

Statistical Models of the Brain

36-759 (CMU)

Spring, 2023

Schedule: TTh 2:00-3:20, MI 130 First class: January 18	Instructor: Rob Kass kass@stat.cmu.edu
	TA: Qi Xin qix@andrew.cmu.edu
Website: Hosted by Canvas	

In 2016 Brent Doiron and I decided to merge my course *Statistical Models of the Brain*, which I'd taught on several previous occasions (beginning in Spring 2011), with his course *Computational Neuroscience*. 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. Our experiences subsequently suggested some modifications of the course, but we were happy with the conception. Because Brent has left Pittsburgh, I reduced somewhat the content he had been teaching, and I also asked several guest lecturers to treat key topics. In addition, in 2021 the course moved to a “flipped classroom” format, in which students are required to watch a short video summary of the material in conjunction with each reading. Since roughly 2015 students have been required to submit a comment or question (a blog post) several hours before each class meeting. Grades are based on the submitted comments, together with class discussion, and answers to a set of short-answer questions. There are no exams.

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, a single course can not provide a comprehensive view of computational neuroscience; the field is too broad. Instead, I have a more modest

goal: I expect students, by studying a series of examples, to gain a sense of the way that computational methods contribute to contemporary understanding of neuroscience. The examples come from published papers, which, like most ideas in core course content, are generally pretty old. In the commentary blog posts, and in our class discussions, students are encouraged to contribute their own thoughts about the success of these papers and the extent to which the ideas should be updated. I would appreciate knowledgeable remarks about existing updates because I do not myself follow most of these areas sufficiently to be aware of them. By the way, in calling this goal “modest” I don’t mean to imply that it isn’t that important. My feeling is the opposite: familiarity with computational methods is essential for all neuroscientists who want to have a reasonably informed understanding of their subject.

A detailed list of topics and assigned readings is at the end of this syllabus.

Course Structure and Logistics

In addition to the lectures (which are almost all in the form of short videos), and class discussions, the course will involve (i) the readings, (ii) student commentary blog posts on readings and/or videos (often, asking questions), (iii) in-class discussion, and (iv) assigned short-answer questions (SAQs). Again, no exam. My primary motivation for this structure is the desire to make it as easy as possible for students to achieve the modest goal I just articulated.

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. Within the course, several days will be devoted to rapid overviews of mathematical and statistical methods, as background. These are aimed mainly at non-computational students, but computational students sometimes gain something of value from a new perspective on familiar material. The material in these background lectures will be covered very quickly, with the primary goal of supplying to non-computational students a conceptual understanding of the main points. Thus, non-computational students will not be expected to know or use the details. Computational students, on the other hand, should know, or study, all of the background methods in full detail, aiming at mastery. Some of the SAQs will be designated as being required only of computational students.

I will provide a video summary lecture on each topic ahead of class (with one class exception, the lecture given by Professor Jon Rubin) so that class time can be devoted to discussion of issues raised in the student commentaries, and anything else that might arise. As I write this syllabus initially, my website <http://www.stat.cmu.edu/~kass/>

`teaching.html#brain` contains videos from Spring 2022. Many of these will be used again this semester, but some will be revised. I will provide the current link to the required video in a message broadcast from Canvas prior to each class.

Grades will be based 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, video lectures, and class discussions. *Please note:* some of the SAQs *will* be based on class discussions.

Students who are unable to attend class may join by Zoom up to 5 times during the semester if they contact both the TA and me **at least 24 hours in advance.**

The course is heavy on readings. I hope that students will spend the time it takes to digest each assigned article thoroughly. However, knowing that time is limited, I *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>.
- Comments on readings **must** be posted on the appropriate discussion forum **no later than 9:30am on the assigned day of class**. Students will have access to commentary by others only *after* they post themselves. I 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 my 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 readings¹:

- 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

¹The scores here are retrospective, for illustration.

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 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 the relevant lecture about the 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 syllabus informs students of the primary learning objectives most relevant to the SAQs by indicating key sections to “pay attention to” (shorthand: “*attention*”). This will help guide students in reading. In addition, the SAQs will be handed out roughly 5 days prior to the opening of the window (except for the first SAQ, which is based only on the first substantive lecture).
- *Because much of the course will move very fast, students should try to read ahead when possible.*
- When I create a new version of a video lecture, I will make it available at least 2 days prior to the class meeting time. The relevant video lecture should be viewed prior to submitting a comment.

