Homework 0: Math Review and Julia Warm-up

CS 205A: Mathematical Methods for Robotics, Vision, and Graphics (Winter 2018)
Stanford University
Due Thursday, Jan 18, before midnight (via gradescope)

> This homework is designed to help students in CS 205A review the basic mathematics that we will be using for the rest of the course. Make sure you are 100% confident in your solution to each problem: We will use these techniques repeatedly for the rest of the course! Depending on your background, you may not have seen some of these topics; make ample use of office hours and section to help fill any gaps in your knowledge or understanding.

Textbook problems: 1.5, 1.7, 1.9, 1.10, 1.15 (15 points each)

Julia Programming Assignment (25 points): This question considers *Gaussian probability distributions* in one and higher dimensions. Consider the following one-dimensional Gaussian probability distribution, also known as the Normal distribution or bell curve distribution:

$$G(x|\mu,\sigma) = \frac{1}{Z}\exp\left[-\frac{1}{2}\frac{1}{\sigma^2}(x-\mu)^2\right] \quad \text{where } Z = \sqrt{\frac{\pi}{\frac{1}{2}\frac{1}{\sigma^2}}}$$

Here, the function $G : \mathbb{R} \mapsto \mathbb{R}$ takes as input a scalar x, and produces as output a scalar. The particular bell shape of G is determined by two parameters: the mean $\mu \in \mathbb{R}$ and the variance $\sigma^2 \in \mathbb{R}$.

(1) (5 points) In a Julia notebook, numerically verify the following identity:

$$\int_{\mathbb{R}} G(x)dx = 1$$

Choose a few different one-dimensional Gaussian functions (by choosing different mean and variance values), plot them, and verify the above identity for each Gaussian function. *Hint: You don't actually need to numerically integrate over all the real numbers. A Gaussian function is close to zero almost everywhere (except near its mean), so just choose a finite integration domain you think is reasonable.*

(2) (5 points) Now consider the more general definition for the Gaussian probability distribution in n dimensions:

$$G(\vec{x}|\vec{\mu},\Sigma) = \frac{1}{Z}\exp\left[-\frac{1}{2}(\vec{x}-\vec{\mu})^T\Sigma^{-1}(\vec{x}-\vec{\mu})\right] \quad \text{where } Z = \sqrt{\frac{\pi^n}{\det(\frac{1}{2}\Sigma^{-1})}}$$

Here, $G : \mathbb{R}^n \mapsto \mathbb{R}$ takes as input a vector $\vec{x} \in \mathbb{R}^n$, and produces as output a scalar. The mean is now a vector $\vec{\mu} \in \mathbb{R}^n$, and the peaky shape and skew shape of the distribution is given by the covariance matrix $\Sigma \in \mathbb{R}^n \times \mathbb{R}^n$. Choose two different two-dimensional Gaussian functions, plot them, and numerically verify the following identity:

$$\int_{\mathbb{R}^n} G(\vec{x})d\vec{x} = 1$$

Hint: Again, just choose a finite integration domain you think is reasonable.

(3) (15 points) Finally, consider a smooth (infinitely differentiable) function $f(\vec{x}) : \mathbb{R}^n \mapsto \mathbb{R}$ that achieves a local maximum at $\vec{x}^\star$. Find a Gaussian function $G^\star$ that approximates f in a small neighborhood around $\vec{x}^\star$. The Gaussian function you find must have the same value as f at $\vec{x}^\star$, must have the same first derivative as f at $\vec{x}^\star$, and must integrate to 1. You can assume that $f(\vec{x}) > 0$ in the neighborhood around $\vec{x}^\star$. Numerically verify the correctness of your reasoning in two dimensions, for the following function:

$$f(\vec{x}) = \frac{1}{4}\left[\sin\left(\frac{1}{2}\pi\vec{x}_1\right) + \cos\left(\pi\vec{x}_2 - \pi\right)\right]$$

Construct your Gaussian approximation around the local maximum $\vec{x}^\star = [1\ 1]^T$. Plot your your results in a way that convincingly shows that your Gaussian correctly approximates f in the neighborhood around $\vec{x}^\star$, and verify that your Gaussian integrates to 1. *Hint: In two and higher dimensions, there is an entire family of Gaussian functions that satisfy the above constraints. Feel free to choose a Gaussian from this family in a way that makes the problem easier. For example, if you assume that your covariance matrix is a scaled version of the identity matrix, the problem has a particularly elegant solution.*

To simplify submission to GradeScope with your other written homework, export a PDF of a clearly documented Julia Notebook that shows your work.