ECE-491 Select Topics In Machine Learning: Machine Learning for Medical Imaging

Signal Processing, Optimization, and Deep-Learning Methods

Fall 2026

Department of Electrical and Computer Engineering

The Cooper Union for the Advancement of Science and Art

Instructor	Nikola Janjušević
Contact	nikola.janjusevic@cooper.edu
Format	Independent study: one hour of lecture per week, followed by assigned readings, homework, and a final project
Class Hours	TBA
Office Hours	By appointment
References	– Prince, J. L., and J. M. Links, *Medical Imaging Signals and Systems*, 2nd ed., Pearson, 2014. – Fessler, J. A., *Image Reconstruction: Algorithms and Analysis*, book draft. – Boyd, S., and L. Vandenberghe, *Convex Optimization*, Cambridge University Press, 2004. – Parikh, N., and S. Boyd, *Proximal Algorithms*, Foundations and Trends in Optimization, 2014. – Chambolle, A., and T. Pock, "An introduction to continuous optimization for imaging," *Acta Numerica*, 2016. – Ongie, G., et al., "Deep learning techniques for inverse problems in imaging," *IEEE J. Sel. Areas Inf. Theory*, 2020. – Monga, V., Y. Li, and Y. C. Eldar, "Algorithm unrolling," *IEEE Signal Processing Magazine*, 2021. – Goodfellow, I., Y. Bengio, and A. Courville, *Deep Learning*, MIT Press, 2016.
Software	Your choice of programming language. Python with PyTorch, or Julia with Lux.jl / Flux.jl, are recommended.
Course website	nikopj.github.io/teaching/ml4medimg26/

Course Description: This course covers medical image reconstruction, restoration, and analysis. A CT scanner measures line integrals and an MRI scanner measures Fourier samples; recovering an image from either is an ill-posed inverse problem, and nearly every method for doing so is an optimization algorithm. The course develops the subject the subject of image/signal recovery from that point of view.

The first half covers the signal processing behind the forward models – sampling, convolution, and the Fourier and Radon transforms – then the matrix algebra and convex optimization needed to invert them, and finally the variational framework of data fidelity plus regularization that underlies sparsity, total variation, and low-rank methods. The second half treats deep learning as a continuation of that framework: algorithm unrolling turns an iterative solver into a network architecture, plug-and-play swaps a proximal operator for a trained denoiser, and diffusion models fit the prior itself. The final weeks cover self-supervised methods, which train a reconstruction network without ground-truth images. Machine learning appears throughout.

The course runs as an independent study. One hour of lecture per week frames the material, and you develop it through assigned readings, problem sets, and a final project. Each problem set asks you to derive a method and then implement it on real imaging data. 3 credits.

Prerequisites: Linear algebra, probability, signals and systems, and programming experience in a language of your choice. Prior exposure to machine learning or convex optimization helps, but Weeks 3–6 develop what you need. Speak with me before registering if you lack this background.

Learning Outcomes: By the end of this course, students will be able to

1. state the forward model of a medical imaging modality as a linear (or linearized) operator and characterize why its inversion is ill-posed;
2. formulate a reconstruction task as a variational problem, and choose a data-fidelity term justified by the measurement noise model;
3. select and implement an appropriate optimization algorithm – gradient-based, proximal, splitting, or primal-dual – for a given problem's smoothness and structure;
4. derive and train model-based deep learning architectures – unrolled networks, plug-and-play, learned priors – and explain their relationship to the underlying iterative algorithm;
5. evaluate a reconstruction or analysis method critically, and communicate the result in a written and oral report.

ABET Outcomes for this Course: ABET is a nonprofit, non-governmental organization that accredits college and university programs in the disciplines of applied science, computing, engineering, and engineering technology. As part of the accreditation process, engineering programs are required to assess student outcomes which are acquired by students who are enrolled in the program. Student outcomes are succinct statements that describe what students are expected to know and be able to do by the time of graduation. These outcomes relate to skills, knowledge and behaviors that students acquire as they progress through the program. The outcomes most closely associated with this course (taken from the ABET website) are:

1. an ability to identify, formulate, and solve complex engineering problems by applying principles of engineering, science, and mathematics

7. an ability to acquire and apply new knowledge as needed, using appropriate learning strategies.

