

CPTS 580: Computer Vision

Instructor: Yan Yan

Fall, 2024

Class Hours: Tue/Thu, 1:30PM - 2:45PM

Class Room: EME B46

Office Hours: Tuesday, Thursday 12:30PM - 1:30PM

Office: EME 123

Contact:

Preferred: Canvas Message

Alternative: Email yan.yan1@wsu.edu

1 Course Description

This course provides an introduction to computer vision (CV), an important sub-field in artificial intelligence (AI). It mainly includes three topics that are useful to construct and analyze CV systems: (i) basic low-level CV techniques, (ii) fundamental machine learning (ML) paradigm, and (iii) modern deep neural networks (DNNs) and their application to CV tasks. These three topics broadly cover both theoretical foundations of basic ML and deep learning, as well as the hands-on implementation for applications of computer vision, which targets the process and understanding of visual data (e.g., images, videos, etc.). Particularly, this course also provides a graduate-level understanding to the connections from fundamental ML theory to practical CV approaches, and from low-level CV techniques to high-level CV system architectures. For the low-level CV techniques, we review classical image filtering, corner detection and hand-crafted features. For the fundamental ML theory, we review standard linear classifier, maximum likelihood estimation, probably approximately correct learning and optimization. For the modern DNNs, this course covers basic elements in neural network structures, and more sophisticated architectures for the CV tasks of object detection, segmentation, video classification and generative models.

2 Course Sections/Topics

- Basic CV
 - Image filtering
 - Corner detection
 - Image descriptors
- Fundamental ML

- Linear classifier
- Maximum likelihood estimation
- Probably approximately correct learning
- Optimization for ML models
- Basic Neural Networks
 - Elementary neural network layers
 - Composition of functions
 - Units in neural networks
 - Backpropagation
 - Softmax classifier
- Basic Convolutional Neural Networks
 - Convolution operation (and image filtering)
 - Structure of CNNs
 - Practical training techniques
- More Advanced CV Tasks
 - Object detection
 - Segmentation (pixel-wise labeling)
 - Video classification
 - Generative model

3 Course Objectives

- Understand machine learning basics: the pipeline of machine learning, what can be learned, how to learn, and statistical/optimization/computational challenges
- Understand the benefits of deep neural networks for learning representations from raw data over traditional machine learning algorithms, and the computational/statistical challenges of learning deep neural networks.
- Understand the key concepts and challenges in computer vision tasks
- Able to build machine learning models or deep neural networks to deal with computer vision problems.

4 Textbooks

- Mehryar Mohri, Afshin Rostamizadeh, and Ameet Talwalkar. “Foundations of Machine Learning”. MIT Press, Second Edition, 2018. Online: <https://cs.nyu.edu/~mohri/mlbook/>.
- Ian Goodfellow, Yoshua Bengio, and Aron Courville. “Deep Learning”. MIT Press. Online: <https://www.deeplearningbook.org/>.
- Aston Zhang, Zachary C. Lipton, Mu Li, Alexander J. Smola. “Dive into Deep Learning”. Online book with code: <https://d2l.ai/>.

5 Course Materials

All course materials can be found on Canvas "Home" page or "Modules" page. Slides will be grouped into sections as listed in Section 2. Assignments for each sections will also be open in the corresponding modules.

6 Grading Policy

6.1 Assignments (60% grade)

- HW1: CV fundamentals, 10% of final grade
- HW2: ML fundamentals, 10% of final grade
- HW3: NN, 10% of final grade
- HW4: CNN, 10% of final grade
- HW5: detection, 10% of final grade
- HW6: segmentation, 10% of final grade

All coding assignments will be based on [Python](#), [Pytorch](#) and [Jupyter notebook](#). This course focuses on neural networks, and does not cover the particular topics about how to program in Python and how to use Pytorch and jupyter notebook. However, there are many good resources as guidelines for learning them. For example, [a tutorial for Python and numpy](#) provided by CS231n from Stanford University (or Google what you need, Appendix B).

When submitting coding assignments:

For coding assignments, it is extremely important to save your running records and make sure they can be displayed in your submissions. Each time when submitting, double check you have uploaded (i) a .ipynb file for Jupyter notebook that can be re-produced by TA, and (ii) a .pdf or .html file converted from the .ipynb file that tracks all your running records to show that you have completely finished this assignment. Note that uploading these 2 files are sufficient, and zipping the 2 files is NOT needed.

