

TO: The Engineering Faculty
FROM: The Faculty of the Elmore Family School of Electrical and Computer Engineering
RE: New undergraduate course – ECE 40801

The Faculty of the Elmore Family School of Electrical and Computer Engineering has approved the following new undergraduate course. This action is now submitted to the Engineering Faculty with a recommendation for approval.

FROM (IF ALREADY OFFERED WITH TEMPORARY NUMBER):

ECE 49595 Computer Vision

3 total credits; 3 credit lecture

Requisites: ECE 20875 and ECE 36800 and ECE 30100 and ECE 30200 and MA 26100 and (MA 26500 or MA 26200)

Fall semesters: 2023, 2024, 2025

TO:

ECE 40801 Introduction to Computer Vision

3 total credits; 3 credit lecture

Requisites: ECE 20875 and ECE 36800 and ECE 30100 and ECE 30200 and MA 26100 and (MA 26500 or MA 26200)

This course is an introduction to computer vision. The course will cover both classical and deep-learning approaches to recognizing objects in images and activities in video. Topics to be covered include scene classification, object localization, segmentation, and activity recognition. The course will cover both the theory behind methods to solve these problems as well as how they are implemented in systems like PyTorch.

Learning Outcomes:

1. construct and train a neural network.
2. construct and train an object detector.
3. construct and train an activity classifier.
4. work in a team to construct a solution to a computer vision problem.

RATIONALE:

This course ran years ago as an experimental course at the graduate-level; however, the need is there for it to be available to undergraduates who are not quite ready for graduate material. It has ran successfully for over three years with increasing enrollment each time.

A handwritten signature in black ink on a light gray background. The signature reads "T.S. Mithuna" followed by a long, sweeping horizontal stroke.

Mithuna Thottethodi

Associate Head of Teaching and Learning

Professor of the Elmore Family School of Electrical and Computer Engineering

Link to Curriculog entry: <https://purdue.curriculog.com/proposal:32691/form>

ECE 40801 Introduction to Computer Vision

Course Information

Course number and title: ECE 40801 Introduction to Computer Vision

Instructional Modality: Face-to-Face

Course credit hours: 3.000

Prerequisites: MA 261 and MA 265 and ECE 20875 and ECE 368 and ECE 301 and ECE 302 or permission of instructor

Prerequisites by topic: multivariate calculus, linear algebra, Python, data structures, rudimentary signal processing including filtering and convolution, probability

Course web page: <https://engineering.purdue.edu/~ece495cv>

Course Brightspace page: <https://purdue.brightspace.com/d21/home/879666>

Instructor(s) Contact Information

Name of the instructor: Jeffrey Mark Siskind

Office Location: BHEE313e

Office Phone Number: 765/496-3197

Purdue Email Address: qobi@purdue.edu

Student Consultation hours, times, and location: T 5:00pm–6:00pm BHEE313e

Course Description

An undergraduate introduction to computer vision. The course will cover both classical and deep-learning approaches to recognizing objects in images and activities in video. Topics to be covered include scene classification, object localization, segmentation, and activity recognition. The course will cover both the theory behind methods to solve these problems as well as how they are implemented in systems like PyTorch.

Learning Resources, Technology & Texts

Computer Vision: Modern Approach 2nd edition, David Forsyth and Jean Ponce

<https://github.com/avinashk94/CVIP/blob/master/David%20A.%20Forsyth%2C%20Jean%20Ponce%20Computer%20vision%20A%20modern%20approach.pdf>

Deep Learning, Ian Goodfellow, Youshua Bengio, Aaron Courville, MIT Press

<https://www.deeplearningbook.org/>

These books are only references. We will not follow these books exactly. Students are not required to purchase these books.

We will cover a recent paper from the literature during each lecture. The papers will be posted at <http://engineering.purdue.edu/~ece495cv> before each lecture. Students should read each paper before the associated lecture and come to class prepared to discuss and ask questions.

