

CS 1674/2074: Intro to Computer Vision

PhD. Nils Murrugarra-Llerena
nem177@pitt.edu



Who am I?



B.S. Computer Science at
**National University of
Trujillo**



M.S. in Computer Science at
**University of São Paulo in
AI**

Who am I?



PhD in Computer Science at
University of Pittsburgh in
Computer Vision



**Research scientist at
Snap Inc.**

Who am I?



**Assistant professor at
Weber State University**




**Teaching Assistant
Professor at University
of Pittsburgh**

[Students' presentations]

Name, hobbies, and mention one thing that you expect to learn in this course 😊

Course intro: Syllabus

- Contact Information
 - Prof. Nils Murrugarra
 - nem177@pitt.edu
 - Please, add prefix “[CS 1674]” in all emails.
 - Website: https://nineil.github.io/courses/fall24_cs1674/
- Lectures:
 - Tue/Thu 11:00 am – 12:15 pm [SENSQ 5313]
- Office hours:
 - Tue/Wed 9:00 am – 11:00 am [SENSQ 5419]



Course intro: Textbook


Computer Vision algorithms and applications

Edition: 2nd

By Richard Szeliski


ISBN: 978-3030343712

Year: 2022



Course intro: What to expect?

- Material is based on previous iterations of [Computer Vision course](#) from Prof. Adriana Kovashka.
- Exams mainly cover this material
- We will do around 7 to 9 programming assignments



Course intro: What to expect?

- There will be a lot of work!
- However, you will learn a lot :). Please, ask questions in class and use my office hours as needed.
- I would like to help you much as possible.



Course intro: What to expect? [Warning #1]

- I've opted for shorter, more manageable HW assignments, but there is a lot of them
- I expect you'd be spending **4-6 hours** on each assignment
- ... But you get to understand algorithms and concepts in detail!




Course intro: What to expect? [Warning #2]

- Some parts will be **hard** and require that you pay close attention!
- **Use instructor's office hours**
- ... You will learn a lot!



Course intro: What to expect?



H/T Kirk Pruhs

Course intro: programming assignments

- We will learn Python programming language
- You get 3 "free" late days counted in minutes, i.e., you can submit a total of 72 hours late. Once you've used up your free late days, you will incur a penalty of 25% from the total assignment credit possible for each late day.



Review Syllabus

Canvas Link:

https://canvas.pitt.edu/courses/288692/files/17497249?module_item_id=5015518

Motivation: Faces and digital cameras



Camera waits for everyone to smile to take a photo [Canon]



Setting camera focus via face detection

Motivation: Face recognition



Motivation: Optical Character Recognition



Motivation: Accurate object detection



The screenshot shows a web browser window titled "YOLO: Real-Time Object Detection" with the URL "pjreddie.com/darknet/yolo/". The main content area features the YOLO logo at the top, followed by the title "YOLO: Real-Time Object Detection". Below this is a video player titled "YOLO Watches Nature Part 2" showing a herd of elephants in a savanna. Blue bounding boxes are drawn around several elephants in the video. To the right of the video is a console window displaying the following output:

```
FPS: 31  
Objects:  
elephant: 0.31  
elephant: 0.33  
elephant: 0.31  
elephant: 0.27  
elephant: 0.41
```

Below the video player, there is a text description: "You only look once (YOLO) is a system for detecting objects on the Pascal VOC 2012 dataset. It can detect the 20 Pascal object classes:" followed by a bulleted list of object classes:

- person
- bird, cat, cow, dog, horse, sheep
- aeroplane, bicycle, boat, bus, car, motorbike, train
- bottle, chair, dining table, potted plant, sofa, tv/monitor

Redmon et al., "You Only Look Once: Unified, Real-Time Object Detection", CVPR 2016

Motivation: Exploring photo collections



Photo Tourism

Exploring photo collections in 3D

Microsoft



(a)



