A Geometric Interpretation of Gauss-Seidel and its Extensions

Take $A = \begin{bmatrix} 2 & 1 \\ 1 & 2 \end{bmatrix}$ and $b = \begin{bmatrix} 3 \\ 3 \end{bmatrix}$. Perform Gauss-Seidel on $A\boldsymbol{x} = b$, knowing that the true solution $\boldsymbol{x} = (1,1)^{\mathrm{T}}$. We have the iterates

$$\begin{aligned} \boldsymbol{x}^{(0)} &= (0,0)^{\mathrm{T}} \\ \boldsymbol{x}^{(1)} &= \left(\frac{3}{2}, \frac{3}{4}\right)^{\mathrm{T}} \\ \boldsymbol{x}^{(2)} &= \left(\frac{9}{8}, \frac{15}{16}\right)^{\mathrm{T}} \end{aligned}$$

Plotting these points in the x_1-x_2 plane, we find that when we travel from $\boldsymbol{x}^{(1)}$ to $\boldsymbol{x}^{(2)}$, we wish we go a little harder so that we actually hit the true solution in this direction. In fact, this direction is characterized by the vector $\boldsymbol{x}^{(2)} - \boldsymbol{x}^{(1)}$. We now want to control how "hard" we move in this direction at every step of the Gauss-Seidel iteration, namely, put a control on $\boldsymbol{x}^{(k+1)} - \boldsymbol{x}^{(k)}$ so that we move further than Gauss-Seidel intends to.

Successive over-Relaxation

Under Gauss-Seidel iteration,

$$\text{(Gauss-Seidel)} \quad \boldsymbol{x}^{(k+1)} = -\left(D+L\right)^{-1} U\boldsymbol{x}^{(k)} + \left(D+L\right)^{-1}\boldsymbol{b},$$

or

$$\left(D+L\right)\boldsymbol{x}^{(k+1)} = \boldsymbol{b} - U\boldsymbol{x}^{(k)} \implies D\boldsymbol{x}^{(k+1)} = \boldsymbol{b} - L\boldsymbol{x}^{(k+1)} - U\boldsymbol{x}^{(k)}.$$

Multiplying D^{-1} on the left of both sides, we have

$$\boldsymbol{x}^{(k+1)} = D^{-1}\left(\boldsymbol{b} - L\boldsymbol{x}^{(k+1)} - U\boldsymbol{x}^{(k)}\right).$$

In practice, we want to control the jump size in the direction of $\boldsymbol{x}^{(k+1)} - \boldsymbol{x}^{(k)}$. We now subtract $\boldsymbol{x}^{(k)}$ from both sides to obtain

$$\boldsymbol{x}^{(k+1)} - \boldsymbol{x}^{(k)} = D^{-1}\left(\boldsymbol{b} - L\boldsymbol{x}^{(k+1)} - D\boldsymbol{x}^{(k)} - U\boldsymbol{x}^{(k)}\right).$$

One can think of the RHS as the **Gauss-Seidel Correction**, in the sense that $\boldsymbol{x}^{(k)}$ goes to $\boldsymbol{x}^{(k+1)}$ in this direction. We label this $\left(\boldsymbol{x}^{(k+1)} - \boldsymbol{x}^{(k)}\right)_{GS}$.

Now, from the example we just saw, going from $\boldsymbol{x}^{(k)}$ to $\boldsymbol{x}^{(k+1)}$ using the exact vector $\left(\boldsymbol{x}^{(k+1)} - \boldsymbol{x}^{(k)}\right)$ may not necessarily yield the "quickest" approach to the true solution. We may need to propel from $\boldsymbol{x}^{(k)}$ a little harder in the direction of $\boldsymbol{x}^{(k+1)} - \boldsymbol{x}^{(k)}$, namely,

$$\boldsymbol{x}^{(k+1)} = \boldsymbol{x}^{(k)} + \omega\left(\boldsymbol{x}^{(k+1)} - \boldsymbol{x}^{(k)}\right)_{GS}, \quad \omega > 1.$$

that is, the true landing spot should be in the same direction of the GS Correction but we step on the gas harder.

