

Final Project: Implementation and Evaluation of Imaging Methods

ECE-491 Machine Learning for Medical Imaging

Fall 2026

Assigned: Week 9 **Due:** Week 16

Overview. Implement and evaluate imaging methods from the research literature, and report the results in a four-page IEEE conference paper and a ten-minute presentation. Two project types are acceptable.

- (a) *Implementation.* Implement a method from a peer-reviewed paper and evaluate it against a course method.
- (b) *Comparison.* Implement and compare several methods on a common imaging problem. Course methods may be included, provided that at least one compared method was not implemented in the homework.

Methods may be classical or deep-learning based and must address a medical imaging problem in reconstruction, restoration, or analysis. The project is individual and accounts for 35% of the course grade.

Method selection. The principal method of an implementation project, and at least one method of a comparison project, must come from a paper that satisfies the following conditions.

- (a) It is published in a peer-reviewed journal or conference.
- (b) Its principal method was not implemented in the homework (Table 1). Course methods not implemented in the homework are eligible with approval.
- (c) It reports at least one reproducible quantitative result.
- (d) Its data are publicly available.
- (e) It is feasible with the available compute, at full scale or at a justified reduced scale.

Submit the project type and paper citations by email in Week 10. Each paper may be selected by one student only.

Set	Methods implemented
HW1	Pseudo-inverse deconvolution; zero-filled MRI reconstruction
HW2+3	Noise covariance estimation and whitening; Tikhonov reconstruction; gradient descent; conjugate gradients
HW4	Total variation by the subgradient method, proximal gradient descent, and ADMM; filtered backprojection
HW5*	SENSE with estimated coil sensitivities; CNN denoiser; U-Net segmentation
HW6*	Plug-and-play ADMM; unrolled reconstruction network; SSDU

Table 1: Methods implemented in the homework. Entries marked * are planned and subject to change.

Paper	Type	Problem	Baseline
GRAPPA, Griswold et al., MRM 2002	Classical	Parallel MRI	SENSE
Dictionary-learning MRI, Ravishankar and Bresler, IEEE TMI 2011	Classical	Undersampled MRI	Wavelet CS
Regularization by denoising, Romano et al., SIAM J. Imaging Sci. 2017	Hybrid	Deblurring, MRI, CT	Plug-and-play
Learned primal-dual, Adler and Öktem, IEEE TMI 2018	Unrolled	Low-dose CT	Chambolle–Pock
Deep equilibrium models, Gilton et al., IEEE TCI 2021	Implicit	MRI, CT	MoDL
Noise2Inverse, Hendriksen et al., IEEE TCI 2020	Self-supervised	CT denoising	Noise2Noise
Equivariant imaging, Chen et al., ICCV 2021	Self-supervised	CT, inpainting	SSDU
NeRP, Shen et al., IEEE TNNLS 2024	Neural field	Sparse-view CT, MRI	Deep image prior
Score-based inverse problems, Song et al., ICLR 2022	Generative	Sparse-view CT, MRI	DPS
VoxelMorph, Balakrishnan et al., IEEE TMI 2019	Unsupervised	MRI registration	Variational registration

Table 2: Example papers. Any paper satisfying the conditions above is eligible.

Schedule. Milestones are listed by course week.

Week	Milestone	Week	Milestone
9	Project assigned	14	Project consultations
10	Project type and citations	15	Presentations
11	Proposal due	16	Paper and code due
13	Progress check-in		

Proposal (Week 11). One page, prepared in L^AT_EX with the IEEEtran conference class, containing

- (a) the project type, the paper citations, and links to any released code;
- (b) the imaging modality, forward model, and noise model;
- (c) for an implementation, the novelty of the method relative to the closest course method; for a comparison, the question the comparison addresses;
- (d) the dataset and the train, validation, and test split;
- (e) the baselines or compared methods;
- (f) the table or figure to be reproduced;
- (g) the evaluation metrics;
- (h) the compute plan and a reduced-scale alternative.

Progress check-in (Week 13). A 15-minute meeting. Present a forward operator that passes the adjoint test, baseline results on the test set, and a preliminary result from at least one implemented method.

Conference paper (Week 16). Submit a single PDF in the format of IEEE ISBI submissions.