6.2 Mid-Term Quiz (10% grade)

- Format: Canvas Quiz, online. Open book: textbooks and slides can be used, but cannot search questions/answers on internet.
- Types of questions: multiple choice (choose one answer), multiple answers (choose at least one answer), ordering, True/False, etc.
- Sections covered: fundamental CV, fundamental ML, fundamental NN.
- Grading method: Canvas auto grading.

6.3 Course Project (4% + 8% + 8% + 5% = 25% grade)

A project on CV tasks for any problems, including the image classification, object detection, segmentation and generative model mentioned in this course, or any other computer vision problems not included by this course. The following assignments are graded components for course project.

- Project proposal (4% of final grade)
 - 3-page written proposal
 - Due on 09/27
- Project mid-term report (8% of final grade)
 - 6-page written report summarizing your mid-term exploration and findings
 - Due on 10/25
- Project final report (8% of final grade)
 - 8-page written report summarizing all your mid-term exploration and findings
 - Due on 11/22
- Project code (if applicable) format: Jupyter-lab file
- Project presentation (5% of final grade) with QA sessions that summarize and highlight your final report

Course project can be done with **a team up to 2 students**. Below list the required components in the written project proposal, report and presentation.

- **Project proposal** (4% grade): (i) determine team members (if applicable), (ii) decide what problem you are going to investigate, and (iii) submit a written proposal to describe your idea and plan.

Proposal format: 3 pages (title, abstract, main text) + unlimited reference page + unlimited appendix page.

Format: PDF (exported from Microsoft Word or compiled from LaTeX, e.g., [a LaTeX template](#)).

What should be included:

- Title: give a name to the project to highlight the significance, what contributions this project can have. The title should be short and relevant to your project.
 - Team members and student IDs. You still need to ensure that you create a group for your team and all team members have been included in one group. Otherwise, team members not included in the group will NOT receive credit. See Appendix A for creating group on Canvas.
 - Abstract: a paragraph that introduce a number of key points at least contain: (i) what is your selected CV problem, (ii) why do you choose this problem for your project, (iii) what unique advantage/improvement will your project have, (iv) what are existing relevant methods for solving your considered problem, (v) what is your plan of the key steps for completing your project.
 - Introduction: a detailed explanation/elaboration of Abstract on the above 5 key points. Typically, one key point can be elaborated to one paragraph. If necessary, you may also need to cite others' paper/work (e.g., to elaborate the key point (iv)) to give a specific comparison between your project and existing methods.
 - Literature review: any references that are important and relevant to the project, e.g., some techniques in the project that are inspired by other people's work including tutorials, Git repository, research publication, blogs, etc. Introduce more about what inspired the project.
 - Technical plan: a detailed elaboration of techniques used in this project (based on the techniques in Section "Literature review", if any). This plan can be described by a flowchart or any other figures.
 - Reference: if anything inspired your project or is relevant to your project, a citation is required.
- **Project mid-term report** (8% grade): (i) determine team members (if applicable), (ii) decide what problem you are going to investigate, and (iii) submit a written proposal to describe your idea and plan.

Proposal format: 6 pages (title, abstract, main text) + unlimited reference page + unlimited appendix page.

Format: PDF (exported from Microsoft Word or compiled from LaTeX, e.g., [a LaTeX template](#)).

What should be included:

- Title (**Can be updated**): give a name to the project to highlight the significance, what contributions this project can have. The title should be short and relevant to your project.
- Team members and student IDs. You still need to ensure that you create a group for your team and all team members have been included in one group. Otherwise, team members not included in the group will NOT receive credit. See Appendix A for creating group on Canvas.
- Abstract (**Can be updated**): a paragraph that introduce a number of key points at least contain: (i) what is your selected CV problem, (ii) why do you choose this problem

for your project, (iii) what unique advantage/improvement will your project have, (iv) what are existing relevant methods for solving your considered problem, (v) what is your plan of the key steps for completing your project.