Learning Outcomes

Outcome i) An ability to construct and train a neural network. [1]

Outcome ii) An ability to construct and train an object detector. [1]

Outcome iii) An ability to construct and train an activity classifier. [1]

Outcome iv) An ability to work in a team to construct a solution to a computer vision problem. [1]

Assignments

There will be no exams. There will be three homework assignments and a term project. Beyond this, there will be a reading assignment for each lecture. A recent paper will be posted on the course web page for each lecture. Students are expected to read the paper prior to lecture and come to class prepared to discuss and ask questions.

The homeworks will be handed out as follows:

Homework 1	neural networks (outcome i)	in class Friday 1 September 2023
Homework 2	object detection (outcome ii)	in class Friday 29 September 2023
Homework 3	activity classification (outcome iii)	in class Friday 27 October 2023

The due dates for the three homework assignments are as follows:

Homework 1	neural networks (outcome i)	5pm ET Friday 29 September 2023
Homework 2	object detection (outcome ii)	5pm ET Friday 27 October 2023
Homework 3	activity classification (outcome iii)	5pm ET Friday 24 November 2023

There will be no extensions.

Each homework assignment should be submitted as a single zip file that I can run on Ubuntu 22.04.3 LTS (Jammy Jellyfish). You should create a directory whose name is your Purdue Career account name, then put everything you wish to submit in that directory, and then package that as a zip file. That directory should have at least a single file named `run` that is an executable `bash` script that runs your homework. In my case, my Purdue Career account name is `qobi`. So I would do something like this:

```
qobi@sapiencia>cd /tmp/
qobi@sapiencia>mkdir qobi
qobi@sapiencia>cd qobi
qobi@sapiencia>chmod a+x run
qobi@sapiencia>pwd
/tmp/qobi
qobi@sapiencia>ls -l
total 4
-rwxr-xr-x 1 qobi qobi 31 Aug 19 12:55 run*
qobi@sapiencia>cat run
#!/bin/bash
echo my submission
qobi@sapiencia>./run
my submission
qobi@sapiencia>cd ..
qobi@sapiencia>zip -r qobi qobi
  adding: qobi/ (stored 0%)
  adding: qobi/run (stored 0%)
qobi@sapiencia>
```

Obviously, there will be more files and the `run` script will be different for a real homework submission. Homeworks will be submitted through Brightspace. What I will do to grade your homework is something like this:

```
qobi@sapiencia>ls -l qobi.zip
-rw----- 1 qobi qobi 335 Aug 19 12:56 qobi.zip
qobi@sapiencia>unzip qobi.zip
Archive: qobi.zip
  creating: qobi/
  extracting: qobi/run
qobi@sapiencia>cd qobi
qobi@sapiencia>./run
my submission
qobi@sapiencia>
```

one at a time for each submission. This should run in less than five minutes on a machine with twelve cores, 3 Titan RTX GPUs, 32GB RAM, and 1TB free disk space.

You are free to prepare your homework on any machine that you wish. But I must be able to run it in less than five minutes on a machine with twelve cores, 3 Titan RTX GPUs, 32GB RAM, and 1TB free disk space running Ubuntu 22.04.3 LTS (Jammy Jellyfish). I will provide courtesy accounts for all students on the above machine. But please note that resources are limited so all enrolled students must share these resources equitably and fairly. I reserve the right to close any student account if they do not share these resources equitably and fairly, or if they act irresponsibly and attempt to break the system. This machine is not provided by and not supported by the university.

I will also request that the university provide accounts for all enrolled students on the RCAC Scholar cluster. However, I do not know anything about the RCAC Scholar cluster and am unable to provide any support.

