







Algorithmic Foundations of Data Science, Spring 25



- This course is intended for grad students and advanced undergrads.
 - The course will be **fast-paced**.
- Contact (alperen.ergur@utsa.edu) if you are in doubt for enrollment.
- This course will **NOT** get you a data scientist job, but it will make you much better at it. Just don't take my word - come and listen industry experts on how this course will benefit you both for academic and industry roles.



Momentum Gradient Descent

Guest Lectures by Industry Experts


Dr. Jens Forsgard, Staff Data Scientist, Reverb

Dr. Enes Bilgin, Staff ML Engineer, Latitude AI


Dr. Burhan Cikili, Lead Data Scientist, Wells Fargo

Dr. Yunus Abdulhameed, Algorithms Engineer, Seno Medical


Dr. Hakan Gunturkun, Head of ML/AI Research, Heretic Ventures



Dual Methods for regularized ERM



Denoising autoencoder



(1) Fundamentals – 6 weeks

Linear Algebra (1 week): Linear regression
Eigen and singular values, Cholesky and QR, a pinch of numerical analysis
Projections, Best fit subspaces: PCA

Basic probs (1.5 weeks): Categorical Distribution, Softmax, and Logistic regression
Gaussians, conditional Gaussians, Hinton diagrams
Gaussian-Bayesian rule, Probabilistic PCA

Basic stats-Fundamentals of Prediction and Learning (1.5 weeks):
Max-likelihood, prior-posterior, conjugate pairs
KL divergence, empirical risk minimization
Neyman-Pearson Lemma, ROC curves,
bias-variance trade-off, VC dimension, Radamacher complexity

Neural Nets (2 weeks): MLP's and Backpropagation algorithm
Convolutions and CNN's
Recurrent neural nets, attention layers, transformers

(2) Optimization Algorithms - 4 weeks

Gradient Descent, Accelerated Gradient Descent (Momentum and Long-Step variants)
Preconditioning gradient descent: Newton's Method and AdaGrad
Stochastic Gradient Descent, Randomized Kaczmarz, Deshuffling, ADAM
Regularization: LASSO and Hinge Regression
Coordinate Descent, Graphical LASSO
Duality, Non-smooth optimization, Regularized Risk Minimization

(3) Probabilistic Algorithms - 2 weeks

Hashing and Bloom filters
Johnson-Lindenstrauss Dimensionality Reduction, Streaming and Sketching
Markov Chains and MCMC

(4) Beyond Supervised Learning - 2 weeks

Kernel Methods: Mercer Kernels, SVM's
Decision Trees, Bagging, Random Forests
Gradient Boosting, XGboost algorithm
Autoencoders and Variational Autoencoders

(5) Bonus Material: Manifold Learning, Topological methods

Industry Speakers:

Dr. Hakan Gunturkun, Head of ML Research @ Heretic Ventures

Talk title: Introduction to Generative AI

Dr. Yunus Abdulhameed, Algorithms Engineer @ Seno Medical Instruments

Talk title: Signal Processing Tools in Data Science

Dr. Burhan Cikili, Lead Data Scientist @ Wells Fargo

Talk title: Causal Modeling and Inference

Dr. Enes Bilgin, Staff Machine Learning Engineer @ Latitude AI

Talk title: Reinforcement Learning with Human Feedback

Dr. Jens Forsgaard, Staff Data Scientist @ Reverb

Talk title: Ranking and Search