Syllabus

Table of contents

🔆 Overview

1	**Semester**	Spring 2025
2	**Course Number**	CSE 524, CEN 524, CSE 494
3	**Course Name**	Machine Learning Acceleration
4	**Modality**	In-person (Face to Face)
5	**Lecture room number**	Tempe - ECGG 224
6	**Class timing**	1:30-2:45pm, T/Th

📌 If you have taken the following courses already, then please do not enroll for this course. The material of these courses overlap significantly. So, you will not learn much new stuff.

- Prof. Jeff Zhang's course titled "CEN/EEE 598: Advanced Hardware and Systems for Machine Learning"
- Prof. Aviral Shrivastava's course titled "CEN/CSE 598: Machine Learning Accelerator Design"

There is also quite a bit of overlap of this course with the following course. So, please look through the contents carefully to decide if you'd like this take this course.

- Prof. Jiaqi Zu's course titled "EEE 598: Algorithm/Hardware Co-Design and Design Automation for Emerging AI Hrdware"

👩‍🏫 Personnel

Instructor

1	**Name**	Aman Arora
2	**Office**	Centerpoint 203-09
3	**Office hours**	On Canvas
4	**Office hours location**	My office or via Zoom
5	**Email**	aman.kbm@asu.edu
6	**Sandwich hour**	Ike's Sandwiches, Thursday 12-1pm

Teaching Assistants and Graders

	Role	Name	Office	Office Hours	Email
1	TA	Kaustubh Mhatre	Centerpoint 203-12BA	On Canvas	kmhatre@asu.edu

🎯 Course Description

Machine learning (ML) has become ubiquitous and is currently a dominant computing workload. This course covers design of hardware and software for training and inference in ML systems. Hardware choices for machine learning include CPUs, GPUs, FPGAs, and ASICs. Tradeoffs in implementing training and inference workloads using these different compute paradigms will be explored. Emerging ML accelerators will be studied. Students will read research papers, present their learnings in class, and complete a project.

Course Objectives

- Review of computer architecture and arithmetic
- Review of ML and Deep Learning (DL) networks
- Computing needs for ML/DL (Characterization of ML/DL workloads), including compute and memory demands for both training and inference
- Overview of acceleration architectures for ML
- Hardware-software co-design and design space exploration of accelerators
- Performance analysis and benchmarking of ML/DL workloads
- Optimization methods for improving the efficiency of ML/DL workloads
- Compilation/mapping of ML/DL workloads
- Case studies of common acceleration platforms
- Case studies of acceleration of common DL networks
- Misc topics (time-permitting): TinyML, Processing-In-Memory, Multi-device execution of ML/DL workloads, Neural Architecture Search, Sustainability

Learning Outcomes

This course will help students gain a solid understanding of the fundamentals of designing machine learning accelerators and relevant cutting-edge topics. After doing this course, students will be able to understand:

- Understand the role and importance of ML acceleration.
- Use libraries and frameworks for designing and programming ML accelerators.
- Efficiency challenges brought by large DL networks
- Characterization of ML/DL workloads
- Differentiate workings of accelerators for different domains such as computer vision, language processing, and recommendation systems.
- The programming challenges brought forth by machine learning accelerators, especially with algorithmic and hardware specializations
- The need for accelerator-aware neural architecture search.
- Various optimizations for efficient inference on accelerators.
- Recent accelerator systems developed by industry and their trade-offs.

⏮️ Enrollment Requirements (Prerequisites)

- Computer Architecture (CSE 420 or CSE 520 or equivalent)
- Programming (e.g., courses like CSE 100/110/205/220). Having programmed in Python before is helpful.
- Machine Learning (e.g., courses like CSE 475, CSE 471, CSE 571, CSE 574, CSE 575, CSE 576, EEE 598/591: Machine/Deep Learning)

📚 Readings

Textbook

	Column 1	Column 2
1	Name	Efficient Processing of Deep Neural Networks
2	Authors	Vivienne Sze , Yu-Hsin Chen , Tien-Ju Yang , Joel S. Emer

3	ISBN	Softcover: 978-3-031-00638-8 eBook: 978-3-031-01766-7
4	Publisher	Springer Nature
5	Link	Efficient Processing of Deep Neural Networks | SpringerLink

This book will be available in the ASU library. You can access it from the course library organizer link on Canvas.

Research papers

Several research papers from top hardware and software conferences will be assigned during this course. Reading them and assimilating them will be essential part of this course.

