

Computational Psychology

Time: Fridays 2:10pm – 3:50pm

Location: 306 Hamilton Hall

Instructor: Dr. Chujun Lin

Email: cl4767@columbia.edu

Office hours and location: By email appointment

This syllabus is subject to changes

COURSE DESCRIPTION

The development of computational tools, such as artificial intelligence, has transformed daily life and many areas of research. For instance, the development of large language models such as ChatGPT, have changed the way people write and think. These computational tools also hold great potential for advancing psychological research. However, they are still unfamiliar to many students and psychologists.

In this class, you will learn about how to apply computational tools to facilitate practical problem solving in the real world and research at different stages, including before, during, and after data collection. Specifically, through a combination of lectures (focusing on the mathematical basics and historical developments of these computational methods), in-class engagements (paper discussions), and coding workshops (applying practical coding skills to implement the computational methods), I will introduce you to methods for handling complex stimuli, validating data quality, and modeling big data. Workshops will be administered via Colab using Python, so please be sure to bring a laptop to every class.

COURSE OBJECTIVES

Over the course of this quarter, I hope you will learn:

- The history and development of a range of computational methods
- The mathematical basics of a range of computational methods
- How to apply these computational methods to solve practical real-world problems
- How to apply these computational methods to advance a range of research questions
- How to think critically about these computational methods
- Practical skills of coding in Python and implementing these advanced models

Don't worry if you are not familiar with Python - the aim of this course is not to test your skills but help you acquire new knowledge. Not every one of the methods I will talk about applies to your career goal or research to the same degree - some may be more relevant, and others may be less. For the ones that are more relevant to your career goal or research, you may want to look beyond the content taught in this class to keep up with the state of the art since computational tools evolve rapidly nowadays. For the ones that are seemingly less relevant to you currently, you are encouraged to an open mind since they may turn out to be relevant to you in the future.

GRADE CRITERIA

- ***Assignments - 40%***

There will be 4 assignments spread throughout the semester. These assignments are intended to help you practice applying these sophisticated computational methods to solve practical problems through coding in Python with Google Colab.

You are not required to have backgrounds in coding or Python, but you should at least reserve 10 hours to complete each assignment. If you do not have any prior experience with coding or Python, you may want to reserve at least 16 hours to complete each assignment. **DO NOT WAIT TILL THE LAST MINUTE TO WORK ON YOUR ASSIGNMENT**. You will have one week to complete each assignment.

You will receive credit for completing each of the assignments accurately and on time (10% each, 40% in total).

- ***Paper Discussion – 40%***

There will be 6 papers for discussion in three of the classes (2 papers per class). Each of you will form a group of 3-4 students to lead the discussion of one paper. Your group will lead the discussion by giving a 45-minute presentation of the paper, including its questions and conclusions, detailed methods, your critique (what are some limitations of the paper and how you think future research can address them), and 3-4 questions for the class to discuss together (one question generated by each of the group members).

In particular, you are responsible to understand every detail of the methods of the paper, including reading the supplementary materials; then you need to figure out how to explain the methods to your classmates in a comprehensible way in your presentation; you will also take any questions that your classmates may have.

For the paper you are presenting, please reserve 17 hours to read the paper inside out (you may need 12 hours) and work on your group presentation (you may need 5 hours). For the rest of the 5 papers you are not presenting, you are still responsible to read the paper before class and engage in the discussion – please reserve 3 hours to read each paper.

You will receive credit for completing each paper presentation, in particular, describing the methods in the paper accurately (20%) and presenting critique and discussion questions of high quality (20%).

- ***Attendance – 20%***

Attendance is mandatory for this class since class engagement is essential for you to understand the methods and apply them in your homework and discussion. There will be 12 classes in total this semester. You can have 2 absences for emergency without

impacting your grades. You will receive credit for attending the rest of the classes (2% each, cap at 20%).

- ***Letter Grade***

Finally, letter grades will be assigned according to the following percentage scale. Percentage grades will be rounded to the nearest point and the stated letter grade cutoffs will be applied without exception.

- A+ if your numeric score $\geq 97.00\%$ the top score
- A if your numeric score = 93.00-96.99% of the top score
- A- if your numeric score = 90.00-92.99% of the top score
- B+ if your numeric score = 87.00 - 89.99% of the top score
- B if your numeric score = 83.00 - 86.99% of the top score
- B- if your numeric score = 80.00 - 82.99% of the top score
- C+ if your numeric score = 77.00 - 79.99% of the top score
- C if your numeric score = 73.00 - 76.99% of the top score
- C- if your numeric score = 70.00 - 72.99% of the top score
- D if your numeric score = 60.00 - 69.99 % of the top score
- F if your numeric score $< 60\%$ of the top score

COURSE POLICIES

- ***Academic Integrity***

All students are expected to adhere to standards of academic integrity. Cheating of any kind on any assignment will not be tolerated, including copying the code from your classmates. It is disrespectful to your peers and the university. Consequences for academic misconduct may include a failing grade in the course and official action (e.g., academic probation or expulsion) by the University. For more information on Columbia's policies on academic integrity, please see here: <http://www.college.columbia.edu/academics/academicintegrity>.

