Bogazici University
Department of Computer Engineering

CMPE58I - Human-Inspired Machine Intelligence

Fall 2025
Course Information

Instructor:	Dr. Şuayb Ş. Arslan	**Lecture Hours:**	MWW 578
Email:	suayb.arslan@bogazici.edu.tr	**Places:**	NH 202|BM B3 | BM B3

Motivation for/Description of the Course:

Human brain and neuroscience research offers valuable insights that can significantly enhance artificial intelligence by incorporating biological details that are crucial to human intelligence. As AI models advance, they not only deepen our understanding of the brain but also uncover the mechanisms that underpin intelligence and our complex mental experiences. However, both neuroscience and AI are vast and intricate fields, each comprising numerous subfields that demand specialized knowledge and skills. Mastering just one subfield often requires years of dedicated study.

This course bridges neuroscience and artificial intelligence by leveraging insights from the human brain to inspire advanced AI models. Focusing on key areas such as vision, auditory, sensorimotor systems or even neurons and biology, the course guides students through brain monitoring technologies such as EEG and fMRI, teaching them to process raw data and translate these insights into the design of computational systems. These systems are based on biological principles, emphasizing the development of deep learning techniques inspired by the brain's architecture, training methods, inference protocols, and dataset processing strategies. Additionally, the course explores some of the fundamentals such as symbol grounding and binding problems as well as emerging research areas that address human traits and abilities that are challenging to replicate, such as consciousness, ethical values, morality, and other complex emergent behaviors. By integrating these advanced topics, the course equips students with the knowledge and skills to contribute to AI development, creating systems that reflect the intricacies of human intelligence and address profound challenges at the intersection of cognition and machine intelligence.

Reference Books: Depending on the level of students, there are three references I frequently resort to.

- "The Principles of Neural Science" by Eric R. Kandel, James H. Schwartz, and Thomas M. Jessell.
- "Computational Neuroscience: A Comprehensive Approach" edited by Jianfeng Feng.
- "Neural Engineering: Computation, Representation, and Dynamics in Neurobiological Systems" by Chris Eliasmith and Charles H. Anderson.

Overall Educational Objectives: To learn the fundamentals of brain functionality, measuring it, translation techniques to digital computation, building architectures and training mechanisms, execution of computational methods on parallel hardware. The course will review and compare neural computation in brain against digital computation with several examples. To do that, students shall learn about processing some of the fundamental brain imaging data such as EEG and fMRI. We also intend to give assignments to use clusters with GPUs to train/test models and help them get hands-on experience. Towards the end of the course, we shall ask students to read recent papers on this topic and guide them through future trends. In addition, some of the philosophical questions and the impact of bio-inspired machine intelligence to various fields of research will be discussed. Major learning outcomes of this course are:

- learn the modular structure and fundamental principles of brain functionality;

- apply signal processing to brain monitoring techniques to derive inspiration;
- translate observation to design computational frameworks;
- develop bio-inspired training, learning rules (optimization), architecture and inference techniques;
- leverage parallel hardware to implement and execute computational models to justify bio-inspired methodology.

Prerequisites: Basic Probability and Linear Algebra. Basic signal processing, Python frameworks, Basic understanding of Machine/Deep Learning.

Tentative Course Outline:

- **Week 1:** Introduction to Neuroscience: Human Vision, Auditory system, etc..
- **Week 2:** Computational Neuroscience: Neurons, Receptive Fields, Modeling, etc.
- **Week 3:** Limitations of classical AI: Bias, Commonsense, etc.
- **Week 4:** Information Extraction (Neural Decoding): Eye Tracking, Online Exp. etc.
- **Week 5:** Brain Imaging Techniques I: EEG, MEG Basics
- **Week 6:** Brain Imaging Techniques II: fMRI and other techniques.
- **Week 7:** Biomedical Signal Processing: Pre-processing, Post-processing (EEG)
- **Week 8:** Project Ideas and Applications (Discussions) [**Midterm Week**]
- **Week 9:** Information Transferability: Deep and Layered Learning Paradigm
- **Week 10:** Human-inspired Deep Learning I: Architecture and Memory
- **Week 11:** Human-inspired Deep Learning II: Training/Learning
- **Week 12:** Human-inspired Deep Learning III: Bio-plausibility, Development/Noise
- **Week 13:** Comparing Humans and Machines: Evaluation
- **Week 14:** Problems in Human-centric AI: Symbol Grounding, Counterfactual thinking, Meta–Learning, etc.

Grading Policy: 10 quizzes(2 pts each) (20%), 3 software assignments(7 pts each)(21%), Midterm (24%), Final Project (35%).

