Problem Set 1: Deblurring and MRI Aliasing

CS-491 Machine Learning for Medical Imaging
Fall 2026

Assigned: Week 2 **Due:** start of class, Week 4

Instructions. Submit a single PDF containing your derivations, figures, and discussion, together with the script or notebook that generated every figure. Any programming language is acceptable. You may use libraries for functions unless specified otherwise.

Problem 1 (Microscopy Deblurring)**.** In this problem, you will simulate the forward observation and reconstruction of a microscopy inverse imaging problem.

Consider a grayscale $N_1 \times N_2$ digital microscopy image $\boldsymbol{x} \in \mathbb{R}^N$ with $N = N_1 N_2$ pixels. To correct for lens blurring, we take an image of a small calibration dot, obtaining the point-spread-function (PSF) of the lens system as a Gaussian-kernel $\boldsymbol{h} \in \mathbb{R}^{N_h}$ with std-dev σ_h-pixels.

(a) Load a grayscale microscopy "ground-truth" image and scale it to a maximum value of 1. Write a function to generate a 2D Gaussian filter with std-dev σ_h-pixels. Make the kernel at least $5\sigma_h$ wide and ensure that the filter sums to 1 to be mean-preserving.

(b) Generate observation $\boldsymbol{y} = \boldsymbol{h} \circledast \boldsymbol{x} + \boldsymbol{\nu} = \boldsymbol{H}\boldsymbol{x} + \boldsymbol{\nu}$, where $\boldsymbol{\nu} \sim \mathcal{N}(0,\, \sigma_n^2 \boldsymbol{I}_N)$, $\sigma_n = \sqrt{N}/2^{16}$, and $\sigma_h = \sqrt{N}/2^7$.

(c) Derive the recovery expression for reconstruction $\boldsymbol{x}^\diamond$ from measurement image $\boldsymbol{y}$ via the pseudo-inverse of the observation matrix. Show both the recovery expression and what the Fourier domain representation of $\boldsymbol{x}^\diamond$ looks like in terms of $\boldsymbol{h}$, $\boldsymbol{x}$, and $\boldsymbol{\nu}$.

(d) Derive an expression for the noise-covariance of the reconstruction and explain its characteristics qualitatively.

(e) Write a function `fpinv` which performs pseudo-inverse recovery of $\boldsymbol{x}$ from $\boldsymbol{y}$ given $\boldsymbol{h}$ using Fourier domain processing. Show the recovered image $\boldsymbol{x}^\diamond$ and its NRMSE ($\|\boldsymbol{x}^\diamond - \boldsymbol{x}\|_2 / \|\boldsymbol{x}\|_2$). How does it look? Is it expected?

(f) Your recovery equation involved a division with values close to zero. To stabilize your results, add a small value λ to the denominator. Sweep values of λ uniformly on a log x-axis scale and generate a graph of NRMSE vs. λ. Display the results (image and NRMSE) for the best performing λ value.

(optional): To obtain a color image, red green and blue color-filters are placed in-between the lens and the sensor to obtain three separate images of the same scene in rapid succession: $\boldsymbol{y}_R, \boldsymbol{y}_G, \boldsymbol{y}_B \in \mathbb{R}^N$*. Stacking these images together into a single vector,* $\boldsymbol{y} = \begin{bmatrix} \boldsymbol{y}_R^\top & \boldsymbol{y}_G^\top & \boldsymbol{y}_B^\top \end{bmatrix}^\top \in \mathbb{R}^{3N}$*, we may express the observation equation in block-matrix form as,*

$$\boldsymbol{y} = \begin{bmatrix} \boldsymbol{H} & & \\ & \boldsymbol{H} & \\ & & \boldsymbol{H} \end{bmatrix} \begin{bmatrix} \boldsymbol{x}_R \\ \boldsymbol{x}_G \\ \boldsymbol{x}_B \end{bmatrix} + \boldsymbol{\nu}, \quad \boldsymbol{\nu} \sim \mathcal{N}(0,\, \sigma_n^2 \boldsymbol{I}_{3N}).$$

(a) Derive the pseudo-inverse recovery expression for $\boldsymbol{x}^\diamond \in \mathbb{R}^{3N}$ *in block-matrix form.*

(b) Repeat your work for grayscale recovery above for an RGB image.

Problem 2 (MRI Aliasing)**.** In this problem we analyze the forward model of MRI and the standard adjoint ("zero-filled") reconstruction. Download the MoDL brain dataset (`dataset.hdf5`, from the download link at `https://github.com/hkaggarwal/modl`), which contains reference brain images ($\boldsymbol{x} \in \mathbb{C}^N$) (with $N = N_1 \times N_2$-pixels) and multicoil sensitivity maps ($\boldsymbol{s} \in \mathbb{C}^{NC}$).

For your Fourier transforms, pre and post-compose your `fft`'s with shifts to correctly center low-frequency k-space at the center of the image, i.e. write a function

```
ft(x) = ifftshift(fft(fftshift(x)))
```

and a corresponding inverse function `ift`. Be careful to note the normalization standards of whatever fft-library you use, as noise-level is referenced w.r.t. a normalized FFT.

