

Pr. 1.

- (a) One use of the “**vec trick**” $\text{vec}(\mathbf{A}\mathbf{X}\mathbf{B}^\top) = (\mathbf{B} \otimes \mathbf{A})\text{vec}(\mathbf{X})$ is computing a 2D **discrete Fourier transform** (DFT) of a 2D signal. The (1D) DFT of a vector $\mathbf{x} \in \mathbb{C}^N$ is the vector $\mathbf{f}_x \in \mathbb{C}^N$ with entries

$$[\mathbf{f}_x]_k = \sum_{n=1}^N x_n \exp\left(\frac{-2\pi i(k-1)(n-1)}{N}\right), \quad k = 1, \dots, N.$$

(Here we follow linear algebra (and Julia) where the first array index is 1, whereas DSP books usually use $0, \dots, N-1$.) We can represent the above computation in matrix-vector form as $\mathbf{f}_x = \mathbf{F}_N \mathbf{x}$, where \mathbf{F}_N is the $N \times N$ DFT matrix with entries

$$[\mathbf{F}_N]_{k,n} = \exp\left(\frac{-2\pi i(k-1)(n-1)}{N}\right), \quad k, n = 1, \dots, N.$$

We can generate \mathbf{F}_N in Julia as follows.

```
# N × N DFT matrix
using FFTW: fft
using LinearAlgebra: I
F = fft(I(N), 1)
```

One can verify that $\mathbf{F}_N' \mathbf{F}_N = N \mathbf{I}_N$, so the (1D) **inverse DFT** of \mathbf{f}_x can be computed as $\mathbf{x} = (1/N) \mathbf{F}_N' \mathbf{f}_x$.

We compute the **2D DFT**, call it \mathbf{S}_X , of the $M \times N$ matrix \mathbf{X} by computing the 1D DFT of each column of \mathbf{X} followed by the 1D DFT of each row of the result (or vice versa).

Explain why the following expression computes the 2D DFT of \mathbf{X} :

$$\mathbf{S}_X = \mathbf{F}_M \mathbf{X} \mathbf{F}_N^\top.$$

- (b) Write $\text{vec}(\mathbf{S}_X)$ as the product of a matrix and $\text{vec}(\mathbf{X})$.
- (c) Show that the following expression computes the **2D inverse DFT** of \mathbf{S}_X :

$$\mathbf{X} = \frac{1}{MN} \mathbf{F}_M' \mathbf{S}_X \overline{\mathbf{F}_N},$$

where $\overline{\mathbf{Y}}$ denotes (elementwise) complex conjugate of matrix \mathbf{Y} .

Hint: Use the fact that $\mathbf{F}_N' \mathbf{F}_N = \mathbf{F}_N \mathbf{F}_N' = N \mathbf{I}_N$.

- (d) Write $\text{vec}(\mathbf{X})$ as the product of a matrix and $\text{vec}(\mathbf{S}_X)$.

Pr. 2.

In the ordinary **least-squares** problem we found the (minimum norm) $\hat{\mathbf{x}}$ that minimized $\|\mathbf{r}\|_2$ where $\mathbf{r} = \mathbf{b} - \mathbf{A}\mathbf{x}$ is the residual. We saw that the optimal \mathbf{x} can be expressed entirely in terms of \mathbf{b} and an SVD of \mathbf{A} . Let \mathbf{A} be an $M \times N$ matrix so that $\mathbf{x} \in \mathbb{F}^N$ and $\mathbf{b} \in \mathbb{F}^M$ and $\mathbf{r} \in \mathbb{F}^M$.

In applications with **heteroscedastic** measurement errors, we prefer to minimize $\|\mathbf{W}\mathbf{r}\|_2$, *i.e.*, a weighted squared error, where \mathbf{W} is a diagonal matrix having diagonal entries $w_1, \dots, w_M \geq 0$. Determine the optimal \mathbf{x} for this **weighted least-squares** problem. (Again, you should express the answer in terms of \mathbf{b} , \mathbf{W} , and an SVD of an appropriate matrix.) Your answer must not have any pseudoinverse in it! (It may have an inverse as long as you are sure that the matrix is invertible and trivial to invert.) You may assume that `W*A == zeros(size(W*A))` is false in Julia.

Pr. 3.