- Introduction (Can be updated): a detailed explanation/elaboration of Abstract on the above 5 key points. Typically, one key point can be elaborated to one paragraph. If necessary, you may also need to cite others' paper/work (e.g., to elaborate the key point (iv)) to give a specific comparison between your project and existing methods.
 - Literature review (Can be updated): any references that are important and relevant to the project, e.g., some techniques in the project that are inspired by other people's work including tutorials, Git repository, research publication, blogs, etc. Introduce more about what inspired the project.
 - Technical plan (Can be updated): a detailed elaboration of techniques used in this project (based on the techniques in Section "Literature review", if any). This plan can be described by a flowchart or any other figures.
 - Intermediate results (NEW): following the plan, what has been done in the project? Elaborate what the team tried for finishing the plan, and what observation and knowledge the team acquired so far. Both theoretical analysis and experimental result can be regarded as the intermediate results.
 - Future work (NEW): to ensure that the proposed plan can be done, what else work remaining for this project? Elaborate what the team plans to try. This section may modify the plan originally proposed in proposal, if any new understanding is established and some adjustment is necessary.
 - Reference: if anything inspired your project or is relevant to your project, a citation is required.
- **Final report (8% grade)**: a report summarizing all the work a team has done, including theoretical or empirical results.

Length limit: 8 pages (title, abstract, main text) + unlimited reference page + unlimited appendix page.

Format: PDF (exported from Microsoft Word or compiled from LaTeX, e.g., [a LaTeX template](#)).

What should be included:

- Title (Can be updated): give a name to the project to highlight the significance, what contributions this project can have. The title should be short and relevant to your project.
- Team members and student IDs. You still need to ensure that you create a group for your team and all team members have been included in one group. Otherwise, team members not included in the group will NOT receive credit. See Appendix A for creating group on Canvas.
- Abstract (Can be updated): a paragraph that introduce a number of key points at least contain: (i) what is your selected CV problem, (ii) why do you choose this problem for your project, (iii) what unique advantage/improvement will your project have, (iv)

what are existing relevant methods for solving your considered problem, (v) what is your plan of the key steps for completing your project.

- Introduction (Can be updated): a detailed explanation/elaboration of Abstract on the above 5 key points. Typically, one key point can be elaborated to one paragraph. If necessary, you may also need to cite others' paper/work (e.g., to elaborate the key point (iv)) to give a specific comparison between your project and existing methods.
- Literature review (Can be updated): any references that are important and relevant to the project, e.g., some techniques in the project that are inspired by other people's work including tutorials, Git repository, research publication, blogs, etc. Introduce more about what inspired the project.
- Technical plan: a detailed elaboration of techniques used in this project (based on the techniques in Section "Literature review", if any). This plan can be described by a flowchart or any other figures.
- Complete Technical Approach (NEW): this new section should present the detailed methods that you explored to deal with the considered problem, so that the key point (iii) can be highlighted, i.e., what unique advantage/improvement your method could accomplish. Figures such as flow charts with rich text captions are highly recommended in this section to illustrate your technical approach. A great example of using all kinds of figures to present technical ideas is ResNet paper¹ by He et al. Moreover, text explanation should be detailed, complete, clear and easy to follow.
- Complete results (NEW): following the plan and explored technical approach, what has been done in the project and what outcomes have been accomplished? Elaborate what the team tried for finishing the plan, and what observation and knowledge the team acquired so far. Both theoretical analysis and experimental result can be regarded as the results. Again, figures are greatly recommended in this section to summarize and highlight the results. You may follow the way how ResNet paper presents their experiment results and uses the results to support their statements.
- Future work after this course (NEW): given what the team tried and learned from the project, is there any idea in the future that can be explored further? Is there any unknown questions and hard-to-understand observation in the task?
- Reference (UPDATED): if anything inspired your project or is relevant to your project, a citation is required, especially when you found new literature that is related to your technical approach.

Note that the final project report can re-use materials from the project proposal and progress report. **NEW** and **UPDATED** sections, i.e., complete technical approach, complete results, future work and reference, should be included in the final report.

What can be *optionally* included:

- Github repository (a URL link) that includes a complete copy of your implementation code (in this way, your code will be open source under a selected license).
- Jupyter-lab notebook that is used to demo how your model can work and display experimental results.

¹arXiv version at <https://arxiv.org/abs/1512.03385>

- Project presentation (5% grade): please present your entire project and summarize your final report in 8 minutes with a 2 minute QA session.

How to select a project?