- ***Class Materials***

The lecture slides and Colab Python coding guides will be made available before each class in which they are presented. These materials are intended to help you take notes during the class, they are no substitute for attendance and engagement in the class.

- ***Special Needs***

If you require special accommodations through the Office of Disability Services (ODS), please let me know as close to the beginning of the semester as possible. More information about registering with ODS can be found here:

<https://www.health.columbia.edu/services/ods>.

- ***Student Wellness***

Academic life can be exceptionally challenging at times. The intention of this course is to enhance your life, not make it miserable. Please always prioritize your health. If you're ever having a tough time, and it's impacting your ability to fully participate in the course, please reach out, and we can figure something out together. The university also has many health and wellness resources available for students:

- <https://www.health.columbia.edu/content/counseling-and-psychological-services>
- <https://blogs.cuit.columbia.edu/nightline/>
- <https://universitylife.columbia.edu/student-resources-directory#!#health>

- ***Other Requests***

If you have other requests or questions, please feel free to talk to the instructor before or after the class, through email, attend the office hours, or schedule a private meeting.

COURSE SCHEDULE

- ***(9/5) Analyzing Naturalistic Text I***

One of the most important types of information people encounter in everyday life is text. For instance, people read news reports, books, and scientific papers to gain new knowledge, message one another on social media to maintain connection, and write comments on the internet to express themselves. In this class, you will learn about the history of text analysis and the mathematical basics of natural language processing. We will discuss how the structure and meaning of language can be captured using statistical and computational methods, and how these methods have evolved from simple word counts to more sophisticated vector representations that underlie modern natural language processing models. We will also consider the types of questions about human behavior and cognition that text analysis can help us answer.

- ***(9/12) Analyzing Naturalistic Text II***

We will continue with analyzing naturalistic text. We will learn about the architecture of more advanced large language models that have been developed in recent years. We will then look at examples of their applications for solving practical psychological problems and advancing psychological research. Finally, we will look at the actual implementation of these models in Python via Google Colab. We will explore how these models represent meaning, generate language, and capture patterns in human communication, and we will begin hands-on work analyzing text data from psychological studies.

- ***(9/19) Paper Discussion I***

Paper 1: Charlesworth, T. E. S., Sanjeev, N., Hatzenbuehler, M. L., & Banaji, M. R. (2023). Identifying and predicting stereotype change in large language corpora: 72 groups, 115 years (1900–2015), and four text sources. *Journal of Personality and Social Psychology*, 125(5), 969–990. <https://doi.org/10.1037/pspa0000354>

Paper 2: Thornton, M. A., Rmus, M., Vyas, A. D., & Tamir, D. I. (2023). Transition dynamics shape mental state concepts. *Journal of Experimental Psychology: General*, 152(10), 2804–2829. <https://doi.org/10.1037/xge0001405>

- ***(9/26) Analyzing Naturalistic Visual Information I***

Another important type of information people encounter in everyday life is visual input. For instance, people try to understand one another's emotions and intentions by observing one another's facial expressions and body language. In this class, you will learn about the history of computer vision and the mathematical basics of computer vision models. We will discuss how visual information can be represented as data, how early vision systems extracted features from images, and how these ideas laid the groundwork for modern vision models. We will also examine how visual processing connects to psychological theories of perception and person understanding.

- ***(10/3) Analyzing Naturalistic Visual Information II***

We will continue with analyzing naturalistic visual information. We will learn about the architecture of more advanced deep convolutional neural networks that have been developed in recent years. We will then look at examples of their applications for solving practical psychological problems and advancing psychological research. Finally, we will look at the actual implementation of these models in Python via Google Colab. This will include working with pre-trained models to extract features from faces and bodies and exploring how these features can be used to study emotional expression, social judgments, and behavioral prediction.

- ***(10/10) Paper Discussion II***

Paper 1: Liu P, Bo K, Ding M, Fang R (2024) Emergence of Emotion Selectivity in Deep Neural Networks Trained to Recognize Visual Objects. *PLOS Computational Biology* 20(3): e1011943. <https://doi.org/10.1371/journal.pcbi.1011943>[Links to an external site.](#)

Paper 2: Khaligh-Razavi SM, Kriegeskorte N (2014) Deep Supervised, but Not Unsupervised, Models May Explain IT Cortical Representation. *PLOS Computational Biology* 10(11): e1003915. <https://doi.org/10.1371/journal.pcbi.1003915>[Links to an external site.](#)

- ***(10/17) Optimization Algorithms I***

After learning about the architecture and applications of models for analyzing naturalistic text and visual information, we now turn to how these models are actually fit to data (how they are trained). In this class, you will learn about optimization algorithms, with a focus on gradient descent—the workhorse behind modern machine learning. We will walk through the mathematical intuition behind gradient descent, explain how it updates model parameters to minimize error, and demonstrate its implementation in Python. This week serves as a foundation for understanding how large-scale models learn from data.