Using the definition of GS Correction, we then have the update

$$\boldsymbol{x}^{(k+1)} = \boldsymbol{x}^{(k)} + \omega D^{-1}\left(\boldsymbol{b} - L\boldsymbol{x}^{(k+1)} - D\boldsymbol{x}^{(k)} - U\boldsymbol{x}^{(k)}\right), \quad \omega > 1.$$

Now, rearranging things a little by dividing both sides by ω and then multiplying on the left by D, we have

$$\begin{aligned} \frac{1}{\omega}D\boldsymbol{x}^{(k+1)} &= \frac{1}{\omega}D\boldsymbol{x}^{(k)} + \boldsymbol{b} - L\boldsymbol{x}^{(k+1)} - D\boldsymbol{x}^{(k)} - U\boldsymbol{x}^{(k)} \\ \implies \left(\frac{1}{\omega}D + L\right)\boldsymbol{x}^{(k+1)} &= \left(\left(\frac{1}{\omega} - 1\right)D - U\right)\boldsymbol{x}^{(k)} + \boldsymbol{b} \\ \overset{\text{multiply by } \omega}{\implies} \left(D + \omega L\right)\boldsymbol{x}^{(k+1)} &= \left(\left(1-\omega\right)D - \omega U\right)\boldsymbol{x}^{(k)} + \omega\boldsymbol{b} \\ \overset{\text{invert } (D+\omega L)}{\implies} \boldsymbol{x}^{(k+1)} &= \left(D+\omega L\right)^{-1}\left(\left(1-\omega\right)D - \omega U\right)\boldsymbol{x}^{(k)} + \left(D+\omega L\right)^{-1}\omega\boldsymbol{b}. \end{aligned}$$

This method is called **Successive over-Relaxation** (**SOR**), where $1 < \omega < 2$, and we note that if $\omega = 1$, we retrieve Gauss-Seidel (check it!).

In the form of an iteration scheme, $\boldsymbol{x}^{(k+1)} = T\boldsymbol{x}^{(k)} + \boldsymbol{c}$, we identify that the iterative operator/matrix of the SOR method is

$$T_{SOR} = (D + \omega L)^{-1} \left((1 - \omega) D - \omega U \right)$$

and

$$\boldsymbol{c} = (D + \omega L)^{-1} \omega \boldsymbol{b}.$$

Your Textbook (Optional Reading)
Perspective from the Residual Vector $\boldsymbol{r}^{(k)} = \boldsymbol{b} - A\boldsymbol{x}^{(k)}$

Again, we consider Gauss-Seidel iteration. In component form, we have

$$\text{(Gauss-Seidel)} \quad x_i^{(k+1)} = \frac{1}{a_{ii}} \left[b_i - \sum_{j=1}^{i-1} a_{ij} x_j^{(k+1)} - \sum_{j=i+1}^{n} a_{ij} x_j^{(k)} \right].$$

Now, when we are updating the i^{th} component of $\boldsymbol{x}^{(k+1)}$, we already have the update

$$\boldsymbol{x}_i^{(k+1)} = \left(x_1^{(k+1)}, x_2^{(k+1)}, \ldots, x_{i-1}^{(k+1)}, x_i^{(k)}, \ldots, x_n^{(k)} \right)$$

since we are using the immediate information from the $(k+1)^{th}$ iteration. Here the subscript i in $\boldsymbol{x}_i^{(k+1)}$ is to indicate that we are within the $(k+1)^{th}$ iteration at the i^{th} component.

Using this vector, we compute the residual vector, now also depends on i – for every component update (in addition to iteration changes), the residual vector changes.

$$\boldsymbol{r}_i^{(k+1)} = \boldsymbol{b} - A\boldsymbol{x}_i^{(k+1)}$$

or in component form,

$$r_{mi}^{(k+1)} = b_m - \sum_{j=1}^{i-1} a_{mj} x_j^{(k+1)} - \sum_{j=i}^{n} a_{mj} x_j^{(k)} = \left[b_m - \sum_{j=1}^{i-1} a_{mj} x_j^{(k+1)} - \sum_{j=i+1}^{n} a_{mj} x_j^{(k)} \right] - a_{mi} x_i^{(k)}. \tag{0.1}$$

Note the term in the square bracket looks a lot like the Gauss-Seidel iteration. In fact, if we choose $m = i$,

$$\begin{aligned} r_{ii}^{(k+1)} &= \left[b_i - \sum_{j=1}^{i-1} a_{ij} x_j^{(k+1)} - \sum_{j=i+1}^{n} a_{ij} x_j^{(k)} \right] - a_{ii} x_i^{(k)} \\ &= a_{ii} x_i^{(k+1)} - a_{ii} x_i^{(k)} \end{aligned}$$

or

$$x_i^{(k+1)} = x_i^{(k)} + \frac{r_{ii}^{(k+1)}}{a_{ii}},$$

which is a very nice relationship between the individual component update law and the residual vector.