- Example reports Previous reports from other courses
 - CS229, Stanford: [Fall 2017](#), [Fall 2018](#), [Fall 2019](#), [Spring 2020](#), [Spring 2021](#) (report + poster), [suggested ideas](#)
 - CS221, Stanford: [Spring 2017](#) (titles only)
 - CS231n, Stanford: [Spring 2017](#), [Spring 2022](#) (report + poster)
- Examples from open competition on [Kaggle competitions](#): deep learning applications
 - X-ray image classification/ranking: [CheXpert competition](#) from Stanford ML group
 - Landmark Retrieval: [Google landmark retrieval competition 2021](#)
 - Find extraterrestrial signals in data from deep space: [SETI Breakthrough Listen - E.T. Signal Search](#)
 - Identify the category of foliar diseases in apple trees: [Plant Pathology 2021 - FGVC8](#)
- Or any other ideas

6.4 Attendance (5% grade)

In-class quizzes (randomly), for attendance check purposes. For absence, please see the [university policy for absences](#). Contact instructor to address your absence under this policy.

6.5 Late Work Grading

Late submissions can still be graded under some circumstances.

6.5.1 Written Assignments And Course Project

For written assignments or written part of assignments in the format of Canvas Quiz, and course project (proposal + mid-term report + final report), there is NO late submission policy, and the submission will be closed after due date (the due date for written assignments is firm).

6.5.2 Coding Assignments

For coding assignments or coding part of assignments, TA/instructor will use the following late work policy.

- After the scheduled due date, you will have a 2-calendar-day non-penalty period, allowing any variations of everyone's situation.
- After the 2-calendar-day non-penalty period, the maximum possible grade is capped by 20% per day.

- Maximum 5-calendar-day late submission (i.e., 2-calendar-day non-penalty and 3-calendar-day penalty) is allowed. After that, assignment submission will be closed.
- Example 1: the scheduled due date is Sep 1, 23:59, PT. I submit my assignment at Sep 1, 23:58, PT (1 minute before the scheduled due date). The maximum possible grade I can get is 100 (full points). No action is taken.
- Example 2: the scheduled due date is Sep 1, 23:59, PT. I submit my assignment at Sep 3, 23:58, PT (1 minute before the end of 2-day non-penalty period). The maximum possible grade I can get is 100 (full points). No action is taken. However, the late submission may influence my schedule for following assignments and give me time constraint (I have 6 x HW + 1 x exams + 1 x project this semester).
- Example 3: the scheduled due date is Sep 1, 23:59, PT. I submit my assignment at Sep 4, 00:03AM, PT (4 minutes after the end of 2-day non-penalty period). The maximum possible grade I can get is 80 (= 100 - 20, capped by 20). If my submission originally gets 81 points out of 100, it will be capped to 80. If my submission originally gets 79 points out of 100, it will not be changed.
- Example 4: the scheduled due date is Sep 1, 23:59, PT. I submit my assignment at Sep 6, 23:59, PT (last minute of the 5-day late submission). The maximum possible grade I can get is 40 (= 100 - 3*20). If my submission originally gets 81 points out of 100, it will be capped to 40.
- Example 5: the scheduled due date is Sep 1, 23:59, PT. I would like to submit my assignment at 00:01AM, PT, on Sep 7. I found the submission has been closed, and my submission will not be graded.

6.6 Extension of Assignments

If more time for finishing an assignment is necessary, it is possible to get extension. For example, you have difficulty for finishing assignments on time due to health condition or other emergency condition. In this case, you can request extension with any supporting document such as doctor's appointment records, doctor's note, medical prescription, etc. Request without reasonable supporting document will not be considered.

6.7 Grading Scale to Letter Grade

Table 1 shows the grading scale from percentage to letter grade. Curving grade is possibly considered only if the final grade is very biased and skewed that cannot reflect the performance in this course.

7 Contact Instructor

7.1 Office Hour

Instructor will create office hour appointments on Canvas Calendar, so it is convenient to make appointments to organize and participate in discussion during office hour. See Appendix A for how to sign up an appointment on Canvas Calendar.

94-100.00%	A
90-93.99%	A-
86-89.99%	B+
82-85.99%	B
78-81.99%	B-
74-77.99%	C+
70-73.99%	C
66-69.99%	C-
62-65.99%	D+
58-61.99%	D
0-57.99%	F

Table 1: Grading scale to letter grade

A walk-in office hour is also possible, but not guaranteed, so signing-up a time slot on Canvas Calendar is highly recommended.

7.2 After Class And Office Hour

A *preferred* way to contact instructor after class is to send a Canvas Message to instructor. Check Appendix A about how to send a Message on Canvas to a user in your course. Most Messages on Canvas will be replied within 3 workdays, so that is the recommended/preferred way to guarantee response on time. (Instructor may sometimes miss your message during some busy period of a semester, so please send a second message if not getting response in 3 workdays).