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/research.html#and](#). *NOTE*: a pdf version of the book is free for both CMU and Pitt students. Also, please check the extensive list of corrections at <http://www.stat.cmu.edu/~kass/KEB/corrections2021.pdf>. Students who have weak backgrounds in neurophysiology should find a basic source on neurons and read it. I 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/>. As of Fall 2022, CaPS is also offering services from TimelyCare, see <https://www.cmu.edu/wellbeing/resources/timely-care.html>. 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

NOTES: (1) Comments are required on all readings unless otherwise indicated. (2) The SAQs will be based on the material indicated as needing special attention (and class discussion). (3) Some details in this syllabus may change, so watch for updates. (4) Several videos concerning differential equation modeling of neurons and circuits were created by guests. (5) The URL for each video will be sent in a Canvas message, and most will be my webpage <http://www.stat.cmu.edu/~kass/teaching.html#brain>.

0. Jan 17 Overview: The nature of statistical models of the brain; as an example, Bayes' theorem and its uses; class structure, including readings, questions, comments, and homework.

1. Jan 19 What is computational neuroscience?

Required reading: Section 1 (Introduction) in Kass, R.E. and 24 others (2018) Computational neuroscience: Mathematical and statistical perspectives, *Ann. Rev. Statist. Appl.*, 5: 183-214.

Additional required video: https://www.youtube.com/watch?v=4iPeV_o5f9Y

Pay special attention: the brain-as-computer metaphor; Marr's three levels of analysis; tuning curve video.

SAQ1 Window: Jan 20-22.

2. Jan 24 Random variables; What is a statistical model? Fitting statistical models to data.

Required reading: Kass, Eden, Brown (KEB), Chapter 1, especially 1.2.1, 1.2.5, 1.2.6; Chapter 3 through Equation (3.1); Section 3.2 through 3.2.3 (reminder to see corrections).

Pay special attention: "Signal" and "noise" in Examples 1.4 and 1.5; Equation (1.4).

Computational students, in addition: Read the rest of Chapter 3, especially 3.2.4 (reminder to see corrections); pay attention to Figure 1.1; Sections 1.2.5, 1.2.6 (see also Section 8.1), 3.2.4.

3. Jan 26 *Background:* Log transformations; random vectors; important probability distributions and the way they model variation in data.

Reading: KEB Ch 2, esp. 2.2.1; Ch 4 through 4.2.2; 4.3.1 through Equation (4.26); 5.1-5.3; 5.4.2.

Attention: Secs 2.2.1, 5.2.1, 5.4.2; Figs 2.5, 2.6.

Comp students: read the rest of 4.3.1 and Ch 5; *Attention:* Secs 4.2.4, 5.5.

4. Jan 31 *Background:* The Law of Large Numbers and the Central Limit Theorem; statistical estimation; least-squares linear regression and the linear algebra concept of a basis.

Reading: KEB Ch 6 through 6.1.1; 6.2.1; 7.1, 7.2, 7.3.1; Introduction to Ch 12; 12.5 through 12.5.1; appendices A.7 and A.9; 12.5.3 through equation (12.57) on p. 342; 12.5.8.

Attention: 6.2.1; Fig 7.2; Introduction to 12.5 and 12.5.8; A.7; Fig 12.9 (which is the same as the bottom of Fig A.2).

Comp students: 6.1.2, 6.3.2; 12.5.5, 12.5.7; A.8; attention to 12.5.7. Secs 7.3.8, 7.3.9 are recommended.

5. Feb 2 Random walk models of integrate-and-fire neurons; effects of noise: balanced excitation and inhibition.

Readings: KEB Sec 5.4.6; Introduction to Ch 19.

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. *Up to Section 2, p. 3877, and concluding remarks*

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. *Only Figure 1 and Figure 2, pp. 390-391.* The histograms are explained by this statement on p. 392, “The ability of the neuron to transmit signals faithfully is only evident after analysing many cycles of the stimulus. However, transmission by a population of neurons, rather than a single neuron, would allow the signal to be evident in real time.” See my video lecture for a bit more on this.

