Statistical Models of the Brain

36-759 (CMU), MATH 3375 (Pitt) Fall, 2017

Schedule: WF 1:30-2:50
CNBC Classroom, MI 130
First class: Aug 30

Instructors: Brent Doiron and Rob Kass
bdoiron@pitt.edu
kass@stat.cmu.edu

Course Website: Hosted by Blackboard

In 2016 the two of us decided to merge the version of *Statistical Models of the Brain*, taught previously by Rob Kass, with *Computational Neuroscience*, taught previously by Brent Doiron. Our primary motivation was to create a course for a broad range of CNBC graduate students that would represent computational neuroscience more accurately than did either of the two predecessors. We also felt that such a course could do double-duty as the first semester of a two-semester sequence for the much smaller group of computationally-oriented CNBC students. Our experience in 2016 has suggested some modifications of the course, but we are generally happy with the conception and nearly all of the topics.

As we prepared the 2016 course, one thing we came to appreciate involves the term *statistical*: as we will explain, the two of us use the word in similar ways but with different emphases. Statistical ideas have been part of neurophysiology since the first probabilistic descriptions of spike trains, and the quantal hypothesis of neurotransmitter release, more than 50 years ago; they have been part of experimental psychology even longer. Throughout the field of statistics, models incorporating random "noise" components are used as an effective vehicle for data analysis. In neuroscience, however, the models also help form a conceptual framework for understanding neural function. In broad stroke, this course will examine a few of the most important methods and claims that have come from applying statistical thinking to the brain. However, some of the topics involve tools typically taught in statistics courses, while other topics involve tools taught in math courses. Most topics will involve modeling of both neural activity, in the sense of neurophysiology, and human behavior. Even at an intuitive level, we are unable to provide a comprehensive view of computational neuroscience; the field is too broad. Instead, we hope that by studying a series of examples, many of them very influential, students will come away with a sense of the way that computational methods contribute to contemporary understanding of neuroscience.

Course Structure and Logistics

Within the course we will provide rapid overviews of these mathematical and statistical methods, as "background" lectures. The material in these background lectures will be covered very quickly and, because we are keenly aware that many biologically-oriented students will be unable to thoroughly understand the methods, we will not require such students to make explicit use of the details. Students must identify themselves as either computational (for instance, if they are getting their Ph.D. in computer science, math, statistics, machine learning, neural computation, or engineering) or non-computational, and we will have different expectations and requirements for these two groups. We expect computational students to know, or study, all of the background methods in full detail, aiming at mastery; to do well on assignments, non-computational students need only grasp the main intuitions. (Rob has used this system successfully in the past.) In addition to the lectures, and class discussions, the course will involve (i) readings, (ii) student commentary (often, asking questions) on readings, which will help generate class discussions, (iii) assigned short-answer questions on readings, (iv) modeling homework, and (v) *for computational students only,* a project. There is no exam. Grades will be based primarily on student commentary (which requires thoughtful engagement with the readings) and short-answer questions (which will be aimed at pulling out the biggest points from the readings and lectures). Specifically, 75% of the grade for non-computational students will be based on commentary and short-answer questions, with the remaining 25% based on modeling homework; 55% of the grade for computational students will be based on commentary and short-answer questions, 20% on the project, and the remaining 25% based on modeling homework. The grading of homework may, at first, sometimes feel severe, but students will have the opportunity to improve their write-ups of homework assignments after the initial feedback (with a small penalty). *However*, students *may be penalized* for handing in assignments late. Computational students *will be penalized* for failing to propose their project on time (see below). The TA for the course is Josue Orellana, juo@andrew.cmu.edu.

A FEW DETAILS:

- The course will be run through the CMU hosting of 36-759 on Blackboard, see `http://www.cmu.edu/blackboard`. All registered students should already have access to this course, including those registered in Pitt MATH 3375.

- Commentary on readings must be posted on the appropriate discussion forum **no later than 10am on the assigned day of class**. Only the instructors will have access to these until some selection of them is made public shortly before the class occurs.

- Short-answer questions, based on readings (and discussion of readings), will require students to submit an answer of roughly 1 to 3 sentences in length. These will be managed and self-graded, with random spot-checks, using a Blackboard tool. The short answer questions for *both* of the 2 classes on a given week will be due by *Tuesday at 10am* the following week.

- Each of the 4 homework assignments will be due 2 weeks after it is assigned, and must be submitted through Blackboard as a pdf file. After the assignment is graded and returned, students will have 1 additional week to hand in a revision, if they choose to do so, in order to improve their score—students are eligible for a bonus of up to 85% of the points originally deducted for that assignment. Students are welcome to seek help from other students, or the TA, but all write-ups have to be done independently.