Supplemental Materials

- Instructor notes/slides available on Canvas
- Deep Learning Systems: Algorithms, Compilers, and Processors for Large-Scale Production
  - Deep Learning Systems: Algorithms, Compilers, an...
- Neural Networks and Deep Learning
  - Neural networks and deep learning

🏆 Grading

Components

Quizzes

There will be quizzes throughout the semester. Quizzes will be based on the assigned readings (papers or book chapters or blogs). Students can expect about 16 quizzes throughout the semester.

> 📌 No makeups will be allowed for quizzes. However, 25% of the quizzes, in which you have the lowest score, will be dropped while calculating the final grade.

Lab exercises

There will be 4 lab exercises throughout the semester. Each lab will be about 2 weeks. Labs will involve hand-on experimentation with DL frameworks, performance analysis, design space exploration, and optimization.

> 📌 For lab assignments, you will get an extra credit of 5%, if you submit at least 24 hours in advance of the deadline. You will get another 5% extra credit if you submit at least 48 hours in advance of the deadline. That is, you can get an overall extra credit of 10% for a lab if you submit 48 hours before the deadline. Note that this is only applicable to labs. The maximum extra credit for a lab is capped at 10%.

> 📌 Late submissions after the deadline are allowed for lab assignments. However, a penalty of 15% per twenty-four hour period elapsed after the deadline will be applied (up to a max of 4 such periods). An exception can be made for medical reasons (including COVID-19) or emergencies, given that an email notice with proof of record is provided in advance.

Paper presentations

During the last month of the semester, the roles will reverse. Each week a few research papers from top ML and computer architecture conferences will be assigned for reading. A team of students will present the paper in class and there will be a follow-up discussion. All students will undertake a quiz (accounted for in the 'Quizzes and Surveys' grade).

Class project or Term paper

Students will choose to either work on a class project or a term paper. Students will work in teams. The class project is for those who are more inclined towards hands-on work, whereas a term paper will involve more reading. Eventually, whether the students choose a class project or a term paper, they will present their learnings/results in front of the class. These presentations will be held during the last week and during the allotted final

exam slot (there will be no final exam).

> 📌 The grading for class project or term paper presentations will include a peer-graded component and an instructor-graded component.

Class participation

Participation in class is highly encouraged. This includes things like attending class regularly, engaging in discussions, listening actively, answering questions, preparing for class by completing readings, etc. A few surveys will also be conducted. Surveys will have questions to get feedback based on your experience during the course. Participating in the surveys will count towards class participation.

Breakdown

	Component	Points
1	Quizzes	20%
2	Lab exercises	40%
3	Paper presentations	10%
4	Class project or Term paper	25%
5	Class participation	5%

Scale

1	**A**	[92.5-100)%
2	**A-**	[85-92.5)%
3	**B+**	[80-85)%
4	**B**	[75-80)%
5	**B-**	[70-75)%
6	**C+**	[65-70)%
7	**C**	[60-65)%
8	**D**	[50-60)%
9	**F**	<50%

> 📌 If you see a discrepancy in your grade for any assignment, you can request regrading upto 1 week (7 calendar days) after the date you received the grade. Any regrading requests after 7 calendar days will not be entertained.

📢 Modality Details

Lectures

Lectures are the primary mode of teaching for this class. By default, all lectures will be held in-person and not on Zoom. The instructor will try to record all lectures. No guarantees, however. Sometimes technical issues can happen. If you have concerns about being recorded, please contact the course instructor. Recordings of lectures (or other additional sessions such as reviews for exam preparation) will be available on Canvas for all students to access for reviewing course materials.

The instructor may prerecord one or two lectures in case of their unavoidable absence. The prerecorded lectures will be made available on Canvas at regular class times.

In case you cannot attend class in person as a result of illness or possible exposure to infectious disease or travel, you must inform the instructor to enable Zoom and/or record the lecture no later than one hour prior to the start of the class meeting. Please contact the instructor and TA via email explaining your situation.

The first 2-2.5 months will involve the instructor lecturing and presenting in the classroom. The next 1 month will involve the students lecturing and presenting in the classroom.

Canvas

The primary mode of communication for the course will be Canvas. All lecture slides, lab documents, paper readings, etc. will be posted on Canvas. Please check the Canvas notifications periodically for any last minute announcements. Many updates may only appear on Canvas. This current document (syllabus) is also linked on Canvas. Any updates to this document will be announced through Canvas announcements.

