PHYS 1970D/2620J: Statistical Physics in Inference and (Deep) Learning

Version: 18th September, 2020
(Location: TBD, Meeting Time: TBD)

Instructor: Jatan Buch
Email: firstname_lastname@brown dot edu

Description

In this course students will explore the statistical physics principles underlying probabilistic inference and various neural network architectures. The course is designed to bridge the gap between approaches to teaching modern statistical physics that are either purely theoretical, or focus largely on its applications in data analysis. To that end, there will be a conscious effort to study topics such as: MaxEnt principle, variational methods, Hebb's rule, bias-variance tradeoff, regularization and others with analytical derivations as well as worked-out code examples in Jupyter notebooks. The course will also provide a space for students to interrogate and reflect on the ethical, political, and policy frameworks that are urgently needed in the age of deep learning.

Prerequisites

PHYS 1530 or equivalent (APMA 1690, for example); prior computational experience will be helpful, but not necessary.

The only difference between the 1000-level and 2000-level of the course is in the number (and complexity) of homework problems: what will be optional for the former will be compulsory for the latter. Regardless, both undergraduate and graduate students are encouraged to register for the advanced version!

Learning outcomes

After taking this course students will be able to:

- appreciate the physical insights that underlie the key methods in probabilistic inference and deep learning.
- identify contexts where these techniques are optimally applicable and judge the validity of the approximations made therein.
- implement these ideas and analysis tools to a variety of problems involving large data sets.
- reflect on how deep learning affects social justice and its potential for reinforcing prevalent systems of discrimination.

Note on teaching format

Since it's unclear at this point whether normal operations will resume in Spring 2021, the class will tentatively be taught in the hybrid format, *i.e.* even during in-person classes there will be a simultaneous live lecture broadcast. Furthermore, to ensure an inclusive learning environment, asynchronous discussions via Piazza (or similar) will be encouraged.

For some topics, recorded lectures will be made available before the class, along with supplemental reading and videos. In those cases, the allotted lecture time will be used for reviewing the material followed by discussions and problem solving in breakout groups.

Besides regular office hours, I will also hold an additional recitation section for each homework to provide helpful insights on the problems and engage in free-form dialog about the course material.

Course Outline

Following is a rough outline of the lesson plan. Also indicated are the approximate duration and abbreviated primary references (see next section) in parentheses and square brackets respectively.

Module 1: Introduction (3 weeks)

i) *Icebreaker*: The Longest run of Heads based on https://www.csun.edu/~hcmth031/tlroh.pdf
ii) Thermodynamics of an Ideal Gas:, Ensembles, Entropy, and Free energy [JS]
iii) Statistics 101: Random variables, examples of probability distributions, law of large numbers, central limit theorem [NGT]
iv) Properties of probability distributions [DH]
v) Information theory and the Maximum Entropy (MaxEnt) principle [EJ, PM, NGT]
vi) Variational methods [DM, PM]

Module 2: Basics of Inference (3 weeks)

i) Bayesian inference [PM, DM, EJ]
ii) Regression I: Fitting a line to data [DH, PM]
iii) Machine learning toolkit: Bias-Variance Decomposition, Cost function, Gradient descent, Regularization [PM]
iv) Supervised versus Unsupervised learning [EB, PM]
v) Regression II: Fitting a plane to data [PM, NGT]
vi) Clustering: k-means, Gaussian Mixture Model (GMM) [PM, DM]

Module 3: Statistical Mechanics of (Deep) Learning (2 weeks)

i) Learning as Inference: The single neuron as a classifier [DM, CKS]
ii) Perceptron learning [CKS]
iii) Hopfield network [DM, PM, CKS]
iv) Restricted Boltzmann Machines (RBMs) [DM, PM]

Module 4: Modern Neural Networks (4 weeks)

i) Feed-forward Deep Neural Networks (DNNs) [PM, IG]
ii) Convolution Neural Networks (CNNs) [PM, IG]
iii) Deep Boltzmann Machines [PM, IG]
iv) Graph Networks (if time permits)

Module 5: Social Justice and Artificial Intelligence

[1 think-pair-share style class discussion + writing component]

References

Since I will provide lecture notes and other relevant material, students do not need to purchase any textbooks for the course. That said, I have mostly relied on open-source material to design the course, and their virtual location is indicated wherever applicable.

