}_k)^T \mathbf{A} (\mathbf{x}_k - \alpha_k \mathbf{g}_k) - (\mathbf{x}_k - \alpha_k \mathbf{g}_k)^T \mathbf{b} \\ &= f(\mathbf{x}_k) - \alpha_k (\mathbf{g}_k^T \mathbf{A} \mathbf{x}_k - \mathbf{g}_k^T \mathbf{b}) + \frac{\alpha_k^2}{2} \mathbf{g}_k^T \mathbf{A} \mathbf{g}_k \\ &= f(\mathbf{x}_k) - \alpha_k \mathbf{g}_k^T \mathbf{g}_k + \frac{\alpha_k^2}{2} \mathbf{g}_k^T \mathbf{A} \mathbf{g}_k \end{aligned}$$

An exact line search chooses  $\alpha_k$  to minimize this expression; the minimum occurs at

$$\alpha_k = \frac{\mathbf{g}_k^T \mathbf{g}_k}{\mathbf{g}_k^T \mathbf{A} \mathbf{g}_k} = \frac{\|\mathbf{r}_k\|_2^2}{\mathbf{r}_k^T \mathbf{A} \mathbf{r}_k}.$$

This is the steepest descent method with exact line search. It can be quite slow, especially when  $\mathbf{A}$  is ill-conditioned; the objective function  $f$  has level sets that are highly-anisotropic ellipsoids and the steepest-descent method jumps back & forth a lot.

Now consider if and how the error decreases. The error dynamics is

$$\mathbf{e}_{k+1} = \mathbf{e}_k - \alpha_k \mathbf{r}_k.$$

Since  $\mathbf{A}$  is SPD, it induces a norm:

$$\|\mathbf{x}\|_A^2 = \mathbf{x}^T \mathbf{A} \mathbf{x}.$$

The steepest-descent method decreases the A-norm of the error at every step:

$$\begin{aligned} \|\mathbf{e}_{k+1}\|_A^2 &= \|\mathbf{e}_k\|_A^2 - 2\alpha_k \mathbf{r}_k^T \mathbf{A} \mathbf{e}_k + \alpha_k^2 \|\mathbf{r}_k\|_A^2 \\ &= \|\mathbf{e}_k\|_A^2 - 2\alpha_k \|\mathbf{r}_k\|_2^2 + \alpha_k^2 \|\mathbf{r}_k\|_A^2 \\ &= \|\mathbf{e}_k\|_A^2 - \frac{\|\mathbf{r}_k\|_2^4}{\|\mathbf{r}_k\|_A^2} < \|\mathbf{e}_k\|_A^2 \end{aligned}$$

This fact by itself is not enough to guarantee convergence, but steepest descent is in fact convergent whenever  $\mathbf{A}$  is SPD. The rate of convergence can be bounded as follows (Saad §5.3.1). The  $\mathbf{A}$ -norm of the error at step  $k + 1$  is bounded by

$$\|\mathbf{x}^{(k+1)} - \mathbf{x}_*\|_A \leq \frac{\lambda_{\max} - \lambda_{\min}}{\lambda_{\max} + \lambda_{\min}} \|\mathbf{x}^{(k)} - \mathbf{x}_*\|_A$$

where  $\lambda$  are eigenvalues of  $\mathbf{A}$ . As usual, the convergence can be slow when  $\mathbf{A}$  is ill-conditioned.

**2** Steepest-descent minimizes  $f$  over a 1D affine subspace. Perhaps we could improve by minimizing over a larger subspace, but which subspace to choose? Steepest descent uses the residual, which is easy to compute. Consider the following:

1. First do a steepest-descent step.
2. Now, compute the new residual, but find the minimum in the affine subspace  $\mathbf{x}_1 + \text{span}\{\mathbf{r}_0, \mathbf{r}_1\}$  instead of just  $\mathbf{x}_1 + \text{span}\{\mathbf{r}_1\}$ .
3. At step  $k$ , find the minimum in the affine subspace  $\mathbf{x}_k + \text{span}\{\mathbf{r}_0, \dots, \mathbf{r}_k\}$ .