- (a) Find (by hand) all solutions of the linear **least-squares** problem $\arg \min_{\mathbf{x} \in \mathbb{R}^2} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2$ when $\mathbf{A} = \begin{bmatrix} 1 & 2 \\ 1 & 2 \end{bmatrix}$ and $\mathbf{b} = \begin{bmatrix} 1 \\ 2 \end{bmatrix}$.
- (b) Describe briefly how you would modify your solution for this variation of the problem: $\arg \min_{\mathbf{x} \in \mathbb{C}^2} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2$.

Pr. 4.

Recall that the **least-squares** problem

$$\hat{\mathbf{x}} = \arg \min_{\mathbf{x}} \frac{1}{2} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2^2,$$

has the solution $\hat{\mathbf{x}} = \mathbf{A}^+ \mathbf{b}$ where \mathbf{A}^+ denotes the **pseudoinverse**. When \mathbf{A} is large, it can be expensive to compute $\hat{\mathbf{x}}$ using the pseudoinverse of \mathbf{A} . In such settings, the **gradient descent** (GD) iteration given by

$$\mathbf{x}_{k+1} = \mathbf{x}_k - \alpha \mathbf{A}'(\mathbf{A}\mathbf{x}_k - \mathbf{b}), \quad (1)$$

will converge to a minimizer of $\|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2$ as iteration $k \rightarrow \infty$ when $0 < \alpha < 2/\sigma_1^2(\mathbf{A})$. Note that the iteration has a fixed point, *i.e.*, $\mathbf{x}_{k+1} = \mathbf{x}_k$ if

$$\mathbf{A}'(\mathbf{A}\mathbf{x}_k - \mathbf{b}) = \mathbf{0},$$

which are exactly the normal equations. So any fixed point minimizes the least-squares cost function.

- (a) Write a function called `lsqd` that implements the above least-squares gradient descent algorithm.

In Julia, your file should be named `lsqd.jl` and should contain the following function:

```
"""
    x = lsqd(A, b ; alpha=0, x0=zeros(size(A,2)), nIters::Int=200)

Perform gradient descent to solve the least-squares problem:
``\argmin_x 0.5 || A x - b ||_2``

# In:
- `A` `m × n` matrix
- `b` vector of length `m`

# Option:
- `alpha` step size to use, and must satisfy ``0 < α < 2 / σ₁(A)²``
  to guarantee convergence,
  where ``σ₁(A)`` is the first (largest) singular value.
  Ch.6 explains a default value for `alpha`
- `x0` is the initial starting vector (of length `n`) to use.
  Its default value is all zeros for simplicity.
- `nIters` is the number of iterations to perform (default 200)

Out:
- `x` vector of length `n` containing the approximate LS solution
"""
function lsqd(A, b ; alpha::Real=0, x0=zeros(size(A,2)), nIters::Int=200)
```

Email your solution as an attachment to eeecs551@autograder.eecs.umich.edu.

The function specification above uses a powerful feature of Julia where functions can have optional arguments with specified default values. If you simply call `lsqd(A,b)` then `alpha`, `x0` and `nIters` will all have their default values. But if you call, say, `lsqd(A, b, nIters=5, alpha=7)` then it will use the specified values for `nIters` and `alpha` and the default for `x0`. Note that these named optional arguments can appear in any order. This approach is very convenient for functions with multiple arguments.

- (b) Use your function (after it passes) to generate a plot of $\log_{10}(\|\mathbf{x}_k - \hat{\mathbf{x}}\|)$ versus $k = 0, 1, \dots, 200$ using $\alpha = 1/\sigma_1^2(\mathbf{A})$ for \mathbf{A} and \mathbf{b} generated as follows.

```
using Random: seed!
m, n = 100, 50; sigma = 0.1
seed!(0) # seed random number generator
A = randn(m, n); xtrue = rand(n); noise = randn(m)
b = A * xtrue + sigma * noise # b and xhat change with σ
```




Figure 1: Photometric stereo example. Left (a): 2D images of a common scene under different lighting conditions. Right (b): 3D surface reconstruction computed from the input images.

Repeat for $\sigma = 0.5, 1, 2$ and submit one plot with all four curves on it. Does $\|\mathbf{x}_k - \hat{\mathbf{x}}\|$ decrease monotonically with k in the plots? Note that $\sigma = \text{sigma}$ here is a noise standard deviation unrelated to singular values.

Pr. 5.

For $\mathbf{A} \in \mathbb{F}^{M \times N}$, $\mathbf{b} \in \mathbb{F}^M$ and $\mathbf{x} \in \mathbb{F}^N$, show that $(\mathbf{I} - \mathbf{A}^+ \mathbf{A})\mathbf{x}$ and $\mathbf{A}^+ \mathbf{b}$ are **orthogonal** vectors.