Now, more astoundingly, let's push up all i indices on the RHS of the component form of $r_{mi}^{(k+1)}$ (Eq. 0.1),

$$r_{m,i+1}^{(k+1)} = b_m - \sum_{j=1}^{i} a_{mj} x_j^{(k+1)} - \sum_{j=i}^{n} a_{mj} x_j^{(k)}$$

and setting $m = i$ yields

$$\begin{aligned} r_{i,i+1}^{(k+1)} &= b_i - \left[\sum_{j=1}^{i} a_{ij} x_j^{(k+1)}\right] - \sum_{j=i+1}^{n} a_{ij} x_j^{(k)} \\ &= b_i - \left[\sum_{j=1}^{i-1} a_{ij} x_j^{(k+1)} + a_{ii} x_i^{(k+1)}\right] - \sum_{j=i+1}^{n} a_{ij} x_j^{(k)} \\ &= \left[b_i - \sum_{j=1}^{i-1} a_{ij} x_j^{(k+1)} - \sum_{j=i+1}^{n} a_{ij} x_j^{(k)}\right] - a_{ii} x_i^{(k+1)} \\ &= a_{ii} x_i^{(k+1)} - a_{ii} x_i^{(k+1)} \\ &= 0. \end{aligned}$$

This implies that Gauss-Seidel can be characterized by choosing each component update $x_{i+1}^{(k)}$ in such a way that the i^{th} component of $\boldsymbol{r}_{i+1}^{(k+1)}$ is 0.

However, to choose $x_{i+1}^{(k)}$ so that we kill the residual vector in the i^{th} component doesn't necessarily reduce the norm of $\boldsymbol{r}_{i+1}^{(k)}$. So, we need to choose how far we "walk" carefully, by implementing a parameter ω such that

$$x_i^{(k+1)} = x_i^{(k)} + \omega \frac{r_{ii}^{(k)}}{a_{ii}}.$$

If we choose $0 < \omega < 1$, we are **under-relaxing**, and the technique is called **under-relaxation methods**. If $\omega > 1$, we are **over-relaxing**, which accelerates convergences. In component form,

$$x_i^{(k+1)} = (1-\omega)\, x_i^{(k-1)} + \frac{\omega}{a_{ii}} \left[b_i - \sum_{j=1}^{i-1} a_{ij} x_j^{(k+1)} - \sum_{j=i+1}^{n} a_{ij} x_j^{(k)}\right].$$

In matrix form,

$$(D + \omega L)\, \boldsymbol{x}^{(k+1)} = [(1-\omega)\, D - \omega U]\, \boldsymbol{x}^{(k)} + \omega \boldsymbol{b}$$

which coincides with the geometric approach.

Spectral Radius and its Relationship to Convergence of Iterative Algorithms

We learned that Gauss-Seidel and Jacobi method will converge if the matrix is strictly diagonally dominant. However, in computing, we also care about **the rate** of convergence. We always seek modifications to existing algorithms so that they help converge faster.

The rate of convergence in fact relatives to something called the **spectral radius** of the algorithm operator, such as $T_J = -D^{-1}(L+U)$, $T_{GS} = -(D+L)^{-1} U$ or

$$T_{SOR} = (D + \omega L)^{-1} ((1-\omega)\, D - \omega U).$$

The **spectral radius** of a matrix A is defined as

$$\rho(A) = \max |\lambda|$$

where λ is an eigenvalue of A (if $\lambda = a + bi$ then $|\lambda| = \sqrt{a^2 + b^2}$).

Theorem. *If A is an $n \times n$ matrix, then*

(1) $\|A\|_2 = \left[\rho\left(A^{\mathrm{T}} A\right)\right]^{1/2}$,
(2) $\rho(A) \leq \|A\|$, *for any induced norm* $\|\cdot\|$.

