UMBC
CMSC 491/691 Robust Machine Learning

Syllabus

Fall 2024, 3 credits

Contents

1 Course Information

1.1 Meetings and Instructors

Lectures	Monday, Wednesday 4:00 PM – 5:15 PM
Location	Math & Psychology Building 106
Course Website:	https://courses.cs.umbc.edu/graduate/691rml/

	Name	Email @umbc.edu	Office Hours
Prof.	Tejas Gokhale	gokhale	ITE 214 Wed 1430 – 1530
TA	TBD		

1.2 Communication

I expect all students to participate in classroom discussions, both by asking questions and by expressing opinions. I also welcome your feedback throughout the semester about how the course is going. In addition to regular office hours, you can schedule an appointment with me if you think one-on-one time would help.

Email. Please send questions related to the class via email to **both** the TA and the instructor using your UMBC email only. Please use the following prefixes in the subject line

- for questions about homework (in this example, homework 2): `[HW2] <your subject>`
- for questions about scribing: `[Scribing 09/18] <your subject>`
- for questions about projects: `[Group 4] <your subject>`
- for questions about lectures: `[Class on 09/20] <your subject>`
- for questions about exams: `[Exam] <your subject>`
- any additional questions: `[CMSC 491/691 Robust Machine Learning] <your subject>`

You may not receive a response for last minute messages or if you fail to follow these instructions; please plan ahead. Questions must follow academic integrity guidelines (Sec 3).

1.3 Prerequisites

CMSC 478 Machine Learning or CMSC 678 Machine Learning or IS 767 Deep Learning or an equivalent class taken at other universities is a prerequisite for the class. We will assume that you have a solid foundation in linear algebra, geometry, probability, and Python programming. Recommended classes at UMBC are: MATH 221 (Linear Algebra), STAT 355 or CMPE 320 (Probability and Statistics), MATH 151 (Calculus and Analytical Geometry). If you are unfamiliar with these topics, you should consider taking both: without these tools, you are likely to struggle with the course. Although we will provide brief math refreshers of these necessary topics, this class should not be your first introduction to these topics.

1.4 Topics

Models that learn from data are widely and rapidly being deployed today for real-world use, but they suffer from unforeseen failures – this course will explore the reasons for these failures and state-of-the-art mitigation techniques.

Part 1: Introduction and ML Refresher. (2 lectures) The course will start with a brief overview of the fundamentals of the standard machine learning pipeline, neural network architectures, and training algorithms.

Part 2: Robustness Challenges and Solutions. (12 weeks) Every week, the first meeting will be a lecture where the instructor will introduce and motivate the problem statement for one robustness challenge (such as distribution shift, adversarial attacks, calibration/uncertainty quantification, continual learning, unsupervised/self-supervised learning, few-shot/test-time learning, explanation/interpretability, dataset biases, etc.).

Every week, the second meeting will have:

- a quiz (worth 2.5%) based on the material of the first meeting
- presentations by three students explaining one paper each that proposes solutions to the specific robustness challenge
- discussions (about the presentation and/or topic)

We will study 12 robustness topics (list subject to change):

1. Domain Adaptation
2. Domain Generalization
3. OOD Detection
4. Adversarial Attacks, Backdoor Attacks
5. Uncertainty and Calibration
6. Online/Continual Learning
7. Self-Supervised/Unsupervised Learning
8. Test-Time Learning, Adaptation
9. Machine Unlearning, Model Editing
10. Interpretability, Explanability, Compositionality
11. Logic, Semantics, and "Commonsense"
12. Robustness Tradeoffs

Part 3: Invited Talks. (1-2 lectures) External speakers will deliver invited talks as part of the class on recent advances and findings in robust machine learning. These talks will be delivered by researchers at the forefront of reliability, robustness, trust, and safety of ML systems.

Please refer to the class website for additional details and updates:
https://courses.cs.umbc.edu/graduate/691rml/

1.5 Course Objectives

This course will introduce students to the robustness and reliability challenges faced by data-driven machine learning systems and the larger implications on the widespread adoption of these systems, their trustworthiness, and issues related to safe use of machine learning models. By taking this class, students will be able to:

- Identify and describe the robustness challenges in machine learning,
- Construct a formalism and definition for robustness in different application domains such as computer vision and natural language processing,
- Evaluate and score robustness of models with Python-based tools and benchmarks
- Apply, experiment with, test and compare several state-of-the-art techniques that degrade or improve an ML model's robustness,
- Examine and assess trade-offs between different robustness dimensions such as distributional robustness, adversarial robustness, fairness and privacy
- Analyze the effect of language guidance on robustness, generalization, and data efficiency of computer vision models
- Propose and design algorithms for domain-specific failure identification and mitigation, rigorously test the technique, and interpret the results to draw conclusions.

