Tutorial 2, Jan 30, 2024

Linear Algebra Refresher

- A vector inner product takes $2n - 1$ flops (n multiplications, $n - 1$ additions), which we often shorten to just $2n$
- A matrix-vector product takes $2nm$ flops for an $m \times n$ matrix
- The *condition number* of a function measures how much its output changes for a small change in the input, i.e. how susceptible the output is to noise
  - For the linear equation $\boldsymbol{A}\boldsymbol{x} = \boldsymbol{b}$, it gives a bound of how inaccurate the solution $\boldsymbol{x}$ will be after approximation; smaller condition numbers are better for stability
  - This is a property of the matrix (not the algorithm or the floating-point representation)
  - It is defined as $\kappa(\boldsymbol{A}) = \left|\frac{\lambda_{\text{max}}}{\lambda_{\text{min}}}\right|$
  - Note $\kappa(\boldsymbol{A}^T\boldsymbol{A}) = (\kappa(\boldsymbol{A}))^2$, so squaring a matrix is very bad for stability
- Once we have solved $\boldsymbol{A}\boldsymbol{x} = \boldsymbol{b}$, we can use it to also solve an updated problem where we perturb $\boldsymbol{A}$ by a small, low-rank $\boldsymbol{UCV}$
  - A low-rank perturbation means that the change in certain dimensions is much less than that in other dimensions
  - Assume we know $\boldsymbol{A}^{-1}$ or that it is easy to compute
  - Then $(\boldsymbol{A} + \boldsymbol{UCV})^{-1} = \boldsymbol{A}^{-1} - \boldsymbol{A}^{-1}\boldsymbol{U}(\boldsymbol{C}^{-1} + \boldsymbol{V}\boldsymbol{A}^{-1}\boldsymbol{U})^{-1}\boldsymbol{V}\boldsymbol{A}^{-1}$ (*matrix inversion lemma*)
  - All we need to do is to calculate $\boldsymbol{C}^{-1}$, which is very efficient for a low-rank $\boldsymbol{C}$
  - A similar formula exists for the determinant