Try it for step 2:

$$\begin{aligned} \mathbf{x}_2 &= \mathbf{x}_1 + \alpha_0 \mathbf{r}_0 + \alpha_1 \mathbf{r}_1 \\ f(\mathbf{x}_2) &= f(\mathbf{x}_1) + \alpha_0 \mathbf{r}_0^T (\mathbf{A} \mathbf{x}_1 - \mathbf{b}) + \alpha_1 \mathbf{r}_1^T (\mathbf{A} \mathbf{x}_1 - \mathbf{b}) + \alpha_0 \alpha_1 \mathbf{r}_0^T \mathbf{A} \mathbf{r}_1 \\ f(\mathbf{x}_2) &= f(\mathbf{x}_1) + \alpha_0 \mathbf{r}_0^T \mathbf{r}_1 + \alpha_1 \mathbf{r}_1^T \mathbf{r}_1 + \alpha_0 \alpha_1 \mathbf{r}_0^T \mathbf{A} \mathbf{r}_1 + \frac{1}{2} (\alpha_0^2 \mathbf{r}_0^T \mathbf{A} \mathbf{r}_0 + \alpha_1^2 \mathbf{r}_1^T \mathbf{A} \mathbf{r}_1) \end{aligned}$$

We want to find the  $\alpha_0$  and  $\alpha_1$  that minimize  $f(\mathbf{x}_2)$ . This is exactly the same kind of problem we started with, but 2D; at step  $k$  it will be a  $k$ -dimensional quadratic minimization problem. How will we solve it? Steepest descents? The same algorithm we're using now? That way lies madness.

**3** The problem with the foregoing algorithm is that at step  $k$  we have to minimize a quadratic in  $k$  variables, which is equivalent to solving a  $k \times k$  SPD linear system, which is basically the problem we were trying to solve originally.

The key to the solution is to change basis. We used the obvious basis for  $\text{span}\{\mathbf{r}_0, \mathbf{r}_1\}$ , but this basis is inconvenient. Every  $n$ -dimensional subspace has an orthogonal basis; orthogonality is always defined with respect to an inner product. Rather than use the standard inner product (dot), use a basis that's orthogonal with respect to the following inner product:  $\langle \mathbf{u}, \mathbf{v} \rangle = \mathbf{u}^T \mathbf{A} \mathbf{v}$ . This is only an inner product when  $\mathbf{A}$  is SPD (in fact, any inner product on  $\mathbb{R}^n$  can be written in this form for some SPD matrix). Vectors that are 'orthogonal' with respect to this inner product are called 'conjugate' or ' $\mathbf{A}$ -conjugate.'

**4** For convenience suppose that  $\mathbf{x}_0 = 0$  so that  $\mathbf{r}_0 = \mathbf{b}$ . The CG algorithm starts with a steepest-descent step

$$\mathbf{x}_1 = \mathbf{x}_0 + a \mathbf{r}_0 = a \mathbf{r}_0.$$

The next step seeks to minimize over the subspace  $\mathbf{x}_2 = \mathbf{x}_1 + b \mathbf{r}_0 + c \mathbf{r}_1 = (a+b) \mathbf{r}_0 + c \mathbf{r}_1$ . Because we started at 0, we're seeking a solution within the subspace spanned by  $\{\mathbf{r}_0, \mathbf{r}_1\}$  (rather than an affine subspace if we started at a nonzero location). At step  $k + 1$  we're seeking  $\mathbf{x}_{k+1} \in \text{span}\{\mathbf{r}_0, \dots, \mathbf{r}_k\}$ .

Suppose  $\{\mathbf{p}_0, \dots, \mathbf{p}_k\}$  is any  $\mathbf{A}$ -conjugate basis for the span of  $\mathbf{r}_0, \dots, \mathbf{r}_k$ . Then

$$\mathbf{x}_{k+1} = \sum_{i=0}^k \alpha_i \mathbf{p}_i$$

I really should have used  $\alpha_{i,k}$ , but you'll see why I don't need to shortly. Look at  $f(\mathbf{x}_{k+1})$  using conjugacy

$$f(\mathbf{x}_{k+1}) = \sum_{i=0}^k \mathbf{p}_i^T \left( \frac{\alpha_i^2}{2} \mathbf{A} \mathbf{p}_i - \alpha_i \mathbf{b} \right).$$

The  $\alpha_i$  that minimize this are

$$\alpha_i = \frac{\mathbf{p}_i^T \mathbf{b}}{\mathbf{p}_i^T \mathbf{A} \mathbf{p}_i}.$$

Clearly this doesn't depend on  $k$ , unless we choose a different set of  $\mathbf{p}_i$  at every step. The next section shows that we don't need to do that: at each new step  $k$  we just add a new  $\mathbf{p}_k$  to our basis. It also means that

$$\mathbf{x}_{k+1} = \mathbf{x}_k + \alpha_k \mathbf{p}_k$$

(because the  $\mathbf{p}$  vectors don't change as you go from step to step). We will also need to note that  $\alpha_k \neq 0$  unless  $\mathbf{x}_k$  is the solution. Since the Hessian of  $f$  is positive definite, expanding the search space will always lead to a decrease in  $f$  unless you're already at the minimizer.