Discord

Discord will be used for discussions. If you have questions, post them on Discord. General questions about labs, paper readings, quizzes, etc. will not be answered by email. Post them on Discord. This allows other students to respond if they know the answer enabling peer-to-peer learning, and also helps other students who may have the same question.

> 📌 Before posting a new question on Discord, search through the existing threads on there. You may find the answer to your question.

If you are not comfortable asking a question on Discord, the next step is to email the TA. The TA may involve the instructor in the discussion if required. If you have a question that requires communicating with the instructor, please send email. I will try to respond to emails within 48 hours.

🗓️ Course Schedule

> 📌 See home page of the course on Canvas

🔬Labs

Lab Administration

- Labs will be released by publishing the lab manual (and supplementary files) on Canvas. The instructor and TA will explain the lab assignment briefly in the lecture.
- Each lab will be released (and will be due) during the weekend (Fri-Sun).

Lab Evaluation

On the due date/time (specified on Canvas), you will upload a report and code (details specified in each lab manual) via Canvas.

Lab Exercises

	Tentative Topic
1	DNN workload characterization and running DNNs on CPU
2	DNN profiling and performance evaluation and running DNNs on GPU
3	Hardware accelerator modeling and design space exploration
4	DNN optimization (sparsity, compression, pruning, quantization)

For lab assignments, you will work with a partner (i.e. you will work in teams or groups of 2). Working in teams is an important part of being an engineer. Lab partners have to contribute equally to the assignment. You will have a single code base, and both of you will submit the files on Canvas.

> 📌 During the submission, you will be asked to mention how your partner contributed and how you contributed. Different grades may be assigned to the team members.

📌 You can choose lab partners on your own or use Discord to look for partners.

Hardware and Software

Any hardware or software required for the lab assignments will be made available online.

They are identical boards. A7 is just a rebranding of 4-DDR.

👩‍💻 Class Project

Students are welcome to propose and choose any project as long as it is relevant to the topics of the course. This involves doing hands-on experiments and generating results. Students will form groups of 3 for the class project. The amount of work expected for the project will be ~1.5-2 lab assignments for each student. A project report is the final deliverable. It is expected to be atleast 8 pages in IEEE two-column style (including figures, but not including references).

The various stages of class project will be:

	Milestone	Grade breakdown	Title	Timeline
1	1	3%	Project interest document (topic + team)	Due about 4 weeks into the semester
2	2	6%	Project proposal	Due about 8 weeks into the semester
3	3	9%	Project progress report	Due 4 weeks before the end of the semester
4	4	27%	Project presentation	Presentations will be done during the last class and the final exam slot
5	5	55%	Project final report	Due 3 days before the grade submission deadline

📌 During the submission, you will be asked to mention how your partners contributed and how you contributed

Here are some suggested ideas for the project:

- Reproducing something from a research paper
- Characterization (performance analysis) of an application/workload on a system
- Performance comparison of two applications on a system.
- Performance comparison of an application on two systems.
- Developing new software (libraries) on existing systems

📜 Term Paper

An alternative to doing class project is to do a term paper. Think of it as writing a tutorial/survey. You will explain the topic in detail (by synthesizing the information you obtain by reading various materials), along with summarizing the state-of-the-art (obtained by surveying the latest research papers on that topic). Students will form groups of 2 for the term paper. The amount of effort spent on the term paper is expected to be ~1.5-2 lab assignments per student. The length of the paper should be at least 10 pages in IEEE two-column style (including figures, but not including references). You should reference all the materials you used to create the term paper. At least 12 references should be research papers, other references can be surveys or blogs.

The various stages of term project will be (grade distribution of each milestone is mentioned in percentages in brackets):

	Milestone	Grade breakdown	Title	Timeline
1	1	3%	Topic interest document	Due about 4 weeks into the semester
2	2	6%	Literature survey list	Due about 8 weeks into the semester
3	3	9%	Paper organization report	Due 4 weeks before the end of the semester
4	4	27%	Paper presentation	Presentations will be done during the last class and the final exam slot
5	5	55%	Final paper submission	Due 3 days before the grade submission deadline

For a term paper, you are not free to choose any topic, but can select one of the following topics. Discuss the topic with the instructor before you decide.