Hint. Use a **compact SVD**.

Pr. 6.

- (a) For $\delta > 0$, determine the solution of the **regularized LS** problem

$$\hat{\mathbf{x}}(\delta) = \arg \min_{\mathbf{x}} \frac{1}{2} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2^2 + \delta^2 \frac{1}{2} \|\mathbf{x}\|_2^2.$$

Express the answer in terms of the SVD of an appropriate matrix.

Use the fact that $\mathbf{B}^+ = (\mathbf{B}'\mathbf{B})^{-1}\mathbf{B}'$ when $\mathbf{B}'\mathbf{B}$ is invertible to simplify the expression.

Hint: Rewrite the cost function so it looks like a usual least-squares problem.

- (b) What does $\hat{\mathbf{x}}(\delta)$ tend to as $\delta \rightarrow \infty$? Does this answer make sense?
- (c) Write (by hand, not code) an iteration based on the **gradient descent** (GD) method such that the iterates converge to the minimizer $\hat{\mathbf{x}}(\delta)$.
- (d) Determine a condition on the step size μ that guarantees convergence of your GD method. Express the condition in terms of the original problem quantities: \mathbf{A} , \mathbf{b} , and δ .

Pr. 7.

(Photometric stereo: compute normals)

In this problem you implement another tool needed for a computer vision method called **photometric stereo**, allowing us (eventually) to reconstruct a 3D object's surface from 2D images of it under different lighting conditions. Figure 1 gives a preview at what you will be able to do after completing all the pieces. Specifically, this version is called “calibrated, far-field photometric stereo.”

Suppose we are in a dark room with an object on a dark table, a camera fixed above it, and a moveable light source. We model the object surface as $z = f(x, y)$, where (x, y) denotes coordinates on the table and z denotes height above the table. Assume that a $m \times n$ sized image $I(x, y)$ is a representation of $f(x, y)$ for each (x, y) tuple. (Given one light source, $f(x, y)$ is the z coordinate of the position where a ray of light hits the surface at (x, y, z) .)

As seen in Figure 1a, the pixel intensity $I(x, y)$ indicates how much light reflects off the surface $f(x, y)$. If our object is diffuse (also called matte or **Lambertian reflectance**), one can derive the relationship

$$I(x, y) = \alpha(x, y) (\boldsymbol{\ell}^T \mathbf{n}(x, y)), \quad (2)$$

where $\boldsymbol{\ell} \in \mathbb{R}^3$ is a unit vector describing the orientation of the incident light rays on the surface, $\mathbf{n}(x, y) \in \mathbb{R}^3$ is the unit-norm surface normal vector of f at $(x, y, f(x, y))$, and $\alpha(x, y) > 0$ is a scaling constant called the **surface albedo**.

Now suppose that we take d images I_1, \dots, I_d of our object, with lighting directions $\ell_1, \dots, \ell_d \in \mathbb{R}^3$. We are going to work on one pixel at a time. For any pixel location (x, y) , we can stack (2) into an **overdetermined** system of equations:

$$\underbrace{\begin{bmatrix} I_1(x, y) \\ \vdots \\ I_d(x, y) \end{bmatrix}}_{\triangleq \mathbf{g}(x, y) \in \mathbb{R}^d} \approx \underbrace{\begin{bmatrix} \ell_1 & \dots & \ell_d \end{bmatrix}^T}_{\triangleq \mathbf{L}^T} \underbrace{(\alpha(x, y) \mathbf{n}(x, y))}_{\triangleq \boldsymbol{\rho}(x, y) \in \mathbb{R}^3}. \quad (3)$$

We could solve (3) for $\boldsymbol{\rho}(x, y)$ for each (x, y) when $d = 3$, but, in practice, when there is noise and our assumptions do not hold exactly, a more robust approach is to take $d > 3$ images in the **least-squares** problem

$$\hat{\boldsymbol{\rho}}(x, y) = \arg \min_{\mathbf{r} \in \mathbb{R}^3} \|\mathbf{g}(x, y) - \mathbf{L}^T \mathbf{r}\|_2^2. \quad (4)$$

We then approximate the surface normal \mathbf{n} by the following normalized estimate:

$$\hat{\mathbf{n}}(x, y) \triangleq \frac{\hat{\boldsymbol{\rho}}(x, y)}{\|\hat{\boldsymbol{\rho}}(x, y)\|_2}. \quad (5)$$

Your task for this problem is to write a function called `compute_normals` that computes the unit-norm surface normal vectors for each pixel in a scene, given $\{I_i(x, y)\}$ and $\{\ell_i\}$, by solving (4). and applying (5).

