

ECE 594N: Special Topics in Pattern Recognition

Geometric and Dynamical Principles of Neural Computation

Instructor: Dr. Nina Miolane

Affiliation: Geometric Intelligence Lab - UC Santa Barbara (<https://gi.ece.ucsb.edu>)

Quarter: Winter 2025

Format: Inverted class format, two lectures per week

Prerequisites: Mathematical maturity with linear algebra, differential equations, and probability, and ability to code with Python

Class webpage + office hours: We will use the [Google drive](#) for sharing class related information. We will also use the [Slack](#) channel for questions and setting up meetings.

Course Information

Welcome! This is a 10-week-long introductory course for researchers from various backgrounds (EE, CS, physics, applied mathematics, neuroscience, and biology) who wish to gain hands-on experience in computational neuroscience. The main premise is to learn about the latest computational tools from geometry and dynamical systems that allow studying computations performed by artificial and biological neural networks. There are three modules:

- Module I provides a historical view of population codes, starting from static decoding approaches of earlier decades to the modeling of dynamics and geometry of the neural code.
- Module II motivates and introduces the theoretical basis of the latent computation framework (LCF), a theory of neural computation that explains several empirically observed phenomena.
- Module III discusses further applications to a broader set of open research directions in artificial and biological neural networks.

Assessment

- There will be one take-home midterm (30%), which will cover the basic topics taught in module I and potentially some fundamental topics from module II.
- The main assessment component of the course is the final project (60%), which can be either
 - (preferred) a research project on a novel topic broadly relevant to the machine learning community and involves some form of dynamics, geometry, or neural computation, or
 - (if desired) a review article of a particular subtopic in line with the topics of the class.
- Active participation in class discussions (10%).

Our mutual agreements as adults

- For all updates on the class, please join: [Slack](#) channel. Furthermore, class material will be disseminated via the [Google drive link](#).

- The class will follow an inverted lecture style. Instructor (Prof. Nina Miolane) will share lecture notes a week in advance (monday of the previous week) and students will read them before coming to class. The first hour of the class will be spent on discussing the content and answering the questions as a group. The second hour will be spent on working on the tutorial problems of the day, with any remaining time spent on discussing research projects.
- We expect everyone to have a final project assignment before the end of week 4 (January 30th, 5pm PT), preferably in collaborative groups of 3 or 4 students. You are welcome to discuss your project ideas with the instructors and involve them in the project, or completely leave them out and collaborate with your favorite group. The only rule is that the collaboration has to be newly formed and be in an exploratory spirit. We will also ask labs around the campus to participate and discuss potential project proposals as part of the course. There will be candidate project ideas posted. More on this later, see the website for regular updates.
- Class participation will be computed following the formula: $15 - \max(\# \text{ of days missed}, 5)$. Basically, you get 5 free days, no questions asked. Please use them wisely and do not ask for exceptions. What is being discussed in the classroom is not available in a textbook, so try to be actively present! You will self-report your participation grade, we will not keep track of your presence.
- The midterm exam will be assigned by week 2 and due until the end of week 5 (February 6th, 5pm PT). It will be open-ended, and allow you to practice asking and answering research questions. You are welcome to work on your own for the exam, but our strong preference is you work on the exam together with your research group. You can submit one exam per group if preferred, but then a contribution statement has to be present and all members should write they understand and agree with the solutions (no effects on the grade, just an honor code requirement).
- Assessments of take-home midterm will be done collaboratively (in week 6, date TBD). The midterm will be assigned weeks in advance and asynchronously, please plan accordingly and do not ask for extensions, as the asynchronous nature should allow accommodating any emergency lasting less than half the quarter. The date assigned for assessment is mandatory for everyone, think of it as the day of the midterm exam.
- The final project has four components.
 - Zeroth component should be a single page description (title, authors, and problem statement + tentative methods) of the proposal (0%, due January 30th 5pm PT).
 - In the first component, you are asked to provide a write-up of your results in NeurIPS format. This is due by the end of week 9 (March 6th, 5pm PT) and will be graded based on a rubric by the instructors (20%).
 - The second component is poster presentation, which will take place in the week of final exams in a designated spot. The poster PDFs are due the morning of the presentation. The grading will be done by instructors using a rubric (20%).
 - The final component is the reviewer feedback, due by the end of week 10 (March 13th, 5pm PT). Each student will provide feedback to at least two other writeups. The quality of the feedback will be graded by the recipients of the feedback and the instructors jointly,

also based on a rubric (20%).

- All rubrics, midterm or final, will be shared with you beforehand. The goal standard for the midterm exam is the demonstration of the collaborative critical thinking ability, whereas for the final project, it is a submission to the NeurIPS 2026 conference

AI usage policy

No AI assistance is allowed in the midterm exam or when providing reviews to your fellow classmate's final project, but it is allowed for any other component of the class. For instance, you can use AI to understand a paper written by your classmate or assist in writing your own paper, but you cannot use AI to suggest how a paper you are assigned to should be reviewed (e.g., summary should be your own words, and similarly for weaknesses or strengths etc.). Any work turned in as part of the class, apart from the reviews and the midterm where it is prohibited, should have a single comprehensive paragraph describing how AI was used.