**5** To move to step  $k + 1$  we need to construct an  $\mathbf{A}$ -conjugate basis for the span of  $\mathbf{r}_0, \dots, \mathbf{r}_{k+1}$ . We could do this using Gram-Schmidt:

$$\mathbf{p}_{k+1} = \mathbf{r}_{k+1} - \frac{\langle \mathbf{p}_0, \mathbf{r}_{k+1} \rangle}{\|\mathbf{p}_0\|_A^2} \mathbf{p}_0 - \dots - \frac{\langle \mathbf{p}_k, \mathbf{r}_{k+1} \rangle}{\|\mathbf{p}_k\|_A^2} \mathbf{p}_k.$$

This will only work if  $\mathbf{r}_{k+1}$  is not in the span of the previous residuals.

Now I will show that  $\mathbf{r}_{k+1}$  is orthogonal (not conjugate) to all previous residuals, which implies that it's not in the span (unless it's 0, in which case we're done:  $\mathbf{x}_{k+1}$  is the solution). First compute the residual at  $k + 1$ :

$$\mathbf{r}_{k+1} = \mathbf{b} - \mathbf{A} \mathbf{x}_{k+1} = \mathbf{b} - \sum_{i=0}^k \alpha_i \mathbf{A} \mathbf{p}_i.$$

Take the dot product with  $\mathbf{p}_j$ :

$$\mathbf{p}_j^T \mathbf{r}_{k+1} = \mathbf{p}_j^T \mathbf{b} - \sum_{i=0}^k \alpha_i \mathbf{p}_j^T \mathbf{A} \mathbf{p}_i$$

Using conjugacy

$$\mathbf{p}_j^T \mathbf{r}_{k+1} = \mathbf{p}_j^T \mathbf{b} - \alpha_j \mathbf{p}_j^T \mathbf{A} \mathbf{p}_j = 0, \quad j = 0, \dots, k$$

Last equality used definition of  $\alpha_j$ . This only shows that  $\mathbf{r}_{k+1}$  is orthogonal to  $\mathbf{p}_i$  for  $i = 1, \dots, k$ , but since the  $\mathbf{p}$  vectors span the same subspace as the  $\mathbf{r}$  vectors this result implies what we want to prove, i.e. that  $\mathbf{r}_{k+1}$  is orthogonal to all previous residuals.

This section shows that either (i)  $\mathbf{r}_{k+1} = 0$  in which case you've found the exact solution, or (ii)  $\mathbf{r}_{k+1}$  is orthogonal to all previous ones so you can use Gram-Schmidt to construct  $\mathbf{p}_{k+1}$ .

**6** We could stop there, but when you implement you'll find that almost all the terms in the Gram-Schmidt construction of  $\mathbf{p}_{k+1}$  are zero, leaving

$$\mathbf{p}_{k+1} = \mathbf{r}_{k+1} - \frac{\langle \mathbf{p}_k, \mathbf{r}_{k+1} \rangle}{\|\mathbf{p}_k\|_A^2} \mathbf{p}_k.$$

The goal of this section is to prove that statement, i.e. to prove that

$$\langle \mathbf{p}_j, \mathbf{r}_{k+1} \rangle = 0, \quad j = 0, \dots, k - 1.$$

- Recall that  $\mathbf{r}_{k+1}$  is orthogonal to the span of the  $\mathbf{p}_j$ ; this implies that  $\mathbf{r}_{k+1}$  must also be orthogonal to all the  $\mathbf{r}_j$  ( $j = 1, \dots, k$ ) (because span of  $\mathbf{p}$ 's is equal to the span of the  $\mathbf{r}$ 's).
- Now notice a recursion for the residuals

$$\mathbf{r}_{k+1} = \mathbf{b} - \mathbf{A} \mathbf{x}_{k+1} = \mathbf{b} - \mathbf{A}(\mathbf{x}_k + \alpha_k \mathbf{p}_k) = \mathbf{r}_k - \alpha_k \mathbf{A} \mathbf{p}_k.$$

Now take the dot product with  $\mathbf{r}_{k+1}$ , using orthogonality noted above

$$\|\mathbf{r}_{k+1}\|_2^2 = -\alpha_k \langle \mathbf{r}_{k+1}, \mathbf{p}_k \rangle \tag{*}$$