The term project (outcome iv) is organized as follows. Student should self organize into teams of two to four enrolled students. Each team should send email to ece49595cv-staff-list@ecn.purdue.edu by 5pm Friday 1 September 2023 with the names and Purdue Career account email addresses of the team members. Each team should then collectively send a one-page pdf project description proposal to ece49595cv-staff-list@ecn.purdue.edu by 5pm Friday 15 September 2023. I will meet with each team during office hours prior to Friday 29 September 2023 to discuss the project description proposal and offer suggestions and guidance. We will have team project presentations with demos in class during the last two or three weeks of class. These will be scheduled later in the semester. You do not need to submit anything for your team term project; you just need to present your work to class in class.

Students will be required to attend all classes with student team term project presentations. Attendance will be taken in those classes.

Grading Scale

On the first day of class, Mr Silver, my eighth grade science teacher, said: "This is an honors class. You all are smart and deserve a good grade. A 90 is a good grade. None of you are perfect yet. So nobody deserves a 100. Thus I will give everyone a 90. Now that we dispensed with grades, we can get down to learning."

I will follow the same policy in this class this semester. The grading policy is simple. This is a 400-level class. You all are smart and deserve a good grade. An A is a good grade. None of you are perfect yet. So nobody deserves an A+. Thus I will give everyone who completes all of the assignments (all three homeworks and term project) by the deadlines, does not violate the anti-plagiarism policy, and attends all of the student presentations an A. Everyone else gets an F. Now that we dispensed with grades, we can get down to learning.

Grades are cheap. An education is expensive. I can manufacture As. When the marginal cost of production goes to zero, the value of the product goes to zero. The As that I give in this course are worthless. I strive to make the education that I give in this course worth the price of tuition. Please don't ask me for a letter of recommendation, to hire you as an RA, or to take you on as a graduate student on the basis of your grade in this course, or even on the basis of the term project/presentation you prepared for this course. You need to do something else that distinguishes you for me to do those things.

Attendance Policy

Nobody ever listens to Grateful Dead studio albums. The Dead existed solely for their live concerts. People went to Dead concerts for two reasons. To hear Jerry Garcia jam, and to get stoned.

I don't condone getting stoned at Purdue or in Indiana. But I do strive to jam like Jerry Garcia. If you want studio albums, you can watch online videos of AI courses taught at Stanford by well-known researchers like Peter Norvig, Sebastian Thrun, Andrew Ng, Daphne Koller, and Feifei Li. Hundreds of thousands, perhaps millions, of people have taken these MOOCs. But with my course, you get to see me jam.

My course is like a concert. You already bought the ticket. I perform. You get out of the class what you put into it. It is your choice whether you come to class (prepared) and do the coursework. If you do, you get an opportunity to learn. If you don't, you wasted your money (or someone else's money).

If you bought a ticket to a Taylor Swift concert, but need to miss the concert, you wouldn't run up on stage and say "Ms. Swift, I'm sorry but I need to miss your concert. But I really like your new songs. So can I come by your home and have you give me a private rendition of the songs I missed at your concert?" So please don't ask me for permission to miss class. And please don't ask me to repeat material from a missed class.

Course Schedule

The homeworks will be handed out as follows:

Homework 1	neural networks (outcome i)	in class Friday 1 September 2023
Homework 2	object detection (outcome ii)	in class Friday 29 September 2023
Homework 3	activity classification (outcome iii)	in class Friday 27 October 2023

The due dates for the three homework assignments are as follows:

Homework 1	neural networks (outcome i)	5pm ET Friday 29 September 2023
Homework 2	object detection (outcome ii)	5pm ET Friday 27 October 2023
Homework 3	activity classification (outcome iii)	5pm ET Friday 24 November 2023

There will be no extensions.

The term project (outcome iv) is organized as follows. Student should self organize into teams of two to four enrolled students. Each team should send email to ece49595cv-staff-list@ecn.purdue.edu by 5pm Friday 1 September 2023 with the names and Purdue Career account email addresses of the team members. Each team should then collectively send a one-page pdf project description proposal to ece49595cv-staff-list@ecn.purdue.edu by 5pm Friday 15 September 2023. I will meet with each team during office hours prior to Friday 29 September 2023 to discuss the project description proposal and offer suggestions and guidance. We will have team project presentations with demos in class during the last two or three weeks of class. These will be scheduled later in the semester. You do not need to submit anything for your team term project; you just need to present your work to class in class.