- Neural Architecture Search
- Acceleration of Graph Neural Networks
- Deep dive into a DNN compiler
- Deep dive into multi-device DNN acceleration
- Fault tolerance + DNN accelerators
- Security + DNN accelerators
- Quantum + DNN acceleration
- Photonics + DNN acceleration
- Federated, on-line and on-device learning
- Sustainability

📽️ Paper Presentations

In the second half of the class, teams of students will make presentations to the class. These presentations will be based on research papers assigned for readings. These papers will be case studies of ML acceleration, spanning two categories: popular hardware accelerators and acceleration of popular DNNs.

Two research papers will be assigned for each lecture. Students will be asked to create teams of three (can be the same teams as class projects). Each team will sign up for a slot to present a paper to the class. This will be a 20-22 minute presentation, followed by a discussion of about 10-12 minutes. The papers will be from top computer architecture and ML conferences. PDFs of the papers will be provided.

Tentative topics on which the papers will be assigned:

- Google TPU (ASIC based)
- NVIDIA Ampere GPU
- Microsoft Brainwave (FPGA based)
- Transformer acceleration
- CNN acceleration
- Recommender system acceleration
- Training acceleration

❓ Quizzes

There will be several quizzes throughout the semester. Some quizzes will be at-home and some quizzes will be in-class. They will be based on the assigned readings - book chapters or research papers. Reading and understanding research papers can be hard in the beginning. But you will

develop the required skills as you go through the semester.

> 📌 There is a paper about how to read a paper. Please read this first: How to read a paper | ACM SIGCOMM Computer Communication Review

While reading research papers, I recommend highlighting important sentences and keep written notes on:

- What's the problem statement and the motivation for this work?
- What are other related works in this area?
- What methodology (tools, benchmarks, etc) did they use?
- What results did they achieve?

Additionally, keep a critical mind. Keep an eye out for:

- Any issues with the evaluation process?
- What can be done to improve this work?
- Any questions that they didn't answer?

🫶 Similar courses at other universities

Here are some links to similar courses at other universities to get an idea of what to expect from this course:

	University	Link
1	Stanford	Hardware Accelerators for Machine Learning (CS 217)
2	Georgia Tech	Hardware Software Co-Design for Machine Learning Systems
3	Cornell Tech	ML Hardware and Systems
4	SFU	SFU Computer Architecture Class Course Information
5	Washington University	Hardware/Software Co-Optimization for Machine Learning

Absence Policies

- Attendance and participation in class activities is an essential part of the learning process, and students are expected to attend lectures and recitation sections regularly. Some absences are, however, unavoidable. Excused absences for classes will be given without penalty in the case of:
  a. a university-sanctioned event [ACD 304-02];
  b. religious holidays [ACD 304-04; a list can be found here Religious Holidays and Observances | Educational ...];
  c. work performed in the line-of-duty according to [SSM 201-18]; and
  d. illness, quarantine or self-isolation related to illness as documented by a health professional.
- Students who request an excused absence must follow the policy/procedure guidelines.
- Excused absences do not relieve students of responsibility for any part of the course work required during the period of absence.
- Anticipated absences for university-sanctioned events, religious holidays, or line-of-duty activity should be communicated to the instructor by email at least 5 working days before the expected absence.
- Extensions on due dates for assignments or taking the midterm outside the scheduled date and time will be allowed only medical reasons. Absences for illness should be documented by a health professional and communicated to the instructor as soon as possible in writing.
- Midterm or final exam will not be offered to be done in an online format, unless it is for medical reasons (with proper proof from a healthcare professional).

🔁 Course Evaluations

This course will be evaluated using the standard ASU course/instructor evaluation forms during the last week of class. In addition, the instructor will use a few other in-class evaluations/surveys.

🚳 Expectations

Students in this class are expected to acknowledge and embrace the FSE student professionalism expectation located at:

ASU Student Professionalism – ASU Engineering Studen...

Classroom behavior

- All cellphones must be turned off and put away. No browsing of the web and non-class related activities using computers or cellphones is permitted. This will be strictly enforced and a student using a cellphone or a laptop doing non-class activities during class will be asked to leave.
- Laptops/Tablets may be used:
  - To take notes
  - To follow along the lecture slides
  - To follow a demonstration by the instructor
  - To answer quizzes/surveys/etc.
- No recording devices are allowed.
- Talking during the lecture and other disruptive behavior may result in expulsion from the classroom. If you have a question, ask the instructor.
- Any violent or threatening conduct by a student in this class will be reported to the ASU Police Department and the Office of the Dean of Students.
- If the instructor is late to the class, they will make a best attempt to inform the students through the TA or graders. However, if there is no communication, the students may leave after waiting for 15 minutes.