The forward model for "fully-sampled" single-coil MRI can be written as $\boldsymbol{y} = \boldsymbol{F}_N \boldsymbol{x} + \boldsymbol{\nu}$ where $\boldsymbol{F}_N$ is the N-pixel unitary 2D DFT matrix and $\boldsymbol{\nu} \sim \mathcal{C}N(0,\, \sigma_n^2 \boldsymbol{I})$ is complex Gaussian i.i.d. noise.

(a) Write a function to load a single slice from a complex brain volume, $\boldsymbol{x} \in \mathbb{C}^N$. Display the magnitude image $|\boldsymbol{x}|$ and the log k-space (log 2D Fourier domain) image side-by-side with colorbars for each image.

(b) To accelerate acquisition, we can skip columns of k-space, effectively making our forward model $\boldsymbol{y} = \boldsymbol{I}_\Omega \boldsymbol{F}_N \boldsymbol{x} + \boldsymbol{\nu}$ where $\boldsymbol{I}_\Omega \in \{0, 1\}^{|\Omega| \times N}$ is the row-removed identity matrix with kept-row indices indicated by index-set Ω. Let Ω be the $2\times$ subsampling set, i.e. the indices of all k-space entries in columns $1, 3, \ldots, N_2 - 1$, where we are only subsampling k-space column-wise. We call $\boldsymbol{x}_{\mathrm{ZF}} = \boldsymbol{F}_N^{\mathrm{H}} \boldsymbol{I}_\Omega^\top \boldsymbol{y}$ the zero-filled reconstruction of $\boldsymbol{y}$. Simulate your observation $\boldsymbol{y}$ with $\sigma_n = \sqrt{N}/2^{11}$. Display the zero-filled reconstruction along with the log-scale zero-filled k-space. Is there aliasing? For a general brain image, under what conditions would the aliasing be completely-removable (i.e. via filtering or cropping)?

(c) Give a description and an image-domain expression for a general $R\times$ subsampling of k-space (columns $1, 1+R, 1+2R, \ldots$) in terms of $\boldsymbol{x}$ (assume $\boldsymbol{\nu} = 0$).[1] Verify your expression by comparing it numerically against different subsampling rates.

(d) It is common to employ a "variable density" subsampling scheme where a fraction of the center k-space frequencies are kept, and the remainder are $R\times$ subsampled. These center-frequencies are often called the Auto-Calibration Signal (ACS). Write a function `gen_uniform_mask` which accepts the number of lines `n`, the subsampling rate `R`, and the center-fraction `cf`, returning a variable density mask. Display your mask for ($R = 2$, $cf = 0.16$) and ($R = 8$, $cf = 0.04$). Display your zero-filled reconstruction and zero-filled k-space side-by-side for the $R = 2$ case. Give a qualitative description compared to the non-variable-density subsampling case.

(e) Give an expression for the variable density mask $\boldsymbol{m} \in \{0,\, 1\}^N$ in terms of the $R\times$ subsampling mask $\boldsymbol{m}_R$ and the ACS mask $\boldsymbol{m}_{\mathrm{ACS}}$.

(f) With $\boldsymbol{M} = \mathbf{diag}(\boldsymbol{m}) \in \{0,\, 1\}^{N \times N}$ derive an expression for the zero-filled reconstruction of masked k-space, $\boldsymbol{x}_{\mathrm{ZF}} = \boldsymbol{F}_N^{\mathrm{H}} \boldsymbol{M} \boldsymbol{F}_N \boldsymbol{x}$ (note $\boldsymbol{M} = \boldsymbol{I}_\Omega^\top \boldsymbol{I}_\Omega$), in terms of image-domain convolutions with $\boldsymbol{x}$. Verify your expression by comparing numerically against the zero-filled reconstruction across acceleration rates and center-frequencies.

(g) Consider the multi-coil MRI forward model,

$$\boldsymbol{y}_c = \boldsymbol{I}_\Omega \boldsymbol{F}_N (\boldsymbol{s}_c \odot \boldsymbol{x}) + \boldsymbol{\nu}_c, \quad \boldsymbol{\nu}_c \sim \mathcal{C}N(0,\, \sigma_c^2 \boldsymbol{I}), \quad c = 1, \ldots, C,$$

where C is the total number of receiver coils used to collect data (i.e. each coil collects its own k-space, but the sampling pattern is shared), and $\boldsymbol{s}_c$ is the c-th coil's sensitivity map. Write the forward model for multi-coil k-space $\boldsymbol{y} = \begin{bmatrix} \boldsymbol{y}_1^{\mathrm{H}} & \boldsymbol{y}_2^{\mathrm{H}} & \cdots & \boldsymbol{y}_C^{\mathrm{H}} \end{bmatrix}^{\mathrm{H}}$ in block-matrix form so that $\boldsymbol{y} = \boldsymbol{A}\boldsymbol{x} + \boldsymbol{\nu}$.

(h) Simulate the $2\times$ subsampled variable-density observation $\boldsymbol{y}$ (with $\sigma_c = \sigma_n$ for every coil) and the corresponding zero-filled reconstruction $\boldsymbol{x}_{\mathrm{ZF}} = \boldsymbol{A}^{\mathrm{H}} \boldsymbol{y}$. Qualitatively describe your results compared to the single-coil case. Are they better, worse, why? Justify your answer in terms of coil-sensitivities, "anti-aliasing filtering", and "anti-imaging filtering".

[1] Depending on how you write it, your expression may only hold when R divides N_2.