1.6 Reading Material

The class does not have a mandatory textbook and I don't expect you to buy one. I am writing a book on Advances in Robust Learning for Computer Vision – relevant chapter drafts may be provided for personal use. Other readings will come from papers from conferences such as CVPR, ICCV, ECCV, NeurIPS, ICLR, ICML etc. – please check the course website https://courses.cs.umbc.edu/graduate/691rml/.

The following books are excellent references for the course content and are available for free online at links provided by the authors.

- Pattern Recognition and Machine Learning, Christopher M. Bishop, Springer, https://www.microsoft.com/en-us/research/people/cmbishop/prml-book/
- Deep Learning, Ian Goodfellow and Yoshua Bengio and Aaron Courville, MIT Press, https://www.deeplearningbook.org/
- Applied Machine Learning, David Forsyth, Springer, https://link.springer.com/book/10.1007/978-3-030-18114-7

2 Evaluation

2.1 Evaluation Components

Paper Presentations	two 15 minute presentations on research papers (in class)	20%
Survey Papers	Review of relevant literature on two topics in robust machine learning (CVPR format; 4 page minimum; CVPR format)	20%
Weekly Quizzes	12 quizzes total. best 10 out of 12 grades.	25%
Project	Course project in groups of 3 (smaller groups only for PhD students with the professor's approval). • 5% for project proposal • 5% for midterm update video • 10% for final presentation • 15% for final report	35%

Extra Credit. Machine learning is a research area that celebrates creativity. We will provide opportunities for extra credit via open-ended questions or tasks that require creativity. Extra credit is capped at 10%.

Grading Scale. Please see the table below for the grading scale. Not that these are minimum grades – your final grade could be higher than this.

If you get at least	your minimum grade will be
90	A
80	B
70	C
65	D

491 vs 691? CMSC 491 is the undergraduate version while CMSC 691 is a graduate level version of the class. There is no difference in lectures and class materials. The main difference lies in the scope of the project. Projects will be judged on the basis of relative growth (from where you start to where you end).

Graduate students should work on a project with an original and unique research hypothesis with a potential for publication.

Undergraduate students can also propose original and unique research hypothesis, but will be allowed to work on:

- an idea provided by the instructors (i.e. you get to skip "ideation"), or
- or innovative applications or combination of existing work.

2.2 Deadlines and Late Submission

Due time for each assignment is 2359 (UMBC time), unless stated otherwise. Submission instructions will be provided with each assigned item. The deadlines and late submission policies for each evaluation component are stated below:

- **Paper Presentations:**
  - In class (sign up sheet will be shared)
  - If you are scheduled to present, but can't make it, send an email to the instructor and the TA to request date-change.
  - *Permissible reasons for date-change:* unforeseen events (illness, injury, emergency), travel to academic conferences, interviews
  - The instructor reserve the right to approve or deny date-change requests
- **Survey Papers:**
  - Due "next Wednesday 2359 UMBC time". *If* Lecture on "Domain Adaptation" is on Monday, Sept 09, *then* survey paper on "Domain Adaptation" is due on Wednesday, Sept 18, 2359 UMBC time.
  - Late submissions: 10% deducted for each late day.
- **Quizzes:**
  - In class. 12 quizzes. Your highest 10 grades will be chosen.
  - If you miss a quiz, you miss a quiz. No late submissions allowed. No make-up quizzes administered.
- **Project:**
  - Each milestone has a fixed deadline that will be announced in class.
  - Late Submissions: 10% deducted for each late day FOR ALL GROUP MEMBERS.

The instructor reserves the right to issue class-wide extensions.

3 Academic Integrity

I take academic honesty seriously. Do not cheat, deceive, plagiarize, improperly share, access or use code, or otherwise engage in academically dishonest behaviors.

- Doing so may result in lost credit, course failure, suspension, or dismissal from UMBC.
- Instances of suspected dishonesty will be handled through the proper UMBC administrative procedures.
- Every member of a group is responsible for the group's submission. If one member is academically dishonest, all members may be sanctioned, regardless of individual actions.