Important Dates:

Quizes/classworks	Weekly (except first two and last two)
Midterm	End of 8th week.
Final Project Presentations	End of Semester.

Quizes: We will have ten quizzes/classworks (2 pts each) to increase student participation (quizes) and group work (classworks).

Software Assignments and Datasets: We shall assign one preprocessing pipeline implementation for EEG data (EEGlab or MNE) and two python implementation assignments on *Google Colab* or another favorite coding platform (both `tensorflow` and `pytorch` are ok) for model training. Students will run their models on benchmark datasets (ImageNet, COCO, CIFAR10, CIFAR100, etc.) or their real-world noisy versions (CIFAR-10N, CIFAR-100N, etc.) and report what they find with a thorough discussion. Upon the availability of resources, ImageNET and EcoSet can also be utilized.

Potential Project Ideas: Here are a few research directions (ideas) for the project teams to pick and investigate. These topics shall be discussed during the lecture hours. Other topics should be fine as long as you check their appropriateness with the instructor's approval.

1. Receptive Fields (RFs) in CNNs: How to calculate them? How to use (biologically inspired) Techniques to improve RF size? What is the upshot?

2. Most of our input is circular or elliptic (not rectangle or square). Our computation platforms are better at processing matrices or vectors i.e., geometry with edges. What would be the effect of translation of circular visual data ingest to matrices for efficient computation? Performance? RFs? Etc.
3. How to use human eye gaze info for designing architecture, training and optimization techniques? Attention guidance in the training phase? Any cost function updates?
4. How to use EEG findings: (1) Spatial Frequency processing, (2) ERP results for face and body processing, (3) Integration of EEG signatures in the training, etc.
5. How to leverage internal and/or external noise to improve generalization performance and model robustness? How does that change the saliency maps? Exploration of computational equivalent of stochastic resonance.
6. How to use human decision-making process in the development of computational frameworks to create a hybrid system? Utilization of online experiments (e.g. multiple choice odd-one out task) and adaptation of human confusion matrices to deep learning systems.
7. The emulation of myelination process in the current SoTA computational techniques. Emergent behaviors and learning based on differential information transfer rates in feedforward networks.
8. Adaptation of biological recurrence to computational models. Critique of current SoTa recurrence implementation in digital hardware (RNNs, LSTMs, etc.). What could be an EEG/MEG-inspired recurrent modeling using topographic maps?.
9. While completing a specific task, how to close the gap between the representations generated by the brain and by the computational systems? What advantages does this approach provide (robustness, fairness, etc.)?
10. Biomimetic training strategies: From gray scale to color, from blur to high spatial resolution, from less noise to more noise, from high to low latency transmission, etc. Effect of training on the choice of architecture based on biological plausibility axis (using novel metrics to measure the plausibility).
11. Do bio-inspired architecture, learning (optimization) or training ensure more ethical, responsible and aligned computational models?

Course and Class Policy Guidelines:

- Attendance is not mandatory, but I strongly recommend active participation as this will help you get good grade, particularly if your preliminary grade falls near a borderline.
- There will not be a make-up exam. In case of health or other personal emergencies as valid reasons for qualification, related questions of the final will be counted as the missing midterm grade.
- We will use *Moodle* to distribute class contents, and as a communication medium between you and the course admin team. If you have any problems accessing the course material on Moodle, please let us know ASAP so we can have such problems fixed on time.
- We will leverage some of the available GPUs in the cloud as a tool for training our models.

Academic Honesty: Lack of knowledge of the academic honesty policy is not a reasonable explanation for a violation. Violators will fail the course and corresponding legal options within the university plagiarism policy will be sought. Students shall avoid all forms of academic dishonesty, including but not limited to:

- *Plagiarism:* The use of another person's words without attribution and without enclosing the words in quotation marks. Plagiarism may also be defined as the act of taking the ideas or expression of ideas of another person and representing them as one's own, even if the original paper has been paraphrased or otherwise modified. A close or extended paraphrase may also be considered plagiarism even if the source is named.
- *Collusion:* When specifically prohibited in advance by the instructor, collaborating with another person in the preparation of notes, themes, reports or other written work offered for credit.
- *Cheating on an examination or quiz:* Giving or receiving information or using prepared material on an examination or quiz.
- *Falsification of data:* Manufacturing data, falsification of information, including providing false or misleading information, or selective use of data to support a particular conclusion or to avoid conducting actual research.

Use of Large Language Models: The use of LLMs is discouraged. Students may occasionally resort to LLMs for text/code completion tasks at the risk of unintelligible and poor–quality source code generation. However, upon detection of such use cases through human examination or automated tools, students shall be panelized strictly based on the similarity/perplexity scores generated using examination results. Note that no LLMs shall be allowed in the Midterms.