Academic Honesty

Collaboration and intellectual exchange are integral parts of this course. Students are encouraged to study and discuss course materials together. However, all submitted work must accurately reflect the individual or collective efforts of those credited. When submitting any assignment, students must explicitly acknowledge all forms of assistance or collaboration received from classmates, instructors, or external sources. There is no restriction on the length of the acknowledgment section, and it will not influence grading in any manner. Failure to acknowledge assistance constitutes a violation of the Honor Code. Academic integrity requires full transparency in all submitted work.

Authorship criteria: Authorship on a final course project does not automatically confer authorship on any subsequent publication derived from that work. Authorship decisions for publications shall be made collectively by all contributors, in accordance with established academic standards for contribution and responsibility. Instructors will not intervene in such matters unless they are themselves coauthors, in which case the same contribution-based principles apply. If instructors or external advisors provide supervision or intellectual input on a project, students bear the responsibility of appropriately involving or notifying them in any subsequent dissemination or submission of the work. Otherwise, authorship disputes are outside the scope of this course.

Course Content

Module I: A Brief History of Population Codes (Weeks 1–4)

Week 1 (January 5–9)

Lecture 1: Introduction to the history of population codes. Hopfield networks. Coding theories of the early 2000s: tuning curves, and population vectors.

Lecture 2: Fisher's discriminant analysis, noise and signal correlations, extractable information limits from large-scale neural populations.

Week 2 (January 12–16)

Lecture 3: A brief introduction to dynamical systems theory, attractors, repellers, stability analysis. An intuitive intro to Hopfield networks.

Lecture 4: Recurrent neural network models of neural computation. Chaos in randomly initialized RNNs. Basic components of network training: task-architecture-learning algorithm.

Important events: Midterm released January 16th, 5pm PT.

Week 3 (January 19–23)

Monday lecture skipped: Jan 19th is a federal holiday.

Lecture 5: Dimensionality reduction and latent variable models in systems neuroscience, a survey of existing methodologies in modeling population dynamics.

Important events: Project ideas released January 23rd, 5pm PT.

Week 4 (January 26–30)

Lecture 6: Dynamics and RNNs: Reverse-engineering computations learned by recurrent neural networks. Low-rank RNNs, designing a bistable dynamical system into RNNs by hand (Hopfield all over again!).

Lecture 7: Geometry and RNNs. Extrinsic, intrinsic, and latent dimensionality. High-dimensional nature of neural responses, a theory of coding efficiency based on neural manifold dimensionality.

Important events: Project proposal due January 30th, 5pm PT.

Module II: Latent Computation Framework (Weeks 5–6)

Week 5 (February 2–6)

Lecture 8: A tale of empirical observations concerning population codes. What new insights will new imaging and intervention technologies bring with large-scale datasets?

Lecture 9: A general model of a biological neuron, and a broad class of biologically plausible neural networks. A new type of encoding-embedding relationship achieving redundancy of neural coding and universality of neural computations.

Important events: Midterm due February 6th, 5pm PT.

Week 6 (February 9–13)

Lecture 10: Revisiting the high-dimensionality of the neural responses from a low-dimensional latent code. Origins of high-dimensionality: Complex tuning curves in the age of latent variables, a tale of two decades unifying geometry and dynamics of population codes.

Lecture 11: Sufficiency of linear decoding from latent variables. Latent origins of information limiting low-rank correlation matrices. Estimating architecture-agnostic latent dimensionality.

Important events: Midterm assessment date TBD.

Module III: Advanced Topics (Weeks 7–10)

Week 7 (February 16–20)

Monday lecture skipped: February 16th monday is a federal holiday.

Lecture 12: An introduction to representational drift in systems neuroscience and its theoretical study with recurrent neural networks.

Week 8 (February 23–27)

Lecture 13: Why RNNs fail to learn long-term dependencies? An analytical theory of abrupt learning using latent computation framework.

Lecture 14: Dynamical phases of latent mechanisms subserving short-term memory in recurrent neural networks. Use of RNNs for digital twin experiments.

Week 9 (March 2–6)

Lecture 15: A brief history of grid and place cells, the Nobel Prize of Medicine in 2014. New approaches to understanding the latent code subserving the spatial navigation.

Lecture 16: Experimental details of how we collect neural datasets in live performing animals. Neural code subserving predictive pursuit (guest lecture by Dr. Andy Alexander).

Important events: Final paper report due March 6th, 5pm PT.

Week 10 (March 9–13)

Lecture 17: A theory of parameter identifiability in recurrent neural networks. Connections to digital twin experiments and reliable modeling practices.

Lecture 18: Measuring the degeneracy of solutions using state-space geometry. Changes in the geometrical properties of neural networks during training.

Important events: Reviews due March 13th, 5pm PT.

Final presentations (March 16–20)

Feedback rounds with the instructors on the final project. The final poster presentation date will be announced by the registrar, the same as the final exam date.

Mandatory Description of the AI Usage

Artificial intelligence tools were used in the preparation of this document to assist with language editing, refinement of paragraph flow, and organization of course topics. The substantive content, structure, and intellectual contributions of the course materials were created by the instructors and represent their original work.