Engineering Professionalism

Picture yourself as a future engineer. Whether you enter academia or industry, your success will depend, in part, on how others perceive you. While your creativity, design abilities, and problem-solving skills will earn recognition, a significant aspect of your achievements will hinge on how your peers view your professionalism alongside your engineering capabilities. Even if you possess outstanding technical skills, failure as an engineer can result if you do not conduct yourself in a professional manner. Here are some key attributes of professional behavior that are relevant to this class:

- **Punctuality**: Consistently arriving late to work can signal a lack of commitment and a lack of respect for others' time. In class, being late can disrupt the learning environment. If you have to leave class early or arrive late, do it in a manner least disruptive to the class.
- **Appropriate conduct in group settings**: Demonstrating courtesy to speakers, giving others the opportunity to voice their opinions, and actively listening are all expected in any professional setting. During lectures, refrain from engaging in disruptive activities or conversations with other students. If you have questions or need clarifications, address them to the instructor.
- **Preparedness**: Imagine the attending an important meeting at work without being familiar with the subject matter. Being prepared for meetings, presentations, or discussions is an essential aspect of engineering professionalism. When you come during office hours (to see a TA or see the instructor), come prepared with a list of questions. Before you arrive for a lecture, spend 10 minutes browsing through the lecture slides.
- **Effective communication**: The ability to communicate clearly, both orally and in writing, is crucial. Whether it is your lab reports or homeworks, make sure they are legible and neat.
- **Comprehensive documentation**: Your work should be accurate, thorough, and neatly organized. Others should be able to understand the content of your documentation easily. Document for others (like homeworks and lab reports), but also for you. Note taking is a essential and powerful skill. I highly encourage that you take notes during class.
- **Honesty**: Contribute your fair share and never take credit for the work of others. Being honest and dedicated to your group's efforts is essential for fostering a positive team dynamic. Do not expect your team mates to do the hard work.
- **Dedication**: Successful teams rely on individuals who are committed to achieving common goals. Demonstrating dedication and doing whatever it takes to accomplish these goals sets apart the successful from the less successful team members.

My objective is to make you a good hardware designer, a good computer engineer and to help you get a job that you enjoy.

👥 Diversity, Equity and Inclusion

It is my intent that students from all diverse backgrounds and perspectives be well served by this course, that students' learning needs be addressed, and that the diversity that students bring to this class can be comfortably expressed and be viewed as a resource, strength and benefit to all students.

This is a safe space, with no judgement related to color of your skin, your ethnicity, your financial status, your accent, etc. Everyone should respect each other, including your classmates, your TA, your graders and the instructor.

Please come to me at any time with any concerns.

Well-Being

Your well-being is a priority in this class – above even learning the content. College and its coursework are difficult for everyone, and this difficulty can become overwhelming. **We all learn differently, and everyone struggles sometimes.** If you are struggling to cope, please reach out to me, the TA, a friend, a family member, or another resource. I'm not trained as a professional counselor, but I can help you get the support you need.

Please do not hesitate to reach out. My job as an instructor is to help ensure everyone in this class learns, feels cared for, and knows that they belong in this course and at ASU.

Student Health and Wellness Services

Student Services | Educational Outreach and Student Services

Career and Professional Development Service

https://eoss.asu.edu/studentservices

ASU Counseling Services | Educational Outreach and Student Services

https://eoss.asu.edu/counseling

EOSS Resources Hub | Educational Outreach and Student Services

Sun Devil life is about connections, traditions and lessons that make your time at ASU unforgettable. Find what helps you to continue being the best you can be.

https://eoss.asu.edu/resources-hub

🏛️ Academic Integrity

All engineering students are expected to adhere to the ASU Student Honor Code and the ASU academic integrity policy, which can be found at Student Policy | Office of the University Provost | A...). Students are responsible for reviewing this policy and understanding each of the areas in which academic dishonesty can occur. If you have taken this course before, you may not reuse or submit any part of your previous assignments without the express written permission from the instructor.

All student academic integrity violations are reported to the Fulton Schools of Engineering Academic Integrity Office (AIO). Withdrawing from this course will not absolve you of responsibility for an academic integrity violation and any sanctions that are applied. The AIO maintains a record of all violations and has access to academic integrity violations committed in all other ASU college/schools.

ALL violations, regardless of their severity, will be reported to the University and will be prosecuted and will result in a E grade on the submitted work and possibly an E as a final grade. This policy applied to all submitted work, i.e., ALL submitted work must be CREATED by you and ONLY you.