- ***(10/24) Optimization Algorithms II***

While gradient descent works well for smooth and differentiable loss functions, many problems in psychology and cognitive modeling involve objectives that are non-differentiable, discontinuous, or contain many local optima. In this class, you will learn about alternative optimization algorithms that are better suited for such cases. We will cover simulated annealing, genetic algorithms, and maximum variation strategies. These methods are particularly useful when the solution space is rugged or when the goal is to explore diverse configurations rather than converge on a single optimal solution. You will also learn how to implement these algorithms in Python and apply them to real psychological modeling problems.

- ***(10/31) Canceled! Enjoy your Halloween!***

- ***(11/7) Paper Discussion III***

Paper 1: Lin, C., Keles, U. & Adolphs, R. Four dimensions characterize attributions from faces using a representative set of English trait words. *Nat Commun* 12, 5168 (2021). <https://doi.org/10.1038/s41467-021-25500-y>

Paper 2: TBD

- ***(11/14) Analyzing Big Data – Handling High-dimensional Data I***

If you want to predict a dependent variable using more independent variables than the number of observations you have, you encounter the problem of overfitting. For instance, in this case, there will be no unique solution to a standard linear regression. In this class, you will learn about ridge regression, a form of regularized regression that addresses overfitting by penalizing large coefficients. We will discuss how ridge regression improves model stability and generalization by trading off bias and variance. You will

also learn how to optimize the regularization strength—a key hyperparameter—using cross-validation. By the end of the class, you will understand how to implement ridge regression in Python, apply it to high-dimensional psychological datasets, and interpret its results in the context of psychological theory.

Sometimes you may assess an outcome variable using multiple measures. For instance, to understand an individual's personality, you might collect self-reports on a list of personality traits. A common approach is to average these traits and use the composite score as the dependent variable in a standard regression. However, this approach can obscure the rich variability across individual traits and the nuanced structure within the data. In this class, you will learn about representational similarity analysis (RSA), a powerful method that allows you to retain high-dimensional outcome variables by comparing patterns of similarity rather than reducing them to a single score. We will explore how to construct and interpret representational dissimilarity matrices (RDMs), apply RSA to both predictors and outcomes, and test the significance of these associations using nonparametric permutation tests. You will also gain hands-on experience implementing RSA in Python, and see how this framework can be applied to behavioral, neural, and perceptual data to uncover deep structure in psychological representations.

- ***(11/21) Analyzing Big Data – Handling High-dimensional Data II***

Building on last week's focus on high-dimensional outcomes, this class introduces artificial neural networks as a flexible modeling approach for simultaneously handling high-dimensional dependent and independent variables. Unlike traditional regression methods, neural networks such as multi-layer perceptrons (MLPs) can model complex, nonlinear relationships across vast numbers of input and output features. In this class, you will learn about the architecture of feedforward neural networks, how information flows through layers of nodes, and how weights are learned through backpropagation. We will walk through how to design, train, and evaluate neural networks using real psychological datasets where the goal is to capture rich multivariate mappings—for example, predicting full behavioral profiles from image or text embeddings. You will also learn how to prevent overfitting in neural networks using techniques such as dropout and early stopping, and how these models can provide new insights into complex psychological processes.

- ***(11/28) Canceled! Enjoy your Thanksgiving!***

- ***(12/5) Analyzing Big Data - Analyzing Big Data – Handling Time-Series Data***

So far, you have learned about several methods that can handle high-dimensional psychological data. However, these methods generally assume that each observation is independent of the others. In many real-world scenarios, this assumption does not hold. Psychological and behavioral data often unfold over time—for example, physiological responses during a conversation, eye-tracking patterns across a screen, or sequences of words in a sentence. In such cases, the value of a given observation depends on those that

came before it. In this class, you will learn about how to model such temporal dependencies using artificial neural networks that incorporate memory, focusing on the Long Short-Term Memory (LSTM) architecture. LSTMs are a type of recurrent neural network (RNN) designed to retain and update information over long sequences, making them well-suited for time-series and sequential data. You will explore how LSTMs work, how they differ from standard feedforward models, and how to implement and train them in Python. We will also discuss practical applications in psychology, such as modeling emotion dynamics, behavioral trajectories, and language processing.