An *alternative* way is to send an email to instructor and TA, but we still recommend using the preferred way for better organization of communication.

8 Course Schedule

The more detailed course schedule can be found on Canvas "Course Schedule" page or "Home" page. The schedule is tentative and subject to change.

8.1 Class Meeting

Week 01, 08/19 - 08/23: Course overview, syllabus, basic CV techniques

Week 02, 08/26 - 08/30: Basic CV techniques

Week 03, 09/02 - 09/06: Basic ML techniques

Week 04, 09/09 - 09/13: Basic ML techniques

Week 05, 09/16 - 09/20: Basic Neural Networks

Week 06, 09/23 - 09/27: Basic Neural Networks

Week 07, 09/30 - 10/04: Basic Neural Networks, CNNs

Week 08, 10/07 - 10/11: CNNs

Week 09, 10/14 - 10/18: CNNs

Week 10, 10/21 - 10/25: Object detection

Week 11, 10/28 - 11/01: Segmentation

Week 12, 11/04 - 11/08: Video classification

Week 13, 11/11 - 11/15: Generative models

Week 14, 11/18 - 11/22: Open research problems in CV

Week 15, 11/25 - 11/29: **No class – Thanksgiving vacation**

Week 16, 12/02 - 12/06: project presentations

8.2 Assignments

- HW1: CV, open on 08/28, due on 09/17
- HW2: ML, open on 09/11, due on 10/01
- HW3: NN, open on 09/25, due on 10/15
- HW4: CNN, open on 10/09, due on 10/29
- HW5: Object detection, open on 10/23, due on 11/12
- HW6: Segmentation, open on 11/06, due on 11/26

8.3 Course Project

- Project proposal due on 09/29
- Project final report due on 11/26

9 Academic Integrity Policy

We must ensure that academic honesty is upheld. All graded assignments/exams must be done independently. If you need any assistance, please contact the instructor or TA. Learn more about [Plagiarism](#), [Unauthorized Assistance](#), [Fabrication](#), [Acts of dishonesty](#), etc., by visiting [WSU's Academic Integrity Policy](#).

Particularly, if the submitted assignments are built on any existing work (e.g., a blog, a tutorial, etc.), a complete and appropriate acknowledgement/citation is required to identify what the original part your submission has done compared to the existing one. Without a complete and appropriate acknowledgement/citation if built on other sources, the submitted assignments will receive failing grade and will be considered as plagiarism.

9.1 Can I Use Generative AI?

Generative AI, such as ChatGPT, should NOT be used to help write the code for any graded assignments/exams. However, generative AI can be used to search any questions that you have when you attempt complete your assignments, analogous to using Google for searching your general questions, such as “what is the network architecture of ResNet”. Such general questions should not be directly related to the specific assignment questions, e.g., “please help me write Python code to build and train a ResNet”.

Generative AI must NOT be used during any exam.

9.2 Can I Use Code from Internet?

You may use others' code in parts of your assignments, if you (i) give a complete and appropriate acknowledgement/citation in your submission to the source where you used the code, including the URL links, etc., and (ii) highlight which part of your submission is your original work that is different from the referred source.

However, TA/instructor will grade your submissions depending on how much major task you complete by your own. Therefore, using too much external code in your submission can still result in bad grades.

A Canvas Tutorials

- Send/reply Canvas Message: [this tutorial](#).
- Canvas appointment sign-up: [this tutorial](#).
- Create a group as a student on Canvas: [this tutorial](#).

B Tips for Coding Assignments

Python, Pytorch, Jupyter notebook are very commonly used tools for building deep neural networks. This course covers principles/rationales of deep neural networks, and does not cover how to use these tools to implement neural networks. 6 assignments will provide this chance to practice hands-on skills using these tools. If you have any specific issues regarding coding, [Google](#) and check out some online tutorials showcasing how these tools can work, such as [Python Numpy Tutorial \(with Jupyter and Colab\)](#) from Stanford CS231n.

C WSU Access Center

Let instructor know if you need special accommodations (such as extended time, make up exams, private or semi-private room for exams, quiet/reduced distraction environment, etc.) for assignments, attendance, midterm exam and final exam as early as possible. You may request special accommodations via visiting [WSU's Access Center](#).