Attention: Shadlen and Newsome, “price of dynamic range is noise,” Fig 2; Stein et al., Fig 2.

Comp students: KEB Secs 19.1-19.2; attention to the theorem in 19.2.1.

6. Feb 7 Population vectors.

Readings: KEB, Sec 12.5.4.

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.

Attention: KEB, Example 12.6; Figure 2 of Georgopoulos et al.; Equation (5) of Black and Donoghue.

7. Feb 9 Information theory in human discrimination.

Background Reading: KEB Section 4.3.2, especially comments about entropy and channel capacity, pp. 95-97, including Examples 4.5 and 4.6.

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

Attention: KEB Example 4.6; Miller, Figure 2.

8. Feb 14 *Background:* Ideas in differential equations and dynamical systems.

Video: https://www.youtube.com/watch?v=p_di4Zn4wz4

Here are some questions that should help you better understand the video:

Q1. Under what circumstances are differential equations used?

Q2. What is the order of a differential equation?

Q3. What is a phase space?

Q4. What is an attracting state?

9. Feb 16 Electrical circuit model of a neuron. Passive synaptic dynamics and phenomenological models of spiking: integrate-and-fire dynamics.

Reading: 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).

Attention: Nernst equation and how it differs from GHK equation; time constant of RC model.

SAQ2 Window: Feb 17-19.

10. Feb 21 The Hodgkin-Huxley model of action potential generation.

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

Attention: voltage clamp; distinction between sodium and potassium conductances.

11. Feb 23 *Background:* Bayes' Theorem; optimality of Bayesian classifiers; mean squared error; Bayes and maximum likelihood.
Reading: KEB Secs 4.3.3-4.3.4 through p. 101; 8.1-8.2; 8.3.3.
Attention: Theorem on p. 182; Equation (8.10); Figure 8.8.
12. Feb 28 *Background:* statistical tests, ROC curves, signal detection theory.
Reading: KEB 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.
Comp students: The rest of Chapter 10.
Attention: Figures 10.3 and 10.4.
13. Mar 2 Neural basis of decision making.
Background reading: KEB, Section 11.1.5 and the discussion of SDT in Section 10.4.4.
Reading: Gold and Shadlen (2007) The neural basis of decision-making, *Ann. Rev. Neuroscience*, 30: 535-574, *only through* the discussion of Figure 5.
Attention: Figures 4 and 5c.
- Mar 7,9 NO CLASSES (CMU Spring Break)
14. Mar 14 Cognition and optimality; ACT-R.
Background reading: KEB, pp. 102-103, through Example 4.9.
Reading: Anderson (2007) *How Can the Human Mind Occur in the Physical Universe?*, Chapter 1.
Attention: Anderson's three "shortcuts"
15. Mar 16 Optimal observers in perception and action.
Background reading: KEB Chapter 16 through equation (16.18) on p. 449, especially Example 16.1; see also Example 8.1.
Comp students: The rest of Chap 16.
Reading: Körding, K.P. and Wolpert, D.M. (2004) Bayesian integration in sensorimotor learning, *Nature*, 427: 244-247.
Attention: KEB, Equations (16.11), (16.12); Körding and Wolpert Figure 2.
16. Mar 21 *Background:* Regression and generalized regression.
Reading: KEB Chapter 14 through 14.1 (can skip 14.1.2, 14.1.5); 15.2 through 15.2.4.

17. Mar 23 Firing rate and neural coding; spike trains as point processes.

Reading: KEB Example 14.5, pp. 410-411; Chapter 19 through page 569.

Comp students: check the rest of Ch 19, and read what interests you; read Sections 1 and 2, and Figures 6 and 9, of Weber and Pillow (2017).

Weber, A.I. and Pillow, J.W. (2017) Capturing the dynamical repertoire of single neurons with generalized linear models, *Neural Comput.*, 29: 3260-3289.

Also of potential interest: The interpretation of Fig 8c in Chen, Y., Xin, Q., Ventura, V., and Kass, R.E. (2018) Stability of point process spiking neuron models, *J. Comput. Neurosci.*, 46:19-42.