- Because much of the course will move very fast, students should try to read ahead in the background material when possible.

- Projects (for computational students, working in small teams) will be handed in as narrated slides, in PowerPoint or Keynote. These voice-over recorded presentations must run between 10 and 15 minutes, in total. All computational students must attend the presentation session, which will be on Wednesday, December 13, and must send their presentation to the TA by 5:00pm December 12. At the session, the presentation will be played and the students will answer questions. (All students are invited.) *Computational students must have their project* *approved by the instructors no later than October 20.* To get approval, students must submit a proposal by email message to the TA (no other document is required) that includes (1) the team of 2 or 3 people who will do the project (all team members must submit their own email); (2) the topic, described in several sentences including a reference to a paper; (3) a statement as to whether any new simulations or other work is involved (not necessary, but might be desirable for some computational students); and (4) what work each student will be responsible for—all students are responsible for the whole finished product, but, for example, only 1 student typically will record the narration.

A key text for statistical tools is *Analysis of Neural Data*, Kass, Eden, and Brown (KEB), published by Springer. Information about the book is at `http://www.stat.cmu.edu/~kass/KEB/index.html`. *NOTE:* a pdf version of the book is free for both CMU and Pitt students.

Students who have weak backgrounds in neurophysiology should find a basic source on neurons and read it. (Such students will have some extra time during the beginning of the semester to do this too, while we are covering basic statistical ideas.) We recommend the first 5 chapters of Bear, Connors, and Paradiso *Neuroscience: Exploring the Brain*, which assumes only high-school biology.

Accommodations for Students with Disabilities

If you have a disability and have an accommodations letter from the Disability Resources office, we encourage you to discuss your accommodations and needs with one of the instructors as early in the semester as possible. We will work with you to ensure that accommodations are provided as appropriate. If you suspect that you may have a disability and would benefit from accommodations but are not yet registered with the Office of Disability Resources, we encourage CMU students to contact them at access@andrew.cmu.edu. Pitt students should contact Disability Resources and Services (DRS), 216 William Pitt Union, (412) 648-7890/(412) 383-7355 (TTY).

Support for Health and Well-being

Take care of yourself. Do your best to maintain a healthy lifestyle this semester by eating well, exercising, avoiding drugs and alcohol, getting enough sleep and taking some time to relax. This will help you achieve your goals and cope with stress. All of us benefit from support during times of struggle. There are many helpful resources available on campus and an important part of the college experience is learning how to ask for help. Asking for support sooner rather than later is almost always helpful. If you or anyone you know experiences any academic stress, difficult life events, or feelings like anxiety or depression, we strongly encourage you to seek support. At CMU, Counseling and Psychological Services (CaPS) is here to help: call 412-268-2922 and visit their website at `http://www.cmu.edu/counseling/`. Consider reaching out to a friend, faculty or family member you trust for help getting connected to the support that can help.

Topics and Readings

Aug 30 BD+RK[1]. Overview: Class structure, including readings, questions, comments, and homework.

Required reading: Kass, Eden, Brown (KEB), Chapter 1; Section 1 (Introduction) in Kass, R.E., ..., Doiron, B., and 23 others (2018) Computational neuroscience: Mathematical and statistical perspectives, *Ann. Rev. Statist.*, to appear; Goldman, M.S. and Fee, M.S. (2017) Computational training for the next generation of neuroscientists, *Current Opinion Neurobiol.*, 46: 25-30.

Sep 1 RK. *Background:* log transformations; random variables and random vectors.

Required, with comment, for non-computational students: KEB Chapter 2; Secs 3.1-3.2; Secs 4.1-4.2.
(*Note serious typo on p. 84:* Following "In the discrete case we have" the quantity $P(X^{(1)} = x)$ should be $P(Y^{(1)} = y)$.)

Sep 6 RK. *Background:* Important probability distributions and the way they model variation in data; least-squares linear regression and the linear algebra concept of a basis.

Background offline: Matlab and R, part 1

Required, with comment, for non-computational students: KEB Secs 4.3.1; 5.1-5.3; 5.4.1-5.4.3; 12.5 through 12.5.1; appendices A.7 and A.9; 12.5.3 through equation (12.5.7) on p. 342.

Sep 8 Josue Orellana. *Background:* Bayes' Theorem and the optimality of Bayes classifiers; the Law of Large Numbers and the Central Limit Theorem; statistical estimation.