- Now expand the recursion

$$\mathbf{r}_{k+1} = \mathbf{b} - \mathbf{A}(\mathbf{x}_{k-1} + \alpha_{k-1}\mathbf{p}_{k-1} + \alpha_k\mathbf{p}_k) = \mathbf{r}_{k-1} - \alpha_{k-1}\mathbf{A}\mathbf{p}_{k-1} - \alpha_k\mathbf{A}\mathbf{p}_k.$$

Now take the dot product with  $\mathbf{r}_{k+1}$

$$\|\mathbf{r}_{k+1}\|_2^2 = -\alpha_{k-1}\langle \mathbf{r}_{k+1}, \mathbf{p}_{k-1} \rangle - \alpha_k\langle \mathbf{r}_{k+1}, \mathbf{p}_k \rangle$$

Using (\*) this implies

$$0 = -\alpha_{k-1}\langle \mathbf{r}_{k+1}, \mathbf{p}_{k-1} \rangle$$

(Recall that we noted that  $\alpha_k \neq 0$  unless  $\mathbf{x}_k$  is the solution.)

- Continuing inductively yields the desired result. I.e. expand again and use previous results to show  $\alpha_{k-j}\langle \mathbf{r}_{k+1}, \mathbf{p}_{k-1} \rangle = 0$ .

**7** We just showed that

$$\mathbf{p}_{k+1} = \mathbf{r}_{k+1} - \frac{\langle \mathbf{p}_k, \mathbf{r}_{k+1} \rangle}{\|\mathbf{p}_k\|_A^2} \mathbf{p}_k.$$

Computing the coefficient requires  $\mathcal{O}(n^2)$  flops for a non-sparse matrix  $\mathbf{A}$ . It turns out that there is an equivalent way of computing the coefficient that only costs  $\mathcal{O}(n)$  flops. To wit,

$$\langle \mathbf{p}_k, \mathbf{r}_{k+1} \rangle = \mathbf{p}_k^T \mathbf{A} \mathbf{r}_{k+1} = \mathbf{r}_{k+1}^T \mathbf{A} \mathbf{p}_k$$

Now use

$$\mathbf{r}_{k+1} = \mathbf{r}_k - \alpha_k \mathbf{A} \mathbf{p}_k \Rightarrow \mathbf{A} \mathbf{p}_k = -\frac{1}{\alpha_k} (\mathbf{r}_{k+1} - \mathbf{r}_k)$$

Plugging in and using that  $\mathbf{r}_{k+1}^T \mathbf{r}_k = 0$

$$\langle \mathbf{p}_k, \mathbf{r}_{k+1} \rangle = \frac{1}{\alpha_k} \mathbf{r}_{k+1}^T \mathbf{r}_{k+1} = -\frac{\|\mathbf{r}_{k+1}\|_2^2}{\alpha_k}.$$

Now look at the denominator

$$\|\mathbf{p}_k\|_A^2 = \mathbf{p}_k^T \mathbf{A} \mathbf{p}_k = (\mathbf{r}_k - \alpha_k \mathbf{p}_{k-1})^T \mathbf{A} \mathbf{p}_k = \mathbf{r}_k^T \left( \frac{1}{\alpha_k} (\mathbf{r}_{k+1} - \mathbf{r}_k) \right) = \frac{\mathbf{r}_k^T \mathbf{r}_k}{\alpha_k} = \frac{\|\mathbf{r}_k\|_2^2}{\alpha_k}$$

(using conjugacy of  $\mathbf{p}_{k-1}$  and  $\mathbf{p}_k$  and orthogonality of  $\mathbf{r}_{k+1}$  and  $\mathbf{r}_k$ ). So we can write

$$\mathbf{p}_{k+1} = \mathbf{r}_{k+1} - \frac{\langle \mathbf{p}_k, \mathbf{r}_{k+1} \rangle}{\|\mathbf{p}_k\|_A^2} \mathbf{p}_k = \mathbf{r}_{k+1} + \frac{\|\mathbf{r}_{k+1}\|_2^2}{\|\mathbf{r}_k\|_2^2} \mathbf{p}_k.$$

The rightmost expression only costs  $\mathcal{O}(n)$  flops to evaluate, and is therefore cheaper than the middle expression.

**8** The CG method will produce the exact solution for any SPD  $\mathbf{A}$  after at most  $n$  steps in exact arithmetic. In finite-precision arithmetic we usually stop the iteration well before  $n$  steps, once the residual is small enough. We will not discuss methods for deciding when to stop the iteration.