Please watch this brief video about integrity, courtesy of Dr. Cynthia Matuszek. This course follows the academic honesty policy from the Office of Undergraduate Education, available at: `https://academicconduct.umbc.edu/`. The following is a concise summary of the policies adopted:

By enrolling in this course, each student assumes the responsibilities of an active participant in UMBC's scholarly community in which everyone's academic work and behavior are held to the highest standards of honesty. Cheating, fabrication, plagiarism, and helping others to commit these acts are all forms of academic dishonesty, and they are wrong. Academic misconduct could result in disciplinary action that may include, but is not limited to, suspension or dismissal.

If you have any questions about what is or is not acceptable, ask first.

We are here to help. If you are struggling with the class content in any way it is your responsibility to communicate it to the instructor and TA so that we can assist you.

3.1 Types of Academic Dishonesty

- Plagiarism: Using a source (for code, text, images, etc.) without appropriate citations and recognition.
- Fabrication: Fabricating sources or any other information in your assignments.
- Aiding and abetting: Providing another student with answers, or helping them to cheat, is an equally serious violation of the principles of academic honesty. A student who commits such an offense is subject to the same penalties.
- Copying: Using another student's work (including students from previous offerings of this or similar classes at UMBC or elsewhere) for an assignment, exam, or project without acknowledgment.

This is not a complete list. To read the full Student Academic Conduct Policy, consult the UMBC Student Handbook, the Faculty Handbook, and the UMBC Policies section of the UMBC Directory (or for graduate courses, the Graduate School website).

3.2 Getting Help and Using Sources

Especially for computer science classes, there are generally questions about what is and is not allowed. You are encouraged to discuss the subject matter with others. The discussion board provides a great forum for this. However, you may not write or complete assignments for another student; allow another student to write or complete your assignments; pair program; copy someone else's work; or allow your work to be copied. (This list is not inclusive.) You are free to use online references like Stack Overflow only for issues that are not the primary aspect of the course and are not directly related to the assignments. For example, you may consult online forums for understanding how to use `numpy` `opencv` etc. or if you're getting a weird compilation error. Don't get stuck fighting your tools. Be sure to properly acknowledge whatever external help—be it from students, third party libraries, or other readings.

3.3 "Study" Sites

There are a number of sites that are primarily designed to help people get through classes without learning the material or doing the supporting work. Despite the self-branding as

"study sites," these are cheating sites, and using them will reduce your ability to learn the course material. In addition, the material on these sites is typically stolen, that is, used without permission of the authors. Uploading any course materials to any external site is is a violation of this class's academic integrity policy, because (1) it risks aiding and abetting, and (2) it is a copyright violation. It goes without saying that getting answers to homeworks, quizzes, etc. from such sites is plagiarism, and is academically dishonest. These violations will be handled like any other.

3.4 Use of "AI" Assistants

The use of "AI" assistants (including but not limited to language models such as ChatGPT) for completing *any* part of assignments in this class is considered cheating. There are no exceptions to this rule.

3.5 Group Work

Some work may be group work, which will be submitted by a group of two or more students. When submitting such an assignment, the same rules apply, except that the submitted work must be the work of the students as a group. By submitting a group assignment, each student is representing that the assignment is the work of the entire group, and each student takes full responsibility for the assignment's originality and content. This means that every member of a group is responsible for the group's submission. If one member is academically dishonest, all members may be sanctioned, regardless of individual actions. There may be additional penalties for failing to contribute to the group as expected or involving your group members in academic misconduct, which may go beyond a zero on the assignment. Note that this means that if a student in a group makes no contribution to the assignment, the rest of the group must not include their name. If someone does not contribute to a project, claiming that they did is aiding and abetting.

3.6 Good Practices

If the integrity of your work in this course is challenged, you are responsible for demonstrating proof that the work submitted is your own. A good starting point is to enable versioning/tracking in Google Docs, Word, Pages, or other software so that your writing activities/progress during the semester can be logged if necessary. Keeping copies of research notes, scribbles, and related material may be helpful, too.

3.7 Viva or Oral Defense of Flagged Submissions

To ensure academic and professional integrity, I reserve the right to hold a one-on-one oral viva (defense) of submissions deemed questionable, to determine your knowledge and mastery of the topic/resources versus the material submitted. Failing that viva will result in an 'F' on the assignment and an Academic Integrity violation report filed with the Graduate School.

3.8 Penalties

Academic misconduct could result in disciplinary action that may include, but is not limited to, suspension or dismissal. The **absolute minimum penalty** for a first offense of academic dishonesty in this course is a grade of zero on the assignment and a one-letter-grade reduction in the final class grade. However, depending on the nature of the offense, the penalty may be more severe, including but not limited to an `F` for the course, suspension, or expulsion. The minimum penalty for a second offense of academic dishonesty is an `F` for the course without possibility of dropping, but may be more severe.