Required, with comment for non-computational students (recommended for all): KEB Secs 4.3.3-4.3.4, through p. 101; 6.1; 6.2.1; 6.3.1; 7.1-7.2.

Sep 13 RK. Random walk models of integrate-and-fire neurons.

[1]BD and RK refer to which instructor will present the lecture. In the case of BD+RK both Brent and Rob will present.

Required background reading: KEB Sec 5.4.6 and Chapter 19 through 19.2.1.

Required, with comment: Shadlen, M.N. and Newsome, W.T. (1998) The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *J. Neurosci.*, 18: 3870–3896. *Required up to Section 2, p. 3877, and concluding remarks*

Homework 1: Simulating random walks and the role of noise in signal transfer

Sep 15 BD. *Background:* Primer on differential equations; introduction to numerical methods with MATLAB.

Required, with comment for non-computational students (recommended for all): Chapter 7 - Moler (2004) Ordinary differential equations; Numerical Computing with MATLAB.

`https://www.mathworks.com/moler/odes.pdf`

Sep 20 BD. Electrical circuit model of a neuron. Passive synaptic dynamics and phenomenological models of spiking: integrate-and-fire dynamics.

Required, with comment: Ermentrout and Terman (2010) *Mathematical Foundations of Neuroscience*, Springer. Secs 1.1-1.5 (an electronic version of this book is freely available to all Pitt and CMU students).

Sep 22 BD. The Hodgkin-Huxley model of action potential generation.

Required, with comment: Ermentrout and Terman (2010) *Mathematical Foundations of Neuroscience*, Springer. Secs 1.7-1.10.

Homework 2: Spike trains using the H-H model

Sep 27 RK. *Background:* Optimality, Bayesian inference and Maximum Likelihood Estimation.

Optimal observers in perception and action.

Required background reading, with comment optional: KEB Secs 7.3.9; 8.1-8.2; 8.3.1-8.3.3; 8.4.3 (Example 5.5 and Figure 8.9); Chapter 16 through equation (16.18) on p. 449.

Required, with comment: Körding, K.P. and Wolpert, D.M. (2004) Bayesian integration in sensorimotor learning, *Nature*, 427: 244–247.

Sep 29 RK. Cognition and optimality; ACT-R.

Required background reading: KEB, pp. 102-103, through Example 4.9.

Required, with comment: Anderson (2007) *How Can the Human Mind Occur in the Physical Universe?*, Chapter 1.

Oct 4 RK. *Background: Regression and generalized regression.*

Background offline: R, part 2

Required, with comment, for non-computational students (recommended for all): KEB Secs 12.5.4-12.5.5; 12.5.8; Chapter 14 through 14.1 (can skip 14.1.2, 14.1.5); 15.2 through 15.2.4.

Oct 6 BD. The mechanics of neuronal variability

Required, with comment: Doiron, B., Litwin-Kumar, A. Rosenbaum, R. Ocker, G. and Josic, K. (2016) The mechanics of state dependent neural correlations. *Nature Neuroscience* 19, 383-393.

Oct 11 RK. ROC curves in human discrimination tasks; ROC curves in neural coding.

Comment is required on some part of the background reading AND/OR the excerpts specified from the Swets paper and Dayan and Abbott book:

KEB Chapter 10 up to the beginning of Sec 10.1.1 (p. 249); Secs 10.2; 10.3.1; 10.3.4-10.3.5; 10.4.1; 10.4.3-10.4.4, especially Figure 10.3; 11.2.1 through p. 298.

Much of this is not very relevant; concentrate on Figures 3 and 6: Swets, J.A. (1986) Form of empirical ROCs in discrimination and diagnostic tasks: Implications for theory and measurement of performance, *Psychol. Bull.*, 99: 181–198.

Dayan, P. and Abbott, L.F. (2001) *Theoretical Neuroscience*, pp. 89-95. This is based on Britten, K.H., Shadlen, M.N., Newsome, W.T., and Movshon, J.A. (1992) The analysis of visual motion: a comparison of neuronal and psychophysical performance, *J. Neurosci.*, 12: 4745-4765.

Oct 13 RK. Firing rate and neural coding; spike trains as point processes.

Required, with comment for non-computational students: KEB Example 14.5, pp. 410-411; Chapter 19 through 19.2.2.

Required, with comment for computational students: KEB Example 14.5, pp. 410-411; Chapter 19.

Homework 3: generalized linear models for psychophysical behavior and hippocampal place cells

Oct 18 RK. Population decoding.