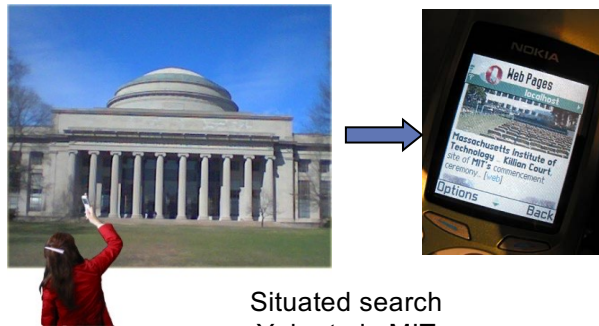
(b)



(c)

Snaveley et al.

Motivation: Linking info with a mobile device



Situated search
Yeh et al., MIT




MSR Lincoln



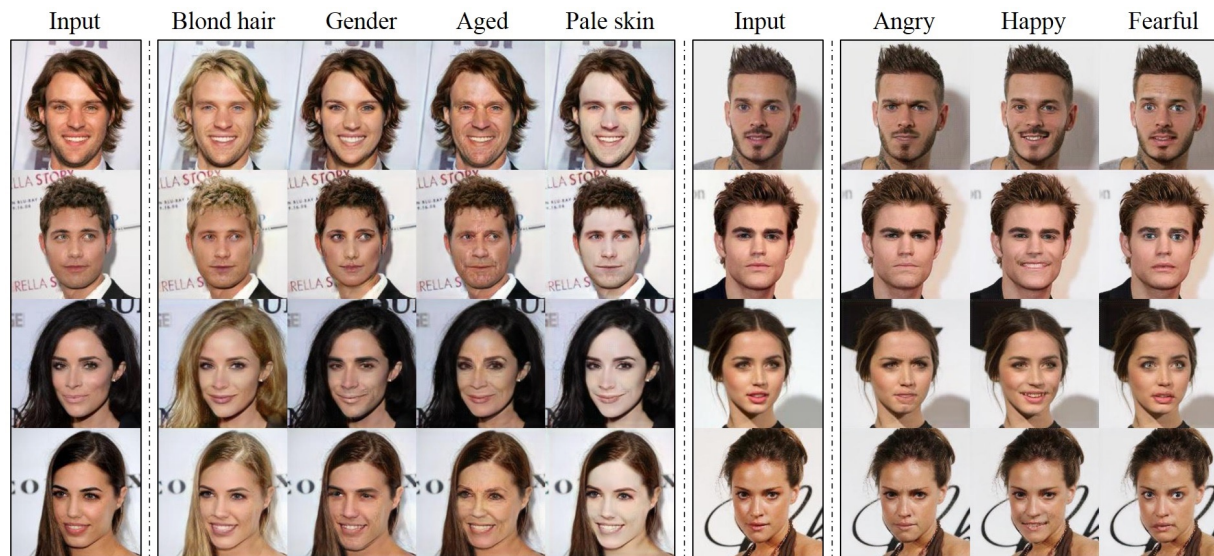
kooba

Motivation: Transferring art styles



DeepArt.io – try it for yourself!

Motivation: Image Generation (faces)



Motivation: Interactive Systems



Shotton et al.



Motivation: Video-based interfaces

[YouTube Link](#)



Human joystick
NewsBreaker Live




Assistive technology systems
Camera Mouse
Boston College

Motivation: Computer Vision for Medicine



Image guided surgery
MIT AI Vision Group




fMRI data
Golland et al.



Motivation: Safety and security




Navigation, driver safety



Monitoring pool (Poseidon)



Pedestrian detection
MERL, Viola et al.



Surveillance

Setup Environment

- Create [Github Account](#)
- Install [Github Desktop](#)
- You may use any IDE for Python
 - I use Pycharm:
<https://www.jetbrains.com/pycharm/>
 - Apply for your educational free license:
<https://www.jetbrains.com/community/education/#students>



Setup Learning Environment

Installation and learning environment:

https://github.com/nineil-pitt/cs1674_2074_fall23