Academic Integrity

For the course, the only requirements are the three homeworks and the term project/presentation. There is a simple requirement: you cannot plagiarize. You must follow standard academic practice and cite the source of all material that you (re)use. If you violate this anti-plagiarism policy you will be prosecuted by the standard University mechanisms for such violations.

Nondiscrimination Statement

Purdue University is committed to maintaining a community which recognizes and values the inherent worth and dignity of every person; fosters tolerance, sensitivity, understanding, and mutual respect among its members; and encourages each individual to strive to reach his or her potential. In pursuit of its goal of academic excellence, the University seeks to develop and nurture diversity. The University believes that diversity among its many members strengthens the institution, stimulates creativity, promotes the exchange of ideas, and enriches campus life. A hyperlink to Purdue's full Nondiscrimination Policy Statement is included in our course Brightspace under University Policies.

Accessibility

Purdue University is committed to making learning experiences accessible. If you anticipate or experience physical or academic barriers based on disability, you are welcome to let me know so that we can discuss options. You are also encouraged to contact the Disability Resource Center at: drc@purdue.edu or by phone: 765-494-1247.

Mental Health/Wellness Statement

If you find yourself beginning to feel some stress, anxiety and/or feeling slightly overwhelmed, try WellTrack. Sign in and find information and tools at your fingertips, available to you at any time.

If you need support and information about options and resources, please contact or see the Office of the Dean of Students. Call 765-494-1747. Hours of operation are M-F, 8 am-5 pm.

If you find yourself struggling to find a healthy balance between academics, social life, stress, etc. sign up for free one-on-one virtual or in-person sessions with a Purdue Wellness Coach at RecWell. Student coaches can help you navigate through barriers and challenges toward your goals throughout the semester. Sign up is completely free and can be done on BoilerConnect. If you have any questions, please contact Purdue Wellness at evans240@purdue.edu.

If you're struggling and need mental health services: Purdue University is committed to advancing the mental health and well-being of its students. If you or someone you know is feeling overwhelmed, depressed, and/or in need of

mental health support, services are available. For help, such individuals should contact Counseling and Psychological Services (CAPS) at 765-494-6995 during and after hours, on weekends and holidays, or by going to the CAPS office on the second floor of the Purdue University Student Health Center (PUSH) during business hours.

Basic Needs Security

Any student who faces challenges securing their food or housing and believes this may affect their performance in the course is urged to contact the Dean of Students for support. There is no appointment needed and Student Support Services is available to serve students 8 a.m.–5 p.m. Monday through Friday. Considering the significant disruptions caused by the current global crisis as it related to COVID-19, students may submit requests for emergency assistance from the Critical Needs Fund

Emergency Preparedness

In the event of a major campus emergency, course requirements, deadlines and grading percentages are subject to changes that may be necessitated by a revised semester calendar or other circumstances beyond the instructor's control. Relevant changes to this course will be posted onto the course website or can be obtained by contacting the instructors or TAs via email or phone. You are expected to read your @purdue.edu email on a frequent basis.

Instructor Email Policies

Please do not send any course-related email to my personal email account. Please send all course-related email to one of the following two mailing lists:

- ece49595cv-staff-list@ecn.purdue.edu
- ece49595cv-students-list@ecn.purdue.edu

The former just goes to me, and any TAs if we have any in the future. The latter goes to me, any TAs if we have any in the future, and to all students enrolled. All students are automatically added to the latter under their Purdue Career account. Please send all course-related email from your Purdue Career account.

Missed or Late Work

The only assignments this semester are the three homeworks and the term project/presentation. These all have strict deadlines as specified above. There will be no exceptions. No late work will be accepted and no credit will be given for missed or late assignments.