- Prepare the paper in $\text{\LaTeX}$ with the unmodified IEEEtran conference class, `\documentclass[conference]{IEEEtran}`.
- At most four pages, including figures, tables, and appendices. An optional fifth page may contain references only.
- Single author.
- Figures must be legible at column width. Compared images share a display window and colorbar.
- Cite all papers, datasets, and code in IEEE style.

The paper contains the sections listed in Table 3.

Section	Contents
Abstract	At most 150 words: problem, methods, and principal quantitative result
Introduction	Problem, its difficulty, the methods studied, and the contributions of this study
Problem formulation	Forward model, noise model, and objective or learning problem, in course notation
Methods	Algorithms or architectures, losses, and pseudocode sufficient for reimplementaion; relation to course methods
Experiments	Data, preprocessing, data split, baselines, metrics, hyperparameter selection, hardware, and runtime
Results	Reproduced results alongside the published values; comparison of methods; qualitative figures
Discussion	Sources of discrepancy, failure cases, and limitations
Conclusion	Conclusions supported by the results

Table 3: Required sections of the conference paper.

The experiments must include

- a reproduction of the target result, or a quantitative account of the discrepancy;
- a comparison of all methods on identical data, with hyperparameters tuned on the validation set, including at least one course method for an implementation project;
- one ablation or sensitivity study, such as the number of unrolled iterations, the sampling rate, the noise level, or the training set size;
- test-set metrics reported as mean $\pm$ standard deviation over test images;
- the parameter count and training time of any learned method.

Presentation (Week 15). Ten minutes, followed by five minutes of questions. Presentations are stopped at ten minutes. Submit slides as a PDF by email before class. A suggested allocation is: problem and forward model (1.5 min), methods (3 min), experimental setup (1 min), results (3 min), and limitations (1.5 min). Attendance at all presentations is required, and questions to presenters count toward participation.

Code (Week 16). Submit a repository link or a zip archive with the paper. The submission must regenerate every figure and table in the paper and include

- (a) a README with setup instructions, data download instructions, and one command per figure or table;
- (b) pinned dependencies;
- (c) fixed random seeds and a description of the training hardware;
- (d) a data download script in place of the data, and trained weights under 100 MB or a link to them;
- (e) an attribution of each file as original, adapted (with source), or third-party.

Implement the core of each method yourself: the algorithm or architecture, the loss, and the iteration or training loop. Libraries may be used for forward operators, data loading, automatic differentiation, and metrics (e.g., PyTorch, Lux.jl, DeepInverse, ODL, ASTRA, SigPy, BART). Library implementations of the selected methods and unmodified author code are not acceptable. Adapted reference code must be cited, per the syllabus.

Grading. The project is graded out of 100 points. A reproduction that falls short of the published results may receive full credit if the discrepancy is quantified and explained.

Component	Criterion	Points
Proposal	Completeness and feasibility	10
Check-in	Forward operator, baseline, and preliminary result	5
Paper	Problem formulation	5
	Description of methods	10
	Reproduction of the target result	15
	Experimental evaluation	10
	Discussion	5
	Writing and format	5
Presentation	Clarity of problem and methods	10
	Results and supporting evidence	8
	Timing and responses to questions	7
Code review	Reproducibility, clarity, and attribution	10
Total		100

Policies. The syllabus policies on academic integrity, generative AI, and data privacy apply. Cite all papers, datasets, and code in both the paper and the code. Data from sources other than public datasets require prior approval.

Resources. Suggested public datasets: fastMRI (multi-coil knee and brain k-space), the MoDL brain dataset (HW1, HW2+3), LoDoPaB-CT (simulated low-dose CT), Low Dose CT Image and Projection Data (paired full-dose and low-dose CT), and the Medical Segmentation Decathlon (CT and MRI segmentation). Suggested software: DeepInverse, ODL, and the ASTRA Toolbox for tomography; SigPy and BART for MRI. For paper selection, begin with the surveys of Ongie et al. (2020) and Monga et al. (2021) listed in the syllabus, and recent issues of IEEE TMI, IEEE TCI, Medical Image Analysis, and the MICCAI and ISBI proceedings.