

## **BIODS 276: Advanced Topics in Computer Vision and Biomedicine**

Fall 2024

Instructor: Serena Yeung (syYeung@stanford.edu)

Location and Time: Mondays and Wednesdays, 10:30-11:50. Units: 3

**Course Description:** This course in artificial intelligence will provide a deep dive into advanced computer vision techniques for reasoning about visual data, and their real-world use in biomedicine. We will cover current cutting-edge models including different families of vision foundation models from representation learners to diffusion and generative models, and both vision-only and vision-language models. We will also cover considerations for real-world use, including model size and computation, training and inference settings, and training data, focusing on applications in biomedicine. Students will be actively engaged in studying and analyzing recent advancements through written analyses, class discussions, and a final project. This course is considered an advanced course and students should be comfortable with deep learning and computer vision at the level of CS231N or BIODS220.

**Prerequisites:** Proficiency in deep learning and computer vision, e.g. CS231N, BIODS220 or equivalent. Prior experience in natural language processing (e.g. CS224N) and biomedical AI (e.g. BIODS260, BIODS220) are helpful but not required.

### **Grading:**

Assignment 1: 20%

Assignment 2: 20%

Discussion Presentation: 15%

Course project: 40%

Class participation: 5%

Students will have 3 total late days in the course, with a maximum of 2 for any assignment. These can be used on A1, A2, the project proposal, and the project milestone report. Grades will be deducted by 25% for each additional late day. Late days *may not* be used on the project final presentation or final project report.

In this class, all deadlines refer to 11:59pm PST on the stated day.

### **Schedule (tentative):**

#### **Lecture Topic**

- 1 (9/23) Course Introduction
- 2 (9/25) Deep Learning for Image Data: From CNNs to Foundation Models
- 3 (9/30) Vision Representation Learners
- 4 (10/2) Vision Representation Learners in Biomedicine
- 5 (10/7) Vision-Language Representation Learners
- 6 (10/9) Vision-Language Representation Learners in Biomedicine

- 7 (10/14) Vision Diffusion and Generative Models
- 8 (10/16) Vision Generative Models in Biomedicine
- 9 (10/21) Vision-Language Generative Models
- 10 (10/23) Vision-Language Generative Models in Biomedicine
- 11 (10/28) A Deeper Dive into GPUs and Compute
- 12 (10/30) Guest Lecture: Advanced Vision Models in Biomedicine
- 13 (11/4) Discussion: Recent Advances in Vision Representation Learners
- 14 (11/6) Discussion: Recent Advances Vision Generative Models
- 15 (11/11) Discussion: Recent Advances in Vision-Language Representation Learners
- 16 (11/13) Discussion: Recent Advances in Vision-Language Generative Models
- 17 (11/18) Discussion: Recent Advances in Computing Efficiency
- 18 (11/20) Guest Lecture: Advanced Vision Models in Industry
- 19 (12/2) Guest Lecture: Responsible and Ethical Deployment of Advanced Vision Models
- 20 (12/6) Course Conclusion: The Next Frontiers of Computer Vision and Biomedicine

**Assignments:** There will be two homework assignments in the class, each worth 20% of the final grade. For each assignment, students will write a report in response to a prompt chosen from a course-provided list, e.g., read and compare / contrast two related papers. The topics will be related to and reinforce the material covered in class.

**Assignment Collaboration Policy:** All assignments must be completed individually. Study groups are allowed, but each student must produce an independent assignment and write the names of their group members on the assignment.

**Discussion Presentation:** There will be a number of class sessions dedicated to discussion of recent advances related to the course topics. These will include presentation of recent works in the field. Students will sign up to make presentations of recent works (from a course-provided list) during these class sessions, as well as to ask questions in other sessions.

**Course Project:** The course project will be worth 40% of the final grade. Students will perform an analysis or investigation of advanced vision models in biomedicine, in an area of their choice. Further details will be provided in class, and students can work in groups of 1-2. The goal is to gain hands-on experience interacting with some of the large vision and vision-language models discussed in class. The project will include a project proposal, milestone report, and final presentation and report.

**Class Participation:** In-class attendance is an important part of the learning experience for this advanced, discussion-centric course. Students are expected to attend at least 15 class sessions during the course (out of 20 in the schedule). Attendance will be recorded through a sign-in sheet at each class, and satisfactory attendance will comprise 5% of the course grade. Any additional absence beyond this will need to be under extenuating circumstances (such as a medical absence justified by a doctor's note), and depending on the circumstances students may be asked to write a one-page make-up report on a provided prompt.