Statistical Models of the Brain

36-759 (CMU), MATH 3375 (Pitt) Fall, 2019

Schedule: WF 1:30-2:50
CNBC Classroom, MI 130
First class: Aug 28

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

Course Website: Hosted by Canvas

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 experiences subsequently have suggested some modifications of the course, but we are generally happy with the conception.

Working together on this, 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 several of the most important methods and claims that have come from applying statistical thinking and modeling to the brain. However, some of the topics use tools typically taught in statistics courses, while other topics use tools taught in math courses. Topics will involve modeling of neural activity in the sense of neurophysiology, neuroimaging, and human behavior; students will be exposed to some of each.

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 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 some 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; on the other hand, to do well on assignments, non-computational students need only grasp the main intuitions.

In addition to the lectures, and class discussions, the course will involve (i) readings, (ii) student commentary (often, asking questions) on readings, (iii) assigned short-answer questions (SAQs) on readings, (iv) for computational students only, a mid-term homework assignment, and (v) a set of presentations by groups of students. 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 the project; 55% of the grade for computational students will be based on commentary and short-answer questions, 20% on the homework, and the remaining 25% based on the project. *Please note:* all deadlines *must* be met: students *will be penalized (possibly severely)* for failing to hand assignments in on time, or for failing to propose their project on time (see below). The TA for the course is Shenghao Wu shenghaw@andrew.cmu.edu.

The course is heavy on readings. We hope that students will spend the time it takes to digest each assigned article thoroughly. However, knowing that time is limited, we *require* only that students (a) post a cogent comment or question on the discussion board and (b) answer the SAQs.

A few details:

- The course will be run through the CMU hosting of 36-759 on Canvas, see `https://canvas.cmu.edu`. All registered students should already have access to this course, including those registered in Pitt MATH 3375.
- Comments on readings must be posted on the appropriate discussion forum **no later than 10am on the assigned day of class**. Students will have access to commentary by others only *after* they post themselves. The instructors will read these posts prior to class, and use them to guide the lecture overview.

  Comments are meant to demonstrate engagement with the material, and will be graded on a 0/1/2 basis, with 1 signifying a minimal response. Comments may consist entirely of questions identifying points not yet clear to the student. In our experience there is a lot of variation in length, but typically a few sentences will suffice. Here are 4 examples of student comments on one of the readings[1]:

  - The building, computer, brain analogy is very instructive. It's interesting to see the shift in perspective where before the trend was to think of the brain as like a computer whereas now the trend is to make a computer operate like the brain. The explanation of three shortcuts made the concept of the cognitive architecture easy to grasp. The modular break up of ACT-R was very informative. The results shown in Figure 1.6 are impressive. I didn't quite catch what figure 1.7 is trying to show. [**SCORE: 2**]
  - Anderson presents a rather attractive metaphor for how he sees it best to approach understanding the brain, one that could be well

[1] Previously we scored 0/1, with 1 for writing anything; the scores here are retrospective, for illustration.

summed up as, "the whole is greater than merely the sum of its parts." That idea that you can't simply deconstruct ad infinitum in one direction and work your way back to the other side seems deeply sobering.

  Taken to its logical conclusion though, I wonder whether if in accepting what could be perceived as Anderson's principal conclusions, one must also find it unsatisfying as it might be that the best that can be achieved is a model of our cognitive architecture, which can only be refined and improved, but that never quite gets there. [**SCORE: 2**]

  - I thought the Anderson chapter was really interesting and easy to read. The example in Figure 1.8 (using module behavior to predict BOLD response) was particularly interesting and really pulled together the concepts of ACT-R and how we can use it to understand brain function. [**SCORE: 2**]
  - Not convinced... too philosophical to be science [**SCORE: 1**]

- Each SAQ 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 the Canvas quiz tool. The SAQs **MUST** be answered by each student working *independently*, and they **MUST** be answered within a specified 48 hour window. Students will be notified when the window opens. The instructors will, in addition, inform students of the learning objectives most relevant to the SAQs *prior to the relevant lecture*. This will help guide students in reading.

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

- The instructors will post the lecture slides in the Slides module of Canvas. Check for them the evening before class, but be aware that there may be edits up until class time.

- Projects by students, working in teams of 3 students (with occasional exceptions in size of team) will be handed in as *narrated slides*, in PowerPoint. These voice-over recorded presentations must run between 10 and 15 minutes, in total. The subject of the project should be a summary of 1 or more papers. All students must attend the presentation sessions, which will be on December 11 and 13, and **must send their presentation to the TA by 5:00pm December 10**. At the session, the presentation will be played and the students will very briefly answer questions. *Students must have their project approved by the instructors no later than Friday October 25.* 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 3 people who will do the project (all team members must submit their own email); (2) the topic, described in several sentences including reference to the paper or papers that will be discussed; and (3) 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 28 RK[2]. Overview: Class structure, including readings, questions, comments, and homework.

Required reading: Kass, Eden, Brown (KEB), Chapter 1, **and** Section 1 (Introduction) in Kass, R.E., ..., Doiron, B., and 23 others (2018) Computational neuroscience: Mathematical and statistical perspectives, *Ann. Rev. Statist. Appl.*, 5: 183-214.

Aug 30 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 4 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.

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 (12.57) on p. 342.

