New class on modeling using machine learning: 6.883J/2.168J/… J

The College of Computing and School of Engineering are developing a new interdisciplinary class for undergraduate and graduate students who wish to delve into the fundamentals of Machine Learning in the service of engineering and applied science outside the realm of mainstream computer science. The class will be offered in Spring 2021 for the first time and will be structured as two core lectures per week, offered by Course 6 faculty and covering the basic principles, methods and proofs; and a third specialized weekly lecture offered separately by Course 1, 2, 3, 10, 15, 20 and 22 faculty covering the use of ML in diverse topics such as forces and waves, sensing and control in complex structures, fluids, living matter, pharmaceutical and cell manufacturing, and system design. Student project topics will be derived from data-intensive research and product development challenges in these specialized disciplines, e.g. imaging through fog by capturing turbid dynamics with a recurrent neural network.

Machine Learning emerged approximately six decades ago at the intersection of neuroscience, mathematical optimization, function representation theory, and statistical physics; its success was fueled by computing systems becoming cheaply and ubiquitously available through the cloud. Extensive use of ML in the physical and life sciences and engineering became apparent only fairly recently, however, as the relevant mathematical ML tools became powerful enough to handle complex constitutive relationships and reveal unexpectedly favorable designs for complex tasks. Publications have proliferated demonstrating ML successes in scientific research such as identifying particle interactions, explaining protein structure, linking therapies to the biomechanics of disease, and tackling the multi-scale spatial and temporal dynamics of climate. The new IAIFI (Institute for AI in Fundamental Interactions), funded by NSF at MIT and Harvard, is poised to lead these developments on the particle physics side, and surely more will follow.

One crucial feature of ML's expansion into these first principles-grounded fields is inherent algorithmic *transparency,* as the data become complementary to governing equations rather than supplant the equations. This gives hope for broader understanding of the ethical implications of ML-driven discoveries and design, and firmer guidance of ML applications toward a fair and diverse society.

MIT has been unique in encouraging a porous boundary between science and engineering, building on the legacy of pioneers such as Norbert Wiener and Vannevar Bush. It is not uncommon for MIT scientists to double up as entrepreneurs, neither is it uncommon for MIT engineers to make fundamental discoveries in their fields. MIT has also been at the forefront of Computer Science and Artificial Intelligence, establishing bustling laboratories long before other universities even began to grasp the significance of these then-fledgling fields.

As we stand at the forefront of a new revolution being delivered by ML to the physical and life sciences and engineering, we need to endow our students with deep understanding of ML's foundations coupled with and complementary to their own specific disciplines. Even though this new class is initially anchored in the School of Engineering, we believe that participation from Science students will inspire and enrich everyone's overall experience, especially in the student projects. We would like to invite students from all Science departments to sign up for the class and look forward to engaging them in an exciting environment of mutual learning.

Class description and core syllabus

The course focuses on modeling with machine learning methods with an eye towards applications in engineering and sciences. Students will be introduced to modern machine learning methods, from supervised to unsupervised models, with an emphasis on newer neural approaches. The course focuses on the understanding of how and why the methods work from the point of view of modeling, and when they are applicable. Using concrete examples, the course covers formulation of machine learning tasks, adapting and extending methods to given problems, and how the methods can and should be evaluated.

Introduction to modeling
Formulation of ML problems, linear models
Feature design, impact, probabilistic prediction
Recommender problems
Non-linear prediction, neural networks
Backpropagation
ML packages
Convolutional models for images, text
Modeling sequences, recurrent neural networks
Neural models on/for graphs
Transfer learning
Overview, computer vision
Unsupervised learning
Latent variable models, mixtures, EM
Complex latent variable models, VAE
Generative adversarial networks
Pre-training, semi-supervision, contrastive estimation
Interpretability and robustness
Reinforcement learning
Deep RL
Learning causal models

Physics and Dynamics Specialized Syllabus (of interest to Courses 2, 6, 8, 12, 18, …)

Dynamics and statistical mechanics – Hopfield and Boltzmann networks, statistical learning
Identification and estimation – from Kalman filters to probabilistic neural networks
Spatial-temporal dynamics – deep neural networks in PDE parameter estimation
Spatial-temporal deep recurrent networks and attention, transformers
Inverse problems – learning sparsity, supervised, unsupervised

Project topics: imaging in turbid media, control of dynamical systems under strongly non-Gaussian uncertainty, control of non-Newtonian fluids, non-equilibrium thermal processes, estimation and control in manufacturing, generative mechanical and optical design. Students may also propose their own topics.

Biochemically Motivated Specialized Syllabus (of interest to Courses 20, 10, 7, 6-7, 5, …)

- Clustering methods for high-throughput molecular data (transcriptomics, proteomics, etc.)
- Design of new molecules
- Engineering of molecular interactions

- Causal modeling of biological processes

Project topics: Designing proteins to neutralize SARS-CoV2, designing inhibitors of oncogenes, recovering developmental pathways from single-cell data.