Hint: Julia's `normalize` and `mapslices` functions are useful.

In Julia, your file should be named `compute_normals.jl` and should contain the following function:

```
"""
    normals = compute_normals(data, L)

# In:
- `data` `M × N × d` array whose `d` slices contain `M × N` images
  of a common scene under different lighting conditions

- `L` `3 × d` matrix whose columns are the lighting direction vectors
  for the images in `data`, with `d ≥ 3`

# Out:
- `normals` `M × N × 3` array containing the unit-norm surface normal vectors
  for each pixel in the scene
"""
function compute_normals(data, L)
```

Email your solution as an attachment to eeecs551@autograder.eecs.umich.edu.

Note: The surface normals at each pixel are independent, so you can compute them all simultaneously by stacking the solutions to (4) for each pixel into a single matrix expression, using `reshape` to help. This solution can be expressed elegantly without using loops! (But it is OK to use loops if you prefer.)

Hint. You can test your code by applying it to synthetic data like the following.

```
include(ENV["hw551test"] * "compute_normals.jl") # use your own path
using Plots: plot
using Colors: RGB
using MIRTjim: jim

d = 8 # number of test images
phi = range(-1, 1, d) * pi / 3 # light source angles
L = [sin.(phi) 0 * phi cos.(phi)]' # light unit vectors (3, d)
m, n = 40, 30 # image size
r0 = 0.8 # cylinder radius
```

```

x = range(-1, 1, m); y = range(-1, 1, n)
afun(x,y) = 0.9 * (abs(x) < 0.5*r0) # albedo alpha(x,y); avoids occlusion
zfun = x -> (abs(x) < r0) ? sqrt(r0^2 - x^2) : 0 # cylinder surface
normal(x,y) = (abs(x) < r0) ? [x/r0, 0, zfun(x)/r0] : zeros(3) # surface normal
ntrue = [normal(x,y) * (afun(x,y) > 0) for x in x, y in y]
ntrue = cat([map(x -> x[i], ntrue) for i in 1:3]..., dims=3) # m x n x 3

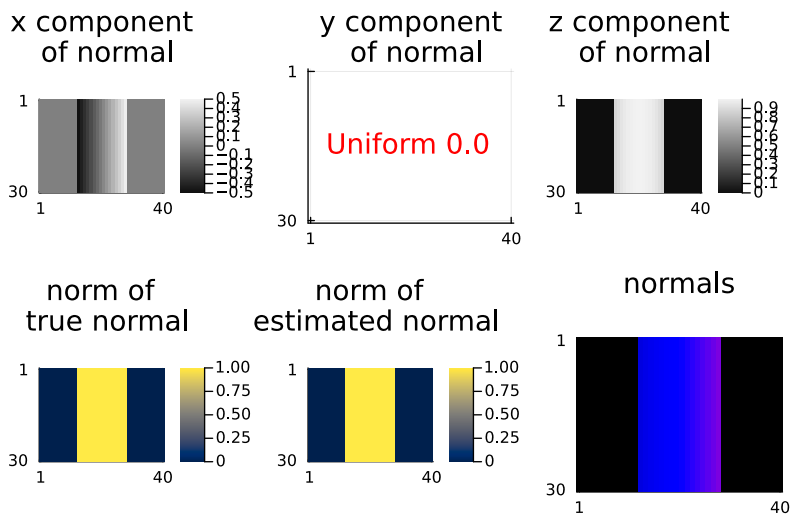
data = zeros(Float32, m, n, d)
for l in 1:d
    data[:, :, l] .= [L[:, l] * normal(x,y) * afun(x,y) for x in x, y in y]
end
p1 = jim(data, "$d images for different light positions")
nhath = compute_normals(data, L)
nhath[isnan.(nhath)] .= 0
p2 = jim(cat(ntrue, nhath, nhath-ntrue, dims=4),
    "true and estimated normals and error")

norm_tru = sum(abs2, ntrue, dims=3)
norm_hat = sum(abs2, nhath, dims=3)

p3 = plot(
    jim(nhath[:, :, 1], "x component\n of normal"; clim=(-0.5, 0.5)),
    jim(nhath[:, :, 2], "y component\n of normal"),
    jim(nhath[:, :, 3], "z component\n of normal"),
    jim(norm_tru, "norm of\n true normal"; color=:cividis),
    jim(norm_hat, "norm of\n estimated normal"; color=:cividis),
    jim(RGB.(ntrue[:, :, 1], ntrue[:, :, 2], ntrue[:, :, 3]), "normals")
) # savefig("hp073.pdf")

```

This code simulates imaging part of the surface of a cylinder with different lighting positions and then applies `compute_normals` to the simulated data. If your code is working properly, then the three estimated normals should look like those in the following figure.