Background offline: Matlab and R

Sep 6 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 11 RK. *Background:* Bayes' Theorem and the optimality of Bayes classifiers; the Law of Large Numbers and the Central Limit Theorem; statistical estimation.

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

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; 7.3.9.

Sep 13 RK+BD. Random walk models of integrate-and-fire neurons; effects of noise: balanced excitation and inhibition.

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*

Required, with comment: Stein, R.B., Gossen, E.R., and Jones, K.E. (2005) Neuronal variability: noise or part of the signal? *Nat. Rev. Neuro.*, 6:389-397. *Required through Figure 2, p. 391.*

Sep 18 RK. Population vectors.

All of the following are required, but the comment may be on any or all:

KEB, Sec 12.5.4 (especially Example 12.6).

Section 1 (Introduction, pp. 2021-2022) of Orellana, J., Rodu, J., and Kass, R.E. (2017) Population vectors can provide near optimal integration of information, *Neural Comput.*, 29: 2021-2029.

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.

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.

Sep 20 RK. Information theory in human discrimination.

Background: 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.

Sep 25 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 27 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.

Oct 2 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/`

Oct 4 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.

Oct 9 RK. *Background:* Maximum likelihood; statistical tests.

Required, with comment for non-computational students (recommended for all): KEB, Secs 8.1-8.2; 8.3.1-8.3.3; 8.4.3 (Example 5.5 and Figure 8.9); Chapter 10 up to the beginning of Sec 10.1.1 (p. 249); Secs 10.4.1; 10.4.3-10.4.4, especially Figure 10.3.

Oct 11 RK. *Background: Regression and generalized regression.*

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.

SAQ1 Available October 13, due October 15

Oct 16 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.

Recommended for computational students: Chen, Y., Xin, Q., Ventura, V., and Kass, R.E. (2018) Stability of point process spiking neuron models, *J. Comput. Neurosci.*, published online (see especially Figure 8c).

NOTE Project proposals must be approved before October 25.

Oct 18 RK. Information theory in neural coding.

Required background reading: KEB, Example 4.6.

Required (background for Jacobs et al.*):* 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.

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.

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.*

HOMEWORK *For computational students*: handed out October 21, due October 28.

Oct 23 RK. Optimal observers in perception and action; neural implementation of Bayesian inference.

Required background reading for everyone, with comment optional: Chapter 16 through equation (16.18) on p. 449.

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

Required for everyone, with comment optional: Salinas, E. (2006) Noisy neurons can certainly compute, *Nature Neurosci.*, 9: 1349–1350.

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

Required for computational students, with comment optional: Orellana, J., Rodu, J., and Kass, R.E. (2017) Population vectors can provide near optimal integration of information, *Neural Comput.*, 29: 2021-2029.

Oct 25 RK. Neural basis of decision making.

Required background reading: KEB, Section 11.1.5 and the discussion of SDT in Section 10.4.4.

Required, only through the discussion of Figure 5, with comment: Gold and Shadlen (2007) The neural basis of decision-making, *Ann. Rev. Neuroscience*, 30: 535-574.

Oct 30 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.

Nov 1 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.

SAQ2 Available Nov 3, due November 5

Nov 6 **NO CLASS**

Nov 8 BD. Synaptic plasticity.

Required, with comment: Abbott, LF and Nelson, SB. (2000) Synaptic plasticity: taming the beast. *Nature Neurosci.* 3: 1178-1183.

Nov 13 BD. Network dynamics.

Required, with comment: Vogels, TP; Rajan, K; Abbott, LF. (2005) Neural network dynamics. *Ann. Rev. Neurosci.* 28: 357–376.

Required, with comment: Potjans, TC; Diesmann, M. (2012) The cell-type specific cortical microcircuit: relating structure and activity in a full-scale spiking network model. *Cerebral Cortex*, 24:785–806.

Nov 15 RK. Population-wide variability: spike count correlations; dimensionality reduction.

Required background reading: 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 Reviews Neurosci.*, 7: 358-366.

Required, only through p. 1504, up to "Selecting a dimensionality reduction method," with comment: Cunningham, J.P. and Yu, B.M. (2014) Dimensionality reduction for large-scale neural recordings, *Nat. Neurosci.*, 17: 1500-1509.

Nov 20 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.

Nov 22 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.

Nov 27 and 29 **NO CLASSES (Thanksgiving)**

Dec 4 RK. Graphs and networks.

Required, with comment: Bassett, D.S., Zurn, P., and Gold, J.I. (2018) On the nature and use of models in network neuroscience, *Nature Reviews Neurosci.*, 19:566-578. *Required up until "Density of study in this 3D space," p. 571.*

Required, with comment: Bau, G.L., ..., Bassett, D.L., and Satterthwaite, T.D. (2017) Modular segregation of structural brain networks supports the development of executive function in youth, *Current Biol.*, 27: 1561-1572.

Dec 6 RK. Deep learning.

Required, with comment: Kriegeskorte, N. (2015) Deep neural networks: A new framework for modeling biological vision and brain information processing, *Ann. Rev. Vision Sci.*, 1:417-446.

Recommended: Yamins, *et al.* (2014) Performance-optimized hierarchical models predict neural responses in higher visual cortex, *Proc. Nat. Acad. Sci.*, 111: 8619-8624.

SAQ3 Available December 8, due December 10

Dec 11 and 13 BD+RK **PRESENTATIONS**