4 Other Useful Resources

4.1 Accessibility and Disability Accommodations, Guidance and Resources

Accommodations for students with disabilities are provided for all students with a qualified disability under the Americans with Disabilities Act (ADA & ADAAA) and Section 504 of the Rehabilitation Act who request and are eligible for accommodations. The Office of Student Disability Services (SDS) is the UMBC department designated to coordinate accommodations that creates equal access for students when barriers to participation exist in University courses, programs, or activities. If you have a documented disability and need to request academic accommodations in your courses, please refer to the SDS website at `sds.umbc.edu` for registration information and office procedures.

- SDS email: `disAbility@umbc.edu`
- SDS phone: `410-455-2459`

If you will be using SDS approved accommodations in this class, please contact the instructor to discuss implementation of the accommodations. During remote instruction requirements due to COVID, communication and flexibility will be essential for success.

4.2 Sexual Assault, Harassment, Violence and Discrimination

UMBC Policy in addition to federal and state law (to include Title IX) prohibits discrimination and harassment on the basis of sex, sexual orientation, and gender identity in University programs and activities. Any student who is impacted by sexual harassment, sexual assault, domestic violence, dating violence, stalking, sexual exploitation, gender discrimination, pregnancy discrimination, gender-based harassment, or related retaliation should contact the University's Title IX Coordinator to make a report and/or access support and resources. The Title IX Coordinator can be reached at `titleixcoordinator@umbc.edu` or `410-455-1717`.

You can access support and resources even if you do not want to take any further action. You will not be forced to file a formal complaint or police report. Please be aware that the University may take action on its own if essential to protect the safety of the community. If you are interested in making a report, please use the Online Reporting/Referral Form. Please note that, if you report anonymously, the University's ability to respond will be limited.

Faculty and Teaching Assistants are Responsible Employees with Mandatory Reporting Obligations. All faculty members and teaching assistants are considered Responsible Employees, per UMBC's Policy on Sexual Misconduct, Sexual Harassment, and Gender Discrimination. Faculty and teaching assistants therefore required to report all known information regarding alleged conduct that may be a violation of the Policy to the Title IX Coordinator, even if a student discloses an experience that occurred before attending UMBC and/or an incident that only involves people not affiliated with UMBC. Reports are required regardless of the amount of detail provided and even in instances where support has already been offered or received. Faculty are required to report past and present sexual harassment, sexual assault, domestic and dating violence, stalking, and gender discrimination that is shared with them to the Title IX Coordinator so that the University can inform students of their rights, resources, and support. While you are encouraged to do so, you are not obligated to respond to outreach conducted as a result of a report to the Title IX Coordinator. If you need to speak with someone in confidence, who does not have an obligation to report to the Title IX Coordinator, UMBC has a number of Confidential Resources available to support you:

- Retriever Integrated Health (Main Campus): 410-455-2472; Monday – Friday 8:30 a.m. – 5 p.m.; For After-Hours Support, Call 988.
- Center for Counseling and Well-Being (Shady Grove Campus): 301-738-6273; Monday-Thursday 10:00a.m. – 7:00 p.m. and Friday 10:00 a.m. – 2:00 p.m. (virtual) Online Appointment Request Form
- Pastoral Counseling via The Gathering Space for Spiritual Well-Being: 410-455-6795; i3b@umbc.edu; Monday – Friday 8:00 a.m. – 10:00 p.m.
- Women's Center (open to students of all genders): 410-455-2714; womenscenter@umbc.edu; Monday – Thursday 9:30 a.m. – 5:00 p.m. and Friday 10:00 a.m. – 4 p.m.
- Shady Grove Student Resources, Maryland Resources, National Resources.

Child Abuse and Neglect: Please note that Maryland law and UMBC policy require that faculty report all disclosures or suspicions of child abuse or neglect to the Department of Social Services and/or the police even if the person who experienced the abuse or neglect is now over 18.

4.3 Inclusion, Safety, and Important University Resources

All students are entitled to a safe, respectful, and inclusive learning environment both inside and outside the classroom. This includes freedom from harassment, violence, prejudice, and exclusionary behavior toward any group. It also includes a welcoming atmosphere and appropriate accommodations for all situations. Disagreements via respectful responses are welcome and valuable to the discussion; unkind or dismissive comments and personal attacks are not. College can be stressful. Don't add to that stress by engaging in harassing or hostile behaviors. They are not welcome in the classroom and are completely inappropriate. Please see the Office of Equity and Inclusion's website for the most up-to-date information and policies.