Incompletes

No incompletes will be given.

References

- A. G. Baydin, B. A. Pearlmutter, A. A. Radul, and J. M. Siskind. Automatic differentiation in machine learning: a survey. *Journal of Machine Learning Research (JMLR)*, 18(153):1–43, 2018.
- J. Carreira and A. Zisserman. Quo vadis, action recognition? a new model and the Kinetics dataset. In *Computer Vision and Pattern Recognition (CVPR)*, pages 6299–6308, 2017.
- N. Dalal and B. Triggs. Histograms of oriented gradients for human detection. In *Computer Vision and Pattern Recognition (CVPR)*, pages 886–893, 2005.
- D. F. Dementhon and L. S. Davis. Model-based object pose in 25 lines of code. *International Journal of Computer Vision (IJCV)*, 15(1–2):123–141, 1995.
- A. Diba, V. Sharma, and L. Van Gool. Deep temporal linear encoding networks. In *Computer Vision and Pattern Recognition (CVPR)*, pages 2329–2338, 2017.
- C. Feichtenhofer. X3D: Expanding architectures for efficient video recognition. In *Computer Vision and Pattern Recognition (CVPR)*, pages 203–213, 2020.

- C. Feichtenhofer, A. Pinz, and A. Zisserman. Convolutional two-stream network fusion for video action recognition. In *Computer Vision and Pattern Recognition (CVPR)*, pages 1933–1941, 2016.
- C. Feichtenhofer, H. Fan, J. Malik, and K. He. Slowfast networks for video recognition. In *International Conference on Computer Vision (ICCV)*, pages 6202–6211, 2019.
- P. F. Felzenszwalb, R. B. Girshick, D. McAllester, and D. Ramanan. Object detection with discriminatively trained part-based models. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 32(9):1627–1645, 2010.
- R. Girshick. Fast R-CNN. In *Computer Vision and Pattern Recognition (CVPR)*, pages 1440–1448, 2015.
- R. Girshick, J. Donahue, T. Darrell, and J. Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *Computer Vision and Pattern Recognition (CVPR)*, pages 580–587, 2014.
- K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016.
- K. He, G. Gkioxari, P. Dollár, and R. Girshick. Mask R-CNN. In *International Conference on Computer Vision (ICCV)*, pages 2961–2969, 2017.
- G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger. Densely connected convolutional networks. In *Computer Vision and Pattern Recognition (CVPR)*, pages 4700–4708, 2017.
- S. Ioffe and C. Szegedy. Batch normalization: accelerating deep network training by reducing internal covariate shift. In *International Conference on Machine Learning (ICML)*, pages 448–456, 2015.
- A. Karpathy, G. Toderici, S. Shetty, T. Leung, R. Sukthankar, and L. Fei-Fei. Large-scale video classification with convolutional neural networks. In *Computer Vision and Pattern Recognition (CVPR)*, pages 1725–1732, 2014.
- A. Krizhevsky, I. Sutskever, and G. E. Hinton. ImageNet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 1097–1105, 2012.
- Y. Li, B. Ji, X. Shi, J. Zhang, B. Kang, and L. Wang. TEA: Temporal excitation and aggregation for action recognition. In *Computer Vision and Pattern Recognition (CVPR)*, pages 909–918, 2020.
- J. Lin, C. Gan, and S. Han. TSM: Temporal shift module for efficient video understanding. In *International Conference on Computer Vision (ICCV)*, pages 7083–7093, 2019.
- T.-Y. Lin, P. Dollár, R. Girshick, K. He, B. Hariharan, and S. Belongie. Feature pyramid networks for object detection. In *Computer Vision and Pattern Recognition (CVPR)*, pages 2117–2125, 2017.
- W. Liu, D. Anguelov, D. Erhan, C. Szegedy, S. Reed, C.-Y. Fu, and A. C. Berg. SSD: Single shot multibox detector. In *European Conference on Computer Vision (ECCV)*, pages 21–37, 2016.
- D. G. Lowe. Three-dimensional object recognition from single two-dimensional images. *Artificial Intelligence (AIJ)*, 31(3):355–395, 1987.
- B. A. Pearlmutter and J. M. Siskind. Reverse-mode AD in a functional framework: Lambda the ultimate backpropagator. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 30(2):1–36, 2008.
- J. Redmon and A. Farhadi. YOLO9000: better, faster, stronger. In *Computer Vision and Pattern Recognition (CVPR)*, pages 7263–7271, 2017.
- J. Redmon, S. Divvala, R. Girshick, and A. Farhadi. You only look once: Unified, real-time object detection. In *Computer Vision and Pattern Recognition (CVPR)*, pages 779–788, 2016.
- S. Ren, K. He, R. Girshick, and J. Sun. Faster R-CNN: Towards real-time object detection with region proposal networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 91–99, 2015.