Pr. 8.**(Photometric stereo continued)**

(This problem statement looks long, but it is actually quite easy because it uses tools you developed previously. You need not write any new code for it.)

For $f : \mathbb{R}^2 \mapsto \mathbb{R}$, from vector calculus, we can express the **surface normal vector** of f at (x, y) as

$$\mathbf{n}(x, y) = \frac{1}{\sqrt{1 + \left(\frac{\partial}{\partial x} f(x, y)\right)^2 + \left(\frac{\partial}{\partial y} f(x, y)\right)^2}} \begin{bmatrix} -\frac{\partial}{\partial x} f(x, y) \\ -\frac{\partial}{\partial y} f(x, y) \\ 1 \end{bmatrix} \triangleq \begin{bmatrix} n_1(x, y) \\ n_2(x, y) \\ n_3(x, y) \end{bmatrix}, \quad (6)$$

where $\frac{\partial}{\partial x} f$ and $\frac{\partial}{\partial y} f$ denote the partial derivatives of depth $f(x, y)$ with respect to x and y , respectively.

From (6), we can compute the partial derivatives as

$$\frac{\partial f(x, y)}{\partial x} = -\frac{n_1(x, y)}{n_3(x, y)}, \quad \frac{\partial f(x, y)}{\partial y} = -\frac{n_2(x, y)}{n_3(x, y)}. \quad (7)$$

In a previous HW you constructed a matrix \mathbf{A} satisfying $\begin{bmatrix} \mathbf{dfdx} \\ \mathbf{dfdy} \end{bmatrix} = \mathbf{A} \mathbf{fxy}$, where \mathbf{dfdx} and \mathbf{dfdy} denote the vectorized approximations of $\frac{\partial f(x, y)}{\partial x}$ and $\frac{\partial f(x, y)}{\partial y}$ and \mathbf{fxy} denotes the vectorized approximation of $f(x, y)$. Using (7), we can compute \mathbf{dfdx} and \mathbf{dfdy} from our normal vectors, and, using our \mathbf{A} matrix, we can obtain the surface corresponding to our normal vectors by solving the following **least-squares** problem:

$$\mathbf{fxy} = \arg \min_{\mathbf{f} \in \mathbb{R}^{mn}} \left\| \begin{bmatrix} \mathbf{dfdx} \\ \mathbf{dfdy} \end{bmatrix} - \mathbf{A} \mathbf{f} \right\|_2^2.$$

Download the `photometric_stereo_xy` notebook from the `hw05` directory on Canvas, and copy your previous `compute_normals.jl` and `first_diffs_2d_matrix.jl` solution files into the same directory. Now use Julia to run that Jupyter notebook. If you have installed the necessary packages mentioned in the notebook, and if you have working versions of the `compute_normals` and `first_diffs_2d_matrix` functions, then the notebook will run properly and generate a surface view of the object.

After $f(x, y)$ is estimated, one can separately generate a stereolithography file (consisting of a collection of tessellated triangles) that can be rendered on a 3D display or printed by a 3D printer. Figure 2 depicts an actual 3D-print made from a solution to this problem generated with a Cube 3 printer.

Submit a screenshot of the final surface plot (with your username in it) produced by your Jupyter notebook.

It should look something like Fig. 2 but in color.



Figure 2: 3D printed reconstruction of the surface you will reconstruct in Problem 8. The 3D printing quality was intentionally set coarse so you can see how the printer constructed the shape from its level curves.

Pr. 9.**Hand-written digit classification using nearest subspace (discussion task)**

This task illustrates how to use **subspace-based classification** with image features for classifying handwritten digits. For this machine learning problem, we focus on just the two digits “4” and “9,” although the principles generalize to all digits. Previously we focused on “0” and “1” because the previous classification methods were less powerful. With the subspace approach, you can now tackle the harder problem of distinguishing “4” from “9.”

Download the `task-2-classify-subspace.ipynb` jupyter notebook file from Canvas under this week’s discussion folder and follow all instructions to complete the task. You may work individually, but we recommend that you work in pairs or groups of three.