Grading Scheme:

Class Conduct and Participation	5%
Problem Sets	40%
Study Notes	20%
Project	35%

1. **Class Conduct, Lab Conduct, and Participation**: You are expected to contribute to discussions related to course material throughout the semester in a manner respectful to your instructor and classmates. This can be in the form of answering/posing questions during class and/or office hours. Lecture time is short, so come having done the assigned reading, and bring questions on it.

   You are expected to attend all lectures. Please notify me beforehand of any planned absence.

   Deductions may incur for violations of classroom conduct and/or technology policy (see Grading Scale below).

2. **Problem Sets**: Problems are posted on the course website. An assignment is due at the start of class the following week: HW1, listed in Week 2 of the schedule below, is due in Week 3. There are six assignments, each with a written component and a computational component in the language of your choice.

   Each assignment lists three problems of comparable weight, and you choose **two**. Pick the ones you want to spend the week on. I will read a third attempt, though it earns no additional credit.

   Submit a single PDF report together with your code. Include the script or notebook that generated every figure you show.

3. **Study Notes**: With each problem set, submit typed notes on the topics listed for the corresponding weeks of the course schedule. One lecture hour per week cannot reach every listed topic; lecture frames the material and you study the remainder from the assigned readings. These notes verify that work. Bullet points and equations suffice, and LaTeX is preferred but any typed format is acceptable.

   Keep the notes to 2-3 pages per assignment. Transcribing the readings is easy and teaches little. Choosing what survives onto two pages requires that you have ranked the material, which requires that you have understood it. For each topic you keep, state the definition or result and give an argument for why it holds. I grade these on correctness and on the judgment behind what you included.

4. **Final Project**: The project is assigned in Week 9 and due at the end of the semester. Either reproduce and critically evaluate a recent paper in machine learning for medical imaging, or propose and carry out an original method or study. Deliverables are a short proposal (Week 11), a conference-style written report, code, and a presentation in Week 15.

Grading Scale:

90 - 100	A – superior and comprehensive understanding of the course principles
80 - 89	B – good degree of familiarity of the course principles
70 - 79	C – average knowledge of course principles and fair performance
60 - 69	D – minimum working knowledge of the course principles
< 60	F – unsatisfactory understanding of the course principles

The following section is from Associate Dean Shay's ECE-291 Syllabus, which I adopt and endorse for this course.

Issues and Concerns: I try to create course policies that support a fair and equitable classroom environment and set high performance standards for all students. I hope to create an inclusive learning environment where you feel both challenged but also constantly respected and recognized within the course. Please make an appointment with me if you are having any issues related to me, the course, or your fellow students.

While I want you to feel comfortable coming to me with issues you may be struggling with or concerns you have, please be aware that I have reporting requirements that are part of my responsibilities as a member of the faculty. If you inform me of an issue of sexual harassment, sexual assault, or discrimination, I will keep the information as private as I can, but I am required to report the basic facts of the incident to Cooper's Title IX Coordinator. The Cooper Union Title IX policy on sexual misconduct can be found here.

Counseling Services at The Cooper Union are coordinated through the Office of Student Affairs. Cooper's counseling and mental health services website can be found here.

Technology Policy: You may not use electronic devices (laptops, tablets, phones) during lecture without my permission. Laptops are expected for in-class computational demonstrations, which will be announced in advance.

Data and Patient Privacy: All course data is publicly released, de-identified research data. Do not bring identifiable patient data, or data obtained under a clinical agreement that forbids it, into any coursework. Clear any outside data source with me first – a lab, a hospital, an internship.

Accommodations: Students with disabilities or who need special accommodations for this class are required to notify the Dean of Students and meet with me so that arrangements can be made. The Cooper Union has limited resources and extra lead time is required for such arrangements to be feasible.

In order to receive accommodations for an exam, you must notify me in writing at least two weeks before the accommodations are needed and you must also be registered with the Dean of Students. Students will not be afforded any special accommodations retroactively, i.e., for academic work completed prior to disclosure of the disability to me. Support services for students are described here.

Group Work and Academic Integrity: You are encouraged to discuss your homework and projects outside of class time with your classmates, however, all work submitted must be your own. You are required to acknowledge collaboration with other parties (including a website or textbook reference) in arriving at your answers/explanations in homework and projects. Submitted work that is strongly suspected of being misrepresented as your own (i.e. plagiarism, including copying the output of generative AI) will be reported to the Dean's Office and may result in a zero on the assignment/exam and/or a D/F in the course.

This course is about understanding *why* an algorithm works, so derivations must be your own reasoning. Use established numerical, automatic-differentiation, and medical imaging libraries freely. You may adapt published reference implementations if you cite them explicitly and can explain every line you submit.

