Supplemental Exercises: Unit 1
Scientific Computing with Case Studies
Dianne P. O'Leary
SIAM Press, 2009

---

1. Suppose we have a computer that uses (single-precision) IEEE Standard floating point arithmetic: 24 digits to represent the mantissa, and exponents in the range $[-126, 127]$.

(a) Consider evaluating the expression `c = a * b` on this machine. Give a machine-representable (finite) value for `a` and a machine-representable (finite) value for `b` for which the computed value `c=INF` because of overflow.

(b) What is the distance between 2^{20} and the next larger floating point number?

---

2. If we type `cos(pi/2)` in MATLAB, the computed answer is `6.1232e-17`. Why doesn't MATLAB return the correct value, 0?

---

3. Let $\boldsymbol{A}$ be an $m \times n$ matrix, and define

$$\|\boldsymbol{A}\|_\infty = \max_i \sum_{j=1}^{n} |a_{ij}| \,.$$

Write an efficient column-oriented algorithm to compute $\|\boldsymbol{A}\|_\infty$. (Don't use any MATLAB function calls except `abs` and `max`.)

Note: In practice, for maximal efficiency, we would just say `normA = norm(A,inf)` to access a column-oriented algorithm.