Required, with comment optional: Georgopoulos, A.P., Lurito, J.T., Petrides, M., Schwartz, A.B., and Massey, J.T. (1989) Mental rotations of the neuronal population vector, *Science*, 243: 234–236.

Required, with comment: Black, M.J. and Donoghue, J.P. (2007) Probabilistically modeling and decoding neural population activity in motor cortex, in G. Dornhege, J. del R. Millan, T. Hinterberger, D. McFarland, K.-R. Muller (eds.), *Toward Brain-Computer Interfacing*, MIT Press, pp. 147–159.

Oct 20 Josue Orellana. Population-wide variability: spike count correlations.

Required background reading: KEB Section 4.3.2; Example 6.1, p. 141.

Required, with comment: Averback, B.B., Latham, P.E., and Pouget, A. (2006) Neural correlations, population coding, and computation, *Nature Review Neurosci.*, 7: 358-366.

Oct 25 RK. Information theory in human discrimination.

Background: return to KEB Section 4.3.2, comments about entropy and channel capacity, pp. 95-97, including Example 4.5.

Required, with comment: Miller, G.A. (1956) The magical number seven, plus or minus two, *Psychol. Rev.*, 63: 343-355.

Oct 27 RK. Information theory in neural coding.

Required background reading: KEB, Example 4.6.

Required, with comment: Nirenberg, S., Carcieri, S.M., Jacobs, A.L. and Latham, P.E. (2001) Retinal ganglion cells act largely as independent encoders, *Nature*, *411:* 698–701.

Optional reading: Rieke, F., Warland, D., de Ruyter van Steveninck, R., Bialek, W. (1997) *Spikes: Exploring the Neural Code*, MIT Press. *Read pages 101–113, 148–156.*

Nov 1 RK. Rate coding versus temporal coding in the retina.

Required, with comment: Jacobs, A.L., Fridman, G., Douglas, R.M., Alam, N.M., Latham, P.E., Prusky, G.T., and Nirenberg, S. (2009) Ruling out and ruling in neural codes, *Proc. Nat. Acad. Sci.*, 106: 5936–5941.

Nov 3 Josue Orellana. Neural implementation of Bayesian inference.

Background: KEB, Section 14.1.6.

Required, with comment: Ma, W.J., Beck, J.M., Latham, P.E., and Pouget, A. (2006) Bayesian inference with probabilistic population codes, *Nature Neurosci.*, 9: 1432–1438.

Recommended: Salinas, E. (2006) Noisy neurons can certainly compute, *Nature Neurosci.*, 9: 1349–1350.

Nov 8 BD. *Background:* Dynamical systems and qualitative analysis of nonlinear systems.

Required, with comment: Strogatz (1994) *Nonlinear Dynamics and Chaos*, Westview Press, Secs 5.0-5.2; 6.0-6.5

`https://westviewpress.com/books/nonlinear-dynamics-and-chaos/`

Nov 17 RK. Neural basis of decision making.

Required background reading: KEB, Section 11.1.5.

Required, with comment: Gold and Shadlen (2007) The neural basis of decision-making, *Ann. Rev. Neuroscience*, 30: 535-574.

Nov 29 BD. Firing-rate models: The inhibitory stabilized cortical network.

Required, with comment: Ozeki, Finn, Schaffer, Miller, Ferster (2009) Inhibitory stabilization of the cortical network underlies visual surround suppression. *Neuron*, 578-592.

Homework 4: Binocular Rivalry in population models.

Dec 1 BD. Network models of working memory and decision-making.

Required, with comment: Machens, C.K., Romo, R. Brody, C.D. (2005) Flexible control of mutual inhibition: a neural model of two-interval discrimination *Science*, 307: 1121–1124.

Recommended: Polk, A., Litwin-Kumar, A. and Doiron, B. (2012) Correlated neural variability in persistent state networks *PNAS*, 109: 6295–6300.

Dec 6 BD. The Hopfield model.

Required, with comment: Hopfield, J. (1982) Neural networks and physical systems with emergent collective computational abilities, *PNAS* 79: 2554–2558.

Recommended: Hertz, Krogh, Palmer *Introduction to the Theory of Neural Computation*, Westview press, Chapter 2.

Dec 8 RK. Reinforcement learning.

Required, with comment: Glimcher, P. (2011) Understanding dopamine and reinforcement learning: The dopamine reeward prediction error hypothesis, *PNAS*, 108: 15647–15654 (with corrections, pp. 17568–17569).

Recommended: Y Niv (2009) Reinforcement learning in the brain *J. Math. Psychol.*, 53: 139–154.