- K. Simonyan and A. Zisserman. Two-stream convolutional networks for action recognition in videos. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 568–576, 2014.
- K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. In *International Conference on Learning Representations (ICLR)*, 2015.
- J. M. Siskind and B. A. Pearlmutter. Nesting forward-mode AD in a functional framework. *Higher-Order and Symbolic Computation (HOSC)*, 21(4):361–376, 2008.
- C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going deeper with convolutions. In *Computer Vision and Pattern Recognition (CVPR)*, 2015.
- C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna. Rethinking the Inception architecture for computer vision. In *Computer Vision and Pattern Recognition (CVPR)*, pages 2818–2826, 2016.
- C. Szegedy, S. Ioffe, V. Vanhoucke, and A. A. Alemi. Inception-v4, Inception-ResNet and the impact of residual connections on learning. In *Conference on Artificial Intelligence (AAAI)*, pages 4278–4284, 2017.
- D. Tran, L. Bourdev, R. Fergus, L. Torresani, and M. Paluri. Learning spatiotemporal features with 3D convolutional networks. In *International Conference on Computer Vision (ICCV)*, pages 4489–4497, 2015.
- L. Wang, Y. Qiao, and X. Tang. Action recognition with trajectory-pooled deep-convolutional descriptors. In *Computer Vision and Pattern Recognition (CVPR)*, pages 4305–4314, 2015.
- L. Wang, Y. Xiong, Z. Wang, Y. Qiao, D. Lin, X. Tang, and L. Van Gool. Temporal segment networks: Towards good practices for deep action recognition. In *European Conference on Computer Vision (ECCV)*, pages 20–36, 2016.
- X. Wang, R. Girshick, A. Gupta, and K. He. Non-local neural networks. In *Computer Vision and Pattern Recognition (CVPR)*, pages 7794–7803, 2018.
- C.-Y. Wu, M. Zaheer, H. Hu, R. Manmatha, A. J. Smola, and P. Krähenbühl. Compressed video action recognition. In *Computer Vision and Pattern Recognition (CVPR)*, pages 6026–6035, 2018.
- J. Yue-Hei Ng, M. Hausknecht, S. Vijayanarasimhan, O. Vinyals, R. Monga, and G. Toderici. Beyond short snippets: Deep networks for video classification. In *Computer Vision and Pattern Recognition (CVPR)*, pages 4694–4702, 2015.
- B. Zhou, A. Andonian, A. Oliva, and A. Torralba. Temporal relational reasoning in videos. In *European Conference on Computer Vision (ECCV)*, pages 803–818, 2018.
- Y. Zhu, Z. Lan, S. Newsam, and A. Hauptmann. Hidden two-stream convolutional networks for action recognition. In *Asian Conference on Computer Vision (ACCV)*, pages 363–378, 2018.
- M. Zolfaghari, K. Singh, and T. Brox. ECO: Efficient convolutional network for online video understanding. In *European Conference on Computer Vision (ECCV)*, pages 695–712, 2018.