The result we will use is item 1 because finally we have connected the elusive matrix 2-norm to the spectral radius. More examples are shown in the homework because the computation is a straightforward process of finding eigenvalues of the matrix $A^{\mathrm{T}} A$.

Another useful result for spectral radius is that repeated operations on a vector $\boldsymbol{x}$ yields the $\mathbf{0}$ vector.

Theorem. $\rho(A) < 1$ *iff* $\lim_{n\to\infty} A^n \boldsymbol{x} = \mathbf{0}$ *for every* $\boldsymbol{x}$.

Now, we use these results to help us understand the general convergence framework of iterative methods. For any iterative algorithms, we identify its the iteration operator/matrix T as

$$\boldsymbol{x}^{(k+1)} = T\boldsymbol{x}^{(k)} + \boldsymbol{c}$$

where $\boldsymbol{x}^{(0)}$ is an arbitrary initial guess.

Let's apply this formula iteratively. Noting that $\boldsymbol{x}^{(k)} = T\boldsymbol{x}^{(k-1)} + \boldsymbol{c}$, we substitute this into the iteration and obtain

$$\begin{aligned}
\boldsymbol{x}^{(k+1)} &= T\left(T\boldsymbol{x}^{(k-1)} + \boldsymbol{c}\right) + \boldsymbol{c} \\
&= T^2\boldsymbol{x}^{(k-1)} + (T+I)\,\boldsymbol{c} \\
&\;\;\vdots \\
&= T^k\boldsymbol{x}^{(0)} + \left(T^{k-1} + T^{k-2} + \cdots + T + I\right)\boldsymbol{c}
\end{aligned}$$

Note by the previous theorem, if $\rho(T) < 1$, we must have $T^k\boldsymbol{x}^{(0)} \to \boldsymbol{0}$ as $k \to \infty$. Meanwhile, the sum

$$\sum_{i=0}^{k-1} T^i \to \sum_{i=1}^{\infty} T^i = (I-T)^{-1}$$

by the geometric series formula (for matrices). Therefore,

$$\lim_{k\to\infty} \boldsymbol{x}^{(k+1)} = \lim_{k\to\infty} T^k\boldsymbol{x}^{(0)} + \left(\sum_{i=1}^{\infty} T^i\right)\boldsymbol{c} = \boldsymbol{0} + (I-T)^{-1}\,\boldsymbol{c},$$

proving that $\boldsymbol{x}^{(k+1)}$ has an explicit limit,

$$\boldsymbol{x} = (I-T)^{-1}\,\boldsymbol{c} \implies (I-T)\,\boldsymbol{x} = \boldsymbol{c} \implies \boldsymbol{x} = T\boldsymbol{x} + \boldsymbol{c}.$$

This procedure proves that for an iterative scheme to converge, a sufficient condition is $\rho(T) < 1$. In fact, one can show that this condition is also necessary. The proof for necessity is shown on pg. 458 of the 9th edition of the textbook.

Theorem. *If $\|T\| < 1$ for any induced norm $\|\cdot\|$ and $\boldsymbol{c}$ is a given vector, then $\boldsymbol{x}^{(k+1)} = T\boldsymbol{x}^{(k)} + \boldsymbol{c}$ converges, for any $\boldsymbol{x}^{(0)} \in \mathbb{R}^n$, to $\boldsymbol{x} \in \mathbb{R}^n$ such that $\boldsymbol{x} = T\boldsymbol{x} + \boldsymbol{c}$ (also known as a fixed point of the map T). Furthermore, we have the following error bounds*

(1) $\left\|\boldsymbol{x} - \boldsymbol{x}^{(k)}\right\| \le \|T\|^k \left\|\boldsymbol{x}^{(0)} - \boldsymbol{x}\right\|$;

(2) $\left\|\boldsymbol{x} - \boldsymbol{x}^{(k)}\right\| \le \frac{\|T\|^k}{1-\|T\|} \left\|\boldsymbol{x}^{(1)} - \boldsymbol{x}^{(0)}\right\|$.

Since we know there is deep connection between the matrix norm $\|T\|$ and the spectral radius $\rho(T)$, we arrive at the following error estimate

$$\left\|\boldsymbol{x} - \boldsymbol{x}^{(k)}\right\| \approx [\rho(T)]^k \left\|\boldsymbol{x}^{(0)} - \boldsymbol{x}\right\|.$$