Teaching Statement

Utkarsh Mall

Artificial Intelligence and Computer Vision have transitioned from being purely foundational research areas to significantly impacting numerous disciplines. Therefore through my teaching, I aim to instill the importance of solving real-world problems through these tools. I believe that the goal of an AI course therefore is to develop a fundamental understanding of concepts as well as learn to apply them and solve meaningful problems in scientific domains such as climate science, environment science, anthropology, and more.

Current and Past Teaching

Mentorship: During my doctoral studies and postdoc, I have mentored 15+ undergraduate and master's students. Several of these students have published and submitted papers in top vision conferences. While mentoring CS undergraduate researchers, I observed that students with a background in introductory ML and vision are often well-versed in modern ML techniques. However, they tend to struggle with tasks that occur before and after applying these methods. For example, while they excel at improving models on established benchmarks, they face challenges in tasks such as collecting and processing new datasets. Similarly, they are less confident in visualizing results in ways that provide insights or reveal weaknesses in a method, especially beyond standard metrics. Teaching these skills in applied courses would close the research loop and result in better-equipped and independent researchers, fostering a more comprehensive approach to tackling real-world problems.

Past teaching: I have given guest lectures and tutorials for several vision courses during my postdoc and doctoral studies. While teaching, I rely on active learning to keep students engaged with the material. I also structure my lectures in a top-down way, starting with a vision problem and showing the solution's effectiveness via applications. During the lecture, I incrementally broke down the solution by interactively asking students about the components needed and how to build those components. I believe that such a goal-oriented approach improves students' research skills. In the top-down structure, students are more likely to come up with good alternative solutions and future ideas on their own. The top-down approach also results in better exploration of interdisciplinary ideas.

During my doctoral studies at Cornell, I was a teaching assistant for graduate computer vision (with Prof. Noah Snavely), where my contributions included revising homework assignments and creating new grading tools and problems. During office hours, I applied a similar philosophy of breaking down assignments top-down and guiding students in building individual components. The feedback indicated that the students liked my teaching style. Indeed, this earned me an Outstanding TA award for the vision course.

Interdisciplinary teaching: I have also mentored and collaborated with students in other fields of interdisciplinary research such as crop science, climate science, fashion anthropology, archaeology, mass communications, and urban planning/engineering. Such collaborations have led to journal publications in their area of research [1, 2]. I also gave a guest lecture at NYU AI Summer School geared toward undergraduate students outside of CS majors. These experiences have deepened my understanding of the gaps in machine learning and computer vision knowledge that exist among researchers outside computer science.

During my time at Cornell, I also TA'ed for an undergraduate course on introduction to visual imaging under Prof. Don Greenberg. This course was geared towards interdisciplinary

students bridging Art, Architecture, Computer Science, and Engineering. I was responsible for tutoring fundamental vision and graphics concepts such as color theory, and rendering. The final project in this course required students to create an interactive 3D virtual environment using the tools learned throughout the course. A significant takeaway for me from this course was the importance of encouraging students to form interdisciplinary teams (CS and Arts majors in this case). This resulted in the most aesthetically pleasing as well as functioning 3D virtual environments. I aim to create a similar class environment bridging the gap between students with different perspectives and fields.

Teaching Plan

I look forward to teaching and developing both undergraduate and graduate fundamental courses centered around applied computer vision and machine learning. I am also interested in teaching courses with a more application and goal-driven approach. I plan to encourage this by designing projects/assignments to focus on dataset curation aspects and fully end-to-end systems. I am also interested in developing courses focusing on the application of AI/Vision in interdisciplinary areas that enable students of diverse academic backgrounds and interests to develop intelligent tools for their applications.

References

[1] Chelsea Butkowski, Lee Humphreys, and **Utkarsh Mall**. Computing colorism: skin tone in online retail imagery. *Visual Communication*, 2022.

[2] Rachel Rose Getman, Denise Nicole Green, Kavita Bala, **Utkarsh Mall**, Nehal Rawat, Sonia Appasamy, and Bharath Hariharan. Machine learning (ml) for tracking fashion trends: Documenting the frequency of the baseball cap on social media and the runway. *Clothing and Textiles Research Journal*, 2020.