Course Schedule (tentative): Below is an outline of the course schedule, subject to change. The italicized topics are the syllabus proper. Each is your responsibility whether or not lecture reaches it, and each is fair game for the study notes you submit with the corresponding problem set.

Week	**Schedule**
	Part I: Forward Models and Signal Processing
1	Medical Imaging as an Inverse Problem *Survey of modalities (CT, MRI, PET/SPECT, ultrasound); the forward model* $\boldsymbol{y} = A\boldsymbol{x} + \boldsymbol{n}$*; well-posedness, ill-conditioning, and why naïve inversion fails*
2	Signal Processing for Imaging *Sampling, aliasing, and the discrete Fourier transform; the Radon transform, projection-slice theorem, and filtered backprojection; MRI k-space and undersampling* **HW1**: Radon/FBP for CT, k-space undersampling artifacts in MRI, fitting the FBP ramp filter from data
	Part II: Optimization Foundations
3	Matrix Algebra and Matrix Calculus *Matrix algebra and the four fundamental subspaces; gradients and Jacobians of matrix expressions; the chain rule and reverse-mode automatic differentiation; the SVD, pseudo-inverse, and least squares; conditioning and Tikhonov regularization*
4	Structured Operators and Iterative Solvers *Matrix algebra of composed and block operators; vectorization and Kronecker products; adjoints of composed operators and the adjoint test; matrix-free operators; the normal equations and conjugate gradients* **HW2**: Matrix calculus and reverse-mode differentiation, Tikhonov via FFT with a learned regularization weight, iterative CT reconstruction
5	Convex Optimization *Convex sets and functions, subgradients, optimality conditions, Lagrange duality; gradient descent, Nesterov acceleration, and stochastic gradient descent; convergence rates*

Week	**Schedule** (cont.)
6	Statistical Estimation and Noise Models *Gaussian, Poisson (CT/PET), and Rician (MRI magnitude) noise; noise-level estimation and whitening; maximum likelihood and MAP estimation; penalized-likelihood reconstruction and the Bayesian view of regularization; fitting prior parameters from data (empirical Bayes)* **HW3**: Noise-level estimation and whitening, MAP estimation for Poisson-likelihood low-dose CT, learning the prior scale from a dataset
	Part III: Variational Methods
7	Non-smooth Optimization and Sparsity *The proximal operator and soft-thresholding; proximal gradient (ISTA) and FISTA; wavelet sparsity, compressed sensing, and accelerated MRI*
8	Splitting Methods and Total Variation *The ROF model and total variation; ADMM and variable splitting; Chambolle–Pock primal-dual; staircasing and the limits of hand-designed priors* **HW4**: Wavelet CS-MRI with FISTA, TV denoising/deblurring via ADMM and primal-dual, learned shrinkage thresholds from a few unrolled ISTA iterations
9	Structured Priors and Multi-Coil Reconstruction *Mixed norms and the nuclear norm; low-rank plus sparse for dynamic MRI; parallel imaging (SENSE, ESPIRiT)* **Final Project Assigned**
	Part IV: Deep Learning
10	Deep Learning Foundations for Imaging *Backpropagation revisited; CNNs, the U-Net, and deep denoisers; direct supervised reconstruction and segmentation; training, validation, and overfitting; loss functions and image quality metrics* **HW5**: Coil sensitivities and SENSE, CNN denoiser, U-Net segmentation
11	Algorithm Unrolling and Plug-and-Play *From iteration to architecture: LISTA, MoDL, and the end-to-end VarNet; learned regularizers and learned proximal operators; replacing the prox with a denoiser (plug-and-play) and conditions for convergence; data efficiency vs. black-box networks* **Project Proposal Due**
12	Unsupervised and Self-Supervised Learning *Training without ground truth: Noise2Noise and Noise2Void, Stein's unbiased risk estimate; SSDU and k-space self-supervision; the deep image prior* **HW6**: Plug-and-play ADMM with your HW5 denoiser, unrolled reconstruction network, self-supervised training (SSDU)
13	Diffusion Generative Methods for Inverse Problems *Score matching and denoising diffusion models; the score, the MMSE denoiser, and Tweedie's formula; diffusion priors for reconstruction (e.g. ImMAP, diffusion posterior sampling); uncertainty quantification*
14	Review and Project Work *Course review; project consultations; slack time*
15	Project Presentations