The first three paragraphs about point processes (Sec 3.6) in the draft paper labelled RevSMBRAIN-PP in the Readings, and also the first figure at the beginning.

Attention: the 3 types of point processes identified in the lecture as needed for the next reading.

Comp students: KEB Figure 19.9 and Weber and Pillow Figure 9.

18. Mar 28 Information theory in neural coding.

Background reading: KEB, Example 4.6.

Readings: 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.

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.

Attention: Figure 3 of Nirenberg *et al.* and Figure 2 of Jacobs *et al.*

Recommended, especially for comp students: 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.* (See Readings)

19. Mar 30 Oscillations and coordinated neural activity.

Background reading: KEB, Chapter 18 through Section 18.2.2 and Section 18.5 through 18.5.1.

Also recommended background: Sections 4.1 and 4.4 of the paper “Oscillating neural circuits: phase, amplitude, and the complex normal distribution,” which is UrbanEtAl in the readings file; and it may help to review the video <https://youtu.be/L6GUocwq3Nk>, minutes 3:00-8:45, concerning Parseval’s theorem and KEB Example 2.2.

Reading: Fries, P. (2005) A mechanism for cognitive dynamics: neuronal communication through neuronal coherence, *TRENDS in Cognitive Sciences*, 9: 474-480.

Attention: Figure 3a.

Also recommended is a 2 minute video by Bob Desimone, MIT: https://www.youtube.com/watch?v=b12aYFv_978&authuser=0

Comp students: The remainder of KEB Sections 18.2 and 18.5 and also read 18.3.

SAQ3 Window: Mar 31-Apr 2.

20. Apr 4 Jon Rubin. Network dynamics.

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

Attention: From network perspective, advantages of ongoing activity due to balanced excitation and inhibition (as in Figure 3 d,e).

21. Apr 6 Chencheng Huang. Network models of decision-making.

Reading: Wang, X.-J. (2008) Decision making in recurrent neuronal circuits, *Neuron*, 60: 215-234, through the section “Recurrent Cortical Circuit Mechanism,” which ends on p. 223.

22. Apr 11 Population-wide variability: spike count correlations; dimensionality reduction.

Background reading: KEB, Example 6.1, p. 141 (review), and Section 17.3.1, especially the examples.

Readings: Averback, B.B., Latham, P.E., and Pouget, A. (2006) Neural correlations, population coding, and computation, *Nature Reviews Neurosci.*, 7: 358-366, *only through* page 360.

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

Attention: Figure 1 of Averbeck, Latham, Yu and Figures 1 and 2 of Cunningham and Yu.

Apr 13 NO CLASS (CMU Break)

23. Apr 18 Reinforcement learning.

Reading: Glimcher, P. (2011) Understanding dopamine and reinforcement learning: The dopamine reward prediction error hypothesis, *PNAS*, 108: 15647–15654 (with corrections, pp. 17568–17569), through the interpretation of Figure 3.

Attention: Figures 2 and 3.

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

24. Apr 20 Artificial neural networks (ANNs).

Reading: Richards, B.A., and many others, Körding, K.P. (2019) A deep learning framework for neuroscience, *Nat. Neurosci.*, 22: 1761-1770.

Attention: The three core component of ANN design, as in Figure 1.

25. Apr 25 Graphs and networks.

Reading: 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.

Attention: Graphical abstract and Figure 7 (the methods are explained, briefly, in the video lecture).

Recommended: 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, *only up until “Density of study in this 3D space,” p. 571.*

NOTE: REQUIRED CLASS AND FEEDBACK Feedback questionnaire, concerning course design, due April 26. Class attendance required April 27, for feedback session with former student Spencer Koerner.

26. Apr 27 What is science?

video: https://www.youtube.com/watch?time_continue=4&v=PwCyvSDkUCY&feature=emb_logo

wrap-up video: https://www.dropbox.com/s/xjnh11gf96mhysz/25_WrapUp.mp4?dl=0

Additional reading, for those interested: Kass (2021) The two cultures: Statistics and machine learning in science, comment to accompany the reprinting of an article by Leo Breiman, *Observational Studies*. (In Readings as KassOnBreiman.pdf)

SAQ4 Window: Apr 28-May 1.