PROFESSIONAL EDUCATION
Short Programs

Modeling and Optimization for Machine Learning
July 16-20, 2018

Instructors:
Prof. Justin Solomon, MIT Department of Electrical Engineering & Computer Science
Dr. Suvrit Sra, principal research scientist, MIT Laboratory for Information and Decision Systems

Monday (7/16/2018)

10:00 Introduction: Overview of optimization (Suvrit)
11:00 Discussion and coffee (15 min)
11:15 Basic notions: Optimality (Justin)
12:15 Lunch break
13:15 Gradient descent, stochastic gradient descent (Suvrit)
14:15 Discussion and coffee (30 min)
14:45 Case study 1: Classification and regression problems in machine learning (Suvrit)
15:30 Practicum
17:00 END

Tuesday (7/17)

9:30 Second-order methods (Newton-type methods, quasi-Newton) (Justin)
10:30 Discussion and coffee (30 min)
11:00 Case study 2: Graph-based learning (Suvrit)
12:00 Lunch break
13:00 Working with constraints in optimization (Justin)
14:00 Discussion and coffee (30 min)
14:30 Case study 3:Assignment problems, optimal transport, and GANs (Justin)
15:30 Practicum: CVX and other solvers, entropy regularization
17:00 END

Wednesday (7/18) -- Justin

9:30 Sparsity, low-rank optimization, smoothness, and other considerations
10:30 Discussion and coffee
11:00 Case study 5: Nonlinear image analysis
12:00 Lunch break
13:00 Advanced methods: Proximal algorithms, splitting, and ADMM
14:00 Discussion and coffee
14:30 Case study 6: Social networks and gossip-based consensus
15:30 Practicum
17:00 END

PROFESSIONAL EDUCATION
Short Programs

Modeling and Optimization for Machine Learning
July 16-20, 2018

Instructors:
Prof. Justin Solomon, MIT Department of Electrical Engineering & Computer Science
Dr. Suvrit Sra, principal research scientist, MIT Laboratory for Information and Decision Systems

Thursday (7/19) -- Suvrit

9:30 Large-scale optimization algorithms: Convex problems
10:30 Discussion and coffee (30 mins)
11:00 Case study 3: Online advertising problem, large-scale logistic regression
12:00 Lunch break
13:00 Large-scale optimization algorithms: Non-convex problems
14:30 Discussion and coffee (30 mins)
15:00 Case study 4: Training neural networks, automatic differentiation
15:30 Practicum: TensorFlow, Torch
17:00 END

Friday (7/20)

9:30 Other nonconvex methods (EM, CCCP) (Suvrit)
10:30 Discussion and coffee
11:00 Case study 7: Clustering, embedding, and visualization (Justin)
12:00 Lunch break
13:00 Practical guide to OPTML (S+J)
14:00 Discussion and coffee
14:15 Optimization and modeling team competition
16:00 Course Closing