Statistical physics

J. Sethna, *Statistical Mechanics: Entropy, Order Parameters, and Complexity Vol. 2,* Oxford University Press (2007) // (available freely on the author's website: link) [JS]

Inference

N. Garcia-Trillos, *Lecture Notes for APMA 1690 and APMA 1720* [NGT]
D.W. Hogg et. al., *Data Analysis Recipes,* arXiv:1008.4686 [DH]
E.T. Jaynes, *Probability Theory: The Logic of Science,* Cambridge University Press (2003)

Deep learning

A.C.C. Coolen et. al., *Theory of Neural Information Processing Systems,* Oxford University Press (2005) [CKS]
P. Mehta, *A high-bias, low-variance introduction to Machine Learning for physicists,* arXiv:1803.08823 [PM]
D.C. Mackay, *Information Theory, Inference and Learning Algorithms,* Cambridge University Press (2003) || (available freely on the author's website: link) [DM]
I. Goodfellow et. al., *Deep Learning* || link [IG]
M.A. Nielsen, *Neural networks and deep learning* || link

Ethics of Deep learning

J. Buolamwini and T. Gebru, *Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification* || paper link || TED talk
C. O'Neil, *Weapons of Math Destruction*, Talk at Google || link
T. Simonite, *Should Data Scientists Adhere to a Hippocratic Oath?,* Wired || link
P. Mehta*, Big Data's Radical Potential,* Jacobin || link
K. Taylor, *The Robots are Coming*, Boston Review || link

Classic papers

A. Barron, *Entropy and the Central Limit Theorem,* Ann. Probab. **14** *(1986), 336*
E.T. Jaynes, *Information Theory and Statistical Mechanics,* Phys. Rev. **106**, 620
R. Cousins, *Why isn't every physicist a Bayesian?,* Am. J. Phys. **63** (1995), 398
D. Geman and S. Geman, *Stochastic Relaxation, Gibbs Distributions, and the Bayesian Restoration of Images,* IEEE Transactions on Pattern Analysis and Machine Intelligence, PAMI-6, 1984
S. Duane et. al., *Hybrid Monte Carlo,* Phys. Lett. B. **195**, 216
D.O. Hebb, *The organization of behavior; a neuropsychological theory*, Wiley, 1949
J.J. Hopfield, *Neural networks and physical systems with emergent collective computational abilities,* PNAS 1982 79 (8) 2554-2558
E.L. Bienstock et. al., *Theory for the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex,* J Neurosci. **2** (1982), 32
D.H. Ackley et. al.*, A Learning Algorithm for Boltzmann Machines*, Cognitive Science Volume 9, Issue 1,147
T. Poggio and F. Girasi, *A Theory of Networks for Approximation and Learning ||* link

Useful multimedia links

An intuitive proof of Bayes' theorem || link
Interactive gallery of various MCMC algorithms || link
Gaussian processes demo || link
Visualization of a neural network in action || link

Assignments and Expected Effort

Grade breakdown: 4 HW sets: 40% + Midterm: 20% + Final project: 40%

Students will have opportunities to regain the points they have lost on a HW set through a combination of optional problems and progressive grading practices. Each HW set will also include a writing component, where students will reflect on the social implications of deep learning in a short essay (~500 words).

For the final project, 2-3 students will form a group and apply a subset of the techniques they have learned in the course to a realistic data set. I will provide a list of potential project ideas and corresponding open data resources, but creative, self-motivated projects are always encouraged. Based on student feedback, I may list graduate student 'mentors' from other departments if there is an interest in working with non-physics data sets. Another potential route might be to create visualizations of a neural network in action and connecting them to physical concepts in order to understand their dynamics.

Following is the breakdown of the number of hours the students will devote to the course over the course of a (truncated) semester, or 12 weeks:

Classroom hours: 3 hours/week (36 hours)
Reading: 4 hours/week (48 hours)
Bi-weekly HWs: 5 hours/week (60 hours)
Final project: 25 hours

Total: ~170 hours

Accessibility and Accommodations

Brown University is committed to full inclusion of all students. Please inform me early in the term if you have a disability or other conditions that might require accommodations or modification of any of these course procedures. You may speak with me after class or during office hours.

For more information, please contact Student and Employee Accessibility Services at 401-863-9588 or SEAS@brown.edu. Students in need of short-term academic advice or support are encouraged to contact one of the deans in the Dean of the College office. Students seeking psychological support services should contact Counseling and Psychological Services.