When you are finished, upload your solutions to gradescope. Note that the submission for the task is separate from the rest of the homework because the task allows you to submit as a group. Only upload one submission per group! Whoever uploads the group submission should add all group members in gradescope, using the “View or edit group” option on the right-hand sidebar after uploading a PDF and matching pages. Make sure to add all group members, because this is how they will receive credit.

Non-graded problem(s) below

(Solutions will be provided for self check; do not submit to gradescope.)

Pr. 10.

Let \mathcal{S} and \mathcal{T} denote subspaces of a vector space \mathcal{V} . Prove or disprove the following.

- (a) If \mathcal{S} is a subset of \mathcal{T} or vice versa, then the **union of subspaces** $\mathcal{S} \cup \mathcal{T}$ is a **subspace** of \mathcal{V} .
- (b) If the **union of subspaces** $\mathcal{S} \cup \mathcal{T}$ is a **subspace** of \mathcal{V} , then \mathcal{S} is a subset of \mathcal{T} or vice versa.

Pr. 11.

Suppose $\mathbf{A} \in \mathbb{F}^{M \times N}$ with $M \geq N$ has **full column rank**, i.e., $\text{rank}(\mathbf{A}) = N$.

Show, using an SVD of \mathbf{A} , that $\mathbf{A}^+ = (\mathbf{A}'\mathbf{A})^{-1}\mathbf{A}'$. When \mathbf{A} is known to be full column rank, one could use this direct formula to compute the **pseudoinverse** instead of employing an SVD.

Pr. 12.

Consider the problem of finding the **minimum 2-norm solution** of the linear **least-squares** problem

$$\hat{\mathbf{x}} = \arg \min_{\mathbf{x}} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2 \text{ when } \mathbf{A} = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} \text{ and } \mathbf{b} = \begin{bmatrix} 1 \\ 1 \end{bmatrix}. \text{ The solution is: } \hat{\mathbf{x}} = \mathbf{A}^+\mathbf{b} = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} 1 \\ 1 \end{bmatrix} = \begin{bmatrix} 1 \\ 0 \end{bmatrix}.$$

- (a) Consider a perturbation $\mathbf{E}_1 = \begin{bmatrix} 0 & \delta \\ 0 & 0 \end{bmatrix}$ of \mathbf{A} , where δ is a small positive number. Solve the perturbed version of the above problem: $\mathbf{z}^* = \arg \min_{\mathbf{z}} \|\mathbf{A}_1\mathbf{z} - \mathbf{b}\|_2$, where $\mathbf{A}_1 = \mathbf{A} + \mathbf{E}_1$.

What happens to $\|\hat{\mathbf{x}} - \mathbf{z}^*\|_2$ as δ approaches 0?

- (b) Now consider the perturbation $\mathbf{E}_2 = \begin{bmatrix} 0 & 0 \\ 0 & \delta \end{bmatrix}$ where again δ is a small positive number. Solve the perturbed problem $\mathbf{z}^* = \arg \min_{\mathbf{z}} \|\mathbf{A}_2\mathbf{z} - \mathbf{b}\|_2$ where $\mathbf{A}_2 = \mathbf{A} + \mathbf{E}_2$.

What happens to $\|\hat{\mathbf{x}} - \mathbf{z}^*\|_2$ here as $\delta \rightarrow 0$?

Pr. 13.

Prove any following *necessary* and *sufficient* conditions for the **range** of a matrix product \mathbf{AB} to equal the range of $\mathbf{A} \in \mathbb{F}^{M \times N}$ for $\mathbf{B} \in \mathbb{F}^{N \times K}$.

- (a) $\mathcal{R}(\mathbf{AB}) = \mathcal{R}(\mathbf{A}) \iff \mathcal{R}(\mathbf{B}) + \mathcal{N}(\mathbf{A}) = \mathbb{F}^N$.
- (b) $\mathcal{R}(\mathbf{AB}) = \mathcal{R}(\mathbf{A}) \iff \mathbf{A} = \mathbf{ABC}$ for some $\mathbf{C} \in \mathbb{F}^{K \times N}$.
- (c) Let \mathbf{A} have compact SVD $\mathbf{A} = \mathbf{U}_r \mathbf{\Sigma}_r \mathbf{V}_r'$. Then $\mathcal{R}(\mathbf{AB}) = \mathcal{R}(\mathbf{A}) \iff \mathcal{R}(\mathbf{V}_r' \mathbf{B}) = \mathbb{F}^r$.