AI policy

All assignments should be fully prepared by the student (or students; in case of a lab where lab groups are permitted). This includes writing code for your labs, solving homework problems, etc. Developing strong competencies in the skills associated with this course will prepare you for success in your degree pathway and, ultimately, a competitive career. Therefore, the use of AI tools (e.g. ChatGPT) to complete any aspect of assignments or exams of this course will be considered academic dishonesty and a violation of the ASU Academic Integrity Policy. If you have questions about what constitutes a violation of this statement, please contact the instructor.

When you go work in the industry, you may be allowed to use AI tools to write code. However, you will need to ensure that the code generated by the AI is correct, you will need to iterate over the code with the AI, you will need to modify the code generated by the AI. AI will be able to generate code for small parts of large industrial designs, and you will need to put them together. The concepts you learn in this will provide you skills to excel at these tasks. If you use AI tools in this class, then you won't understand the basic concepts to be able to use AI efficiently in the future.

Most importantly, you are here to learn. And humans learn by doing. If you obtain your code from the AI, then you are not learning.

Another reason to not use AI tools is that if you enter the lab or homework questions into an AI tool like ChatGPT, that content becomes public. However, note that all slides, lab manuals, homeworks, etc. are copyrighted. So, by putting in content into the AI tool, you are violating the copyright.

© Copyright

Course content, including lectures, are copyrighted materials and students may not share outside the class, upload to online websites not approved by the instructor, sell, or distribute course content or notes taken during the conduct of the course (see ACD 304–06 https://policy.asu.edu/, "Commercial Note Taking Services" and ABOR Policy 5-308 F.14 https://public.azregents.edu/Policy Manual/5-308-Student Code of Conduct.pdf for more information).

You must refrain from uploading to any course shell, discussion board, or website used by the course instructor or other course forum, material that is not the student's original work, unless the students first comply with all applicable copyright laws; faculty members reserve the right to delete materials on the grounds of suspected copyright infringement.

🥊 Policy against threatening behavior

(Per the Student Services Manual, https://policy.asu.edu/ SSM 104–02) Students, faculty, staff, and other individuals do not have an unqualified right of access to university grounds, property, or services. Interfering with the peaceful conduct of university-related business or activities or remaining on campus grounds after a request to leave may be considered a crime. All incidents and allegations of violent or threatening conduct by an ASU student (whether on- or off-campus) must be reported to the ASU Police Department (ASU PD) and the Office of the Dean of Students.

♿ Disability Accommodations

Suitable accommodations will be made for students having disabilities. Students needing accommodations must register with the ASU Student Accessibility and Inclusive Learning Services office and provide documentation of that registration to the instructor. Students should communicate the need for an accommodation in sufficient time for it to be properly arranged. See ACD 304-08 (https://policy.asu.edu/) Classroom and Testing Accommodations for Students with Disabilities.

⛔ Harassment and Sexual Discrimination

Arizona State University is committed to providing an environment free of discrimination, harassment, or retaliation for the entire university community, including all students, faculty members, staff employees, and guests. ASU expressly prohibits discrimination, harassment, and retaliation by employees, students, contractors, or agents of the university based on any protected status: race, color, religion, sex, national origin, age, disability, veteran status, sexual orientation, gender identity, and genetic information.

Title IX is a federal law that provides that no person be excluded on the basis of sex from participation in, be denied benefits of, or be subjected to discrimination under any education program or activity. Both Title IX and university policy make clear that sexual violence and harassment based on sex is prohibited. An individual who believes they have been subjected to sexual violence or harassed on the basis of sex can seek support, including counseling and academic support, from the university. If you or someone you know has been harassed on the basis of sex or sexually assaulted, you can find information and resources at https://sexualviolenceprevention.asu.edu/faqs.

As a mandated reporter, I am obligated to report any information I become aware of regarding alleged acts of sexual discrimination, including sexual violence and dating violence. ASU Counseling Services, https://eoss.asu.edu/counseling, is available if you wish discuss any concerns confidentially and privately.

📷 Photo Requirement

Arizona State University requires each enrolled student and university employee to have on file with ASU a current photo that meets ASU's requirements (your "Photo"). ASU uses your Photo to identify you, as necessary, to provide you educational and related services as an enrolled student at ASU. If you do not have an acceptable Photo on file with ASU, or if you do not consent to the use of your photo, access to ASU resources, including access to course material or grades (online or in person) may be negatively affected, withheld or denied.