PHYS 503 – Instrumentation Physics: Applications of Machine Learning

Course description

This course is designed to give students a solid foundation in machine learning applications to physics, positioning itself at the intersection of machine learning and data-intensive science. This course will introduce students to the fundamentals of analysis and interpretation of scientific data, and applications of machine learning to problems common in laboratory science such as classification and regression. There will be two 75-minute classes each week, split into discussions of core principles and hands-on exercises involving coding and data. There will be a few projects throughout semester that will build on the course material and utilize open source software and open data in physics and related fields. The list of topics will evolve, according to the interests of the class and instructors. Material will be clustered into units of varying duration, as indicated below. The lists of suggested readings and references are advisory; a large amount of material of excellent quality is now available on the worldwide web, particularly on the sites of university courses addressing the topics of each unit.

A distinguishing feature of this course is its sharp focus on endeavors in the data-rich physical sciences as the arenas in which modern machine learning techniques are taught. The course uses open scientific data, open source software from data science and physics-related fields, and publically-available information as enabling elements. Research-inspired projects are an important part of the course and students will not only execute them but will play an active role in helping define and shape them. Example projects might include machine learning approaches to searches for new particles or interactions at high-energy colliders; methods of particle tracking and reconstruction; identification, classification and measurement of astrophysical phenomena; novel approaches to medical imaging and simulation using techniques from physics and machine learning; machine learning in quantum information science. Through these projects and the course material, students will learn how large datasets in physics are generated, curated, and analyzed, using machine learning as a tool to generate key insights in both experimental and theoretical science.

Credit and grading

Students must register for this course in the fall semester for a total of 4 credit hours. Grading is by letter.

Meeting Times

Two 75-minute class sessions per week for 14 weeks; optional office hours as necessary/desired.

Learning objectives

As a result of completing this course, students will

- Develop an understanding though hands-on experience of the use of deep learning and machine learning techniques in the analysis of large, complex data sets drawn from fields that include physics, medicine, and agriculture.
- Be able to evaluate the appropriateness of artificial intelligence methods to their analyses of

Physics 523 project data.
- Understand the tradeoffs in clarity, efficacy, and availability of training data sets associated with the use of machine learning methods in technical analyses.

Syllabus: topics covered

Unit 1: Data Science (3 weeks)

- Introduction to data science and machine learning in the physical sciences and related fields
- Scientific python environment
- Notebooks and numerical python
- Handing, visualizing and finding structure in data
- Dimensionality and Linearity
- Adapting linear methods to nonlinear problems
- Kernel functions

Example physics connections and investigations:
- Jet recognition and clustering algorithms at high energy colliders
- Visualization of complex data (and data simulations) from the Large Hadron Collider, the Deep Underground Neutrino Experiment, the Laser Interferometer Gravitational-Wave Observatory, and the Large Synoptic Survey Telescope (LSST)
- Collaborative analysis tools—scientific python, notebooks, and JupyterLab—used by LSST
- Modeling long time-scale beam instabilities in high energy storage rings

Unit 2: Probability and Statistics (1 week)

- Probability Theory
- Probability Density Estimation from data
- Statistical methods

Example physics connections and investigations:
- Framing the concepts: determining the parameters of fundamental physics from data sets with complex backgrounds

Unit 3: Bayesian Inference (3 weeks)

- Introduction to Bayesian statistics
- Stochastic Processes, Markov Chains and Markov Chain Monte Carlo
- Variational Inference
- Optimization
- Computational Graphs and Probabilistic Programming
- Bayesian Model selection
- Learning in a Probabilistic context

Example physics connections and investigations:
- The physics of cardiac neurology: predictable onset of arrhythmias in nonlinear nerve networks
- Framing the concepts: Occam's razor and Bayesian inference in a comparison of Ptolemaic epicycles with a Newtonian heliocentric model

Unit 4: Supervised Learning (1 week)

- Supervised Learning in Scikit-Learn
- Cross Validation

Example physics connections and possible investigations:

- Classification of galaxies based on Dark Energy Survey images
- Bounding the phase space of electromechanical systems in precision atomic physics experiments

Unit 5: Learning and Inference using Artificial Neural Networks (ANNs) (1 week)

- Introduction to ANNs
- Types of Neural Networks
- Loss functions
- Backpropogation and Training

Example physics connections and possible investigations:

- Suppression of quantum chromodynamics backgrounds in rare-process searches
- Incorporating known symmetries of fundamental physics into loss functions

Unit 6: Deep Learning (4 weeks)

- Convolutional Neural Networks
- Unsupervised learning networks
- Autoencoder networks
- Recurrent Networks
- Graph Networks and Graph Neural Networks
- Deep Reinforcement Learning

Example physics connections and possible investigations:

- Deep Neural Networks for classification and regression analysis of diagnostic medical imagery
- Convolutional Neural Network approach to casting particle physics detector data as an image classification problem
- Time-domain anomaly detection in sky surveys using the Open Supernova Catalog
- Fast simulation of calorimeter response in particle physics detectors using variational autoencoder networks
- Using graph neural networks for charged particle tracking in neutrino and collider experiments
- Neural message passing and interaction networks for the quantum properties of organic molecules
- Diagnosis of congenital cardio-pulmonary pathologies through field analysis of acoustic, myographic, and electrical anomalies in infants
- Applications of reinforcement learning in controlling particle beams and confined plasmas

Unit 7: Methods for accelerated machine learning and inference (1 week)

- GPU accelerators;
- Distributed learning;
- Role hardware accelerators in ML inference

Example physics connections and possible investigations:

- The physics and information science issues of multi-spectral system configuration: time-alignment of data from distributed sensor arrays linked by affordable but unreliable networks
- Application of fast ML inference in multi-messenger astrophysics (e.g. supernovae detection)

Readings and other sources

Unit 1 reading and reference material:

- *A Whirlwind Tour of Python,* Jake VanderPlas: free PDF, http://www.oreilly.com/programming/free/files/a-whirlwind-tour-of-python.pdf, notebooks online, http://nbviewer.jupyter.org/github/jakevdp/WhirlwindTourOfPython/blob/master/Index.ipyn;
- *Python Data Science Handbook,* https://jakevdp.github.io/PythonDataScienceHandbook/;
- *Notebooks and numerical python,* https://nbviewer.jupyter.org/github/illinois-mla/syllabus/blob/master/notebooks/JupyterNumpy.ipynb
- *Handling data,* https://nbviewer.jupyter.org/github/illinois-mla/syllabus/blob/master/notebooks/Pandas.ipynb
- *Visualizing data,* https://nbviewer.jupyter.org/github/illinois-mla/syllabus/blob/master/notebooks/Visualization.ipynb
- *Finding structure in data,* https://nbviewer.jupyter.org/github/illinois-mla/syllabus/blob/master/notebooks/Clustering.ipynb
- *Measuring and reducing dimensionality,* https://nbviewer.jupyter.org/github/illinois-mla/syllabus/blob/master/notebooks/Dimensionality.ipynb
- *Adapting linear methods to nonlinear problems,* https://nbviewer.jupyter.org/github/illinois-mla/syllabus/blob/master/notebooks/Nonlinear.ipynb
- *Kernel Functions,* https://nbviewer.jupyter.org/github/illinois-mla/syllabus/blob/master/notebooks/Nonlinear.ipynb

Unit 2 reading and reference material:

- *Probability theory,* https://nbviewer.jupyter.org/github/illinois-mla/syllabus/blob/master/notebooks/Probability.ipynb
- *Estimate probability density from Data,* https://nbviewer.jupyter.org/github/illinois-mla/syllabus/blob/master/notebooks/Density.ipynb
- *Statistical methods,* https://nbviewer.jupyter.org/github/illinois-mla/syllabus/blob/master/notebooks/Statistics.ipynb

Unit 3 reading and reference material:

- *Bayesian statistics,* https://nbviewer.jupyter.org/github/illinois-mla/syllabus/blob/master/notebooks/Bayes.ipynb
- *Markov-chain Monte Carlo in practice,* https://nbviewer.jupyter.org/github/illinois-mla/syllabus/blob/master/notebooks/MCMC.ipynb
- *Stochastic processes and Markov-chain theory,* https://nbviewer.jupyter.org/github/illinois-mla/syllabus/blob/master/notebooks/Markov.ipynb
- *Variational inference,* https://nbviewer.jupyter.org/github/illinois-mla/syllabus/blob/master/notebooks/Variational.ipynb

- *Optimization,* https://nbviewer.jupyter.org/github/illinois-mla/syllabus/blob/master/notebooks/Optimization.ipynb
- *Frameworks for computational graphs and probabilistic programming,* https://nbviewer.jupyter.org/github/illinois-mla/syllabus/blob/master/notebooks/Frameworks.ipynb
- *An introduction to the theory of Markov processes mostly for physics students,* C. Maes, https://fys.kuleuven.be/itf/staff/christ/files/pdf/pub/markovlectures2015.pdf
- *Bayesian model selection,* https://nbviewer.jupyter.org/github/illinois-mla/syllabus/blob/master/notebooks/ModelSelection.ipynb
- *Learning in a probabilistic context,* https://nbviewer.jupyter.org/github/illinois-mla/syllabus/blob/master/notebooks/Learning.ipynb

Unit 4 reading and reference material:

- *Supervised learning in Scikit Learn,* https://nbviewer.jupyter.org/github/illinois-mla/syllabus/blob/master/notebooks/Supervised.ipynb
- *Cross validation,* https://nbviewer.jupyter.org/github/illinois-mla/syllabus/blob/master/notebooks/CrossValidation.ipynb

Unit 5 reading and reference material:

- *Neural networks,* https://nbviewer.jupyter.org/github/illinois-mla/syllabus/blob/master/notebooks/NeuralNetworks.ipynb
- *Types of Neural Networks,* https://towardsdatascience.com/the-mostly-complete-chart-of-neural-networks-explained-3fb6f2367464

Unit 6 reading and reference material:

- *Deep learning,* https://nbviewer.jupyter.org/github/illinois-mla/syllabus/blob/master/notebooks/DeepLearning.ipynb
- *Relational inductive biases, deep learning, and graph networks*, http://arxiv.org/abs/1806.01261
- *DeepMind's GNN library,* http://github.com/deepmind/graph_nets
- *Graph Neural Networks: A Review of Methods and Applications,* http://arxiv.org/pdf/1812.08434.pdf
- *An Introduction to Deep Reinforcement Learning,* http://arxiv.org/abs/1811.12560
- *A Beginner's Guide to Deep Reinforcement Learning,* http://pathmind.com/wiki/deep-reinforcement-learning

Unit 7 reading and reference material:

- *The Fast Machine Learning Laboratory,* http://fastmachinelearning.org;
- *Fast Inference of Deep Neural Networks in FPGAs for Particle Physics,* https://iopscience.iop.org/article/10.1088/1748-0221/13/07/P07027
- *Real-time Reinforcement Learning,* http://arxiv.org/abs/1911.04448