CS 2770: 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 ☺

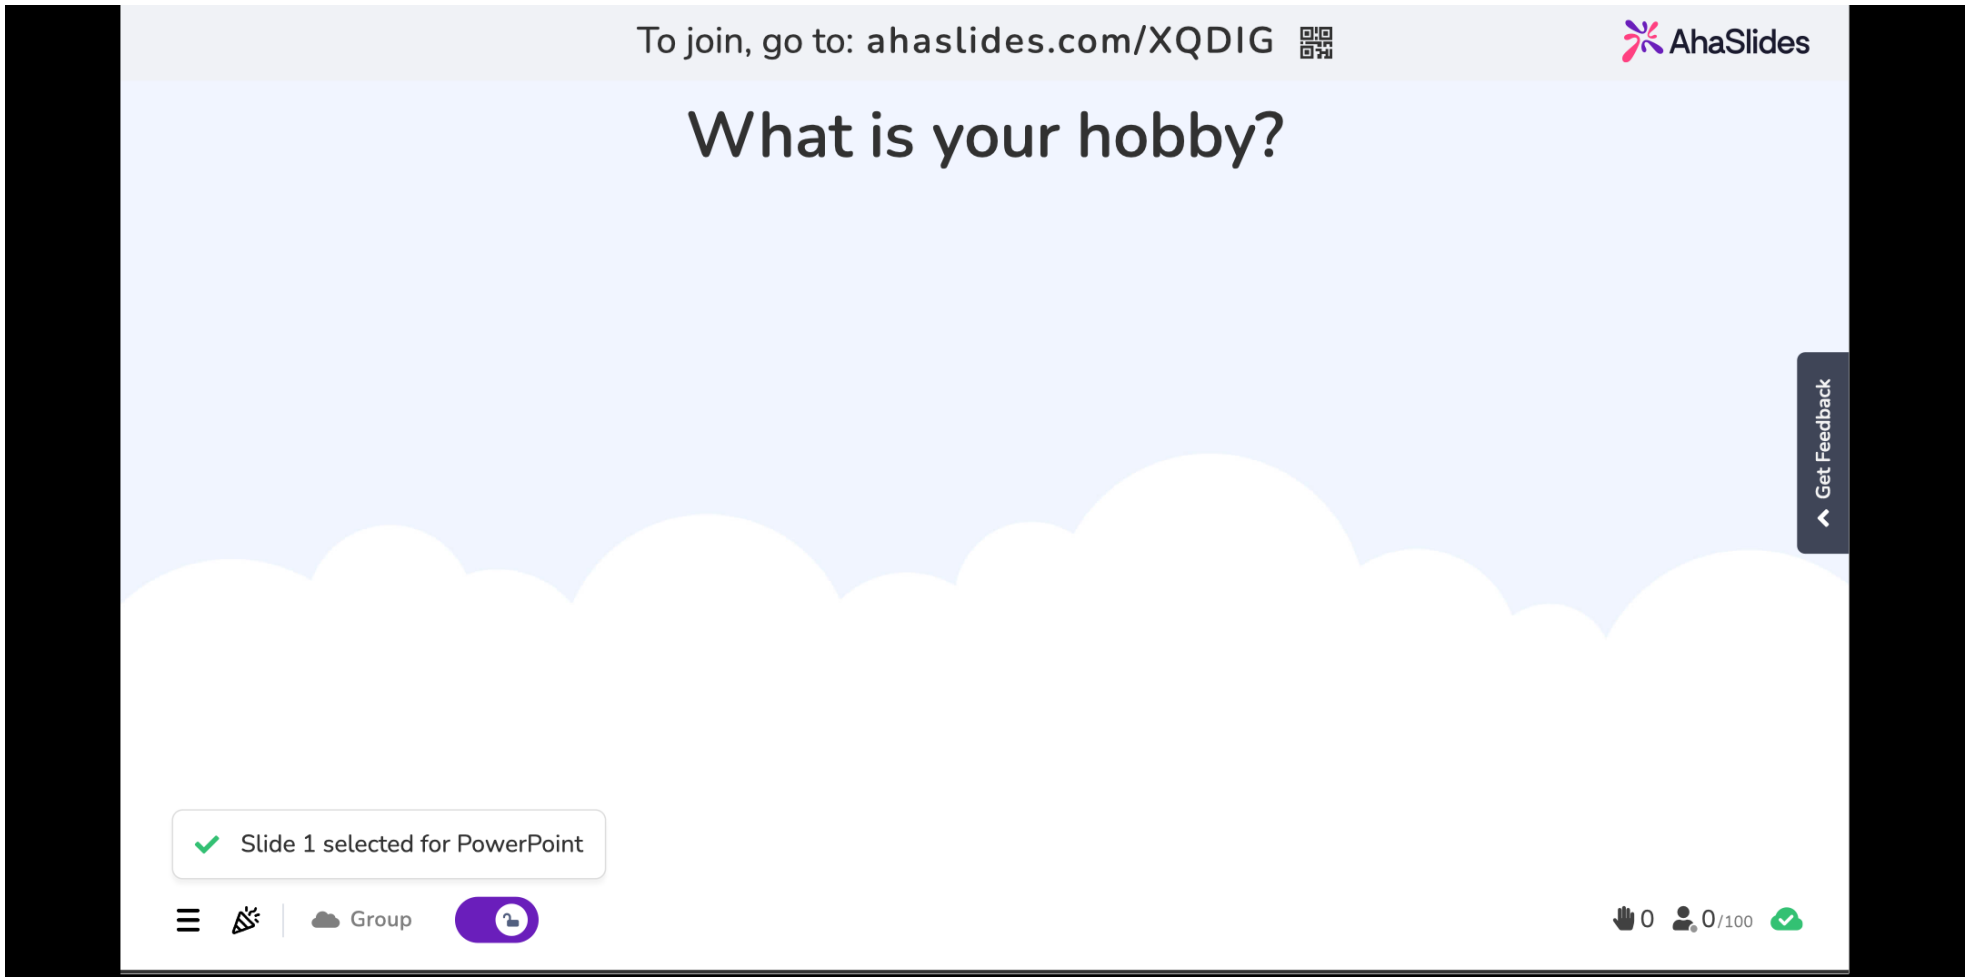
To join, go to: ahaslides.com/XQDIG
AhaSlides
What is your hobby?
Get Feedback
Slide 1 selected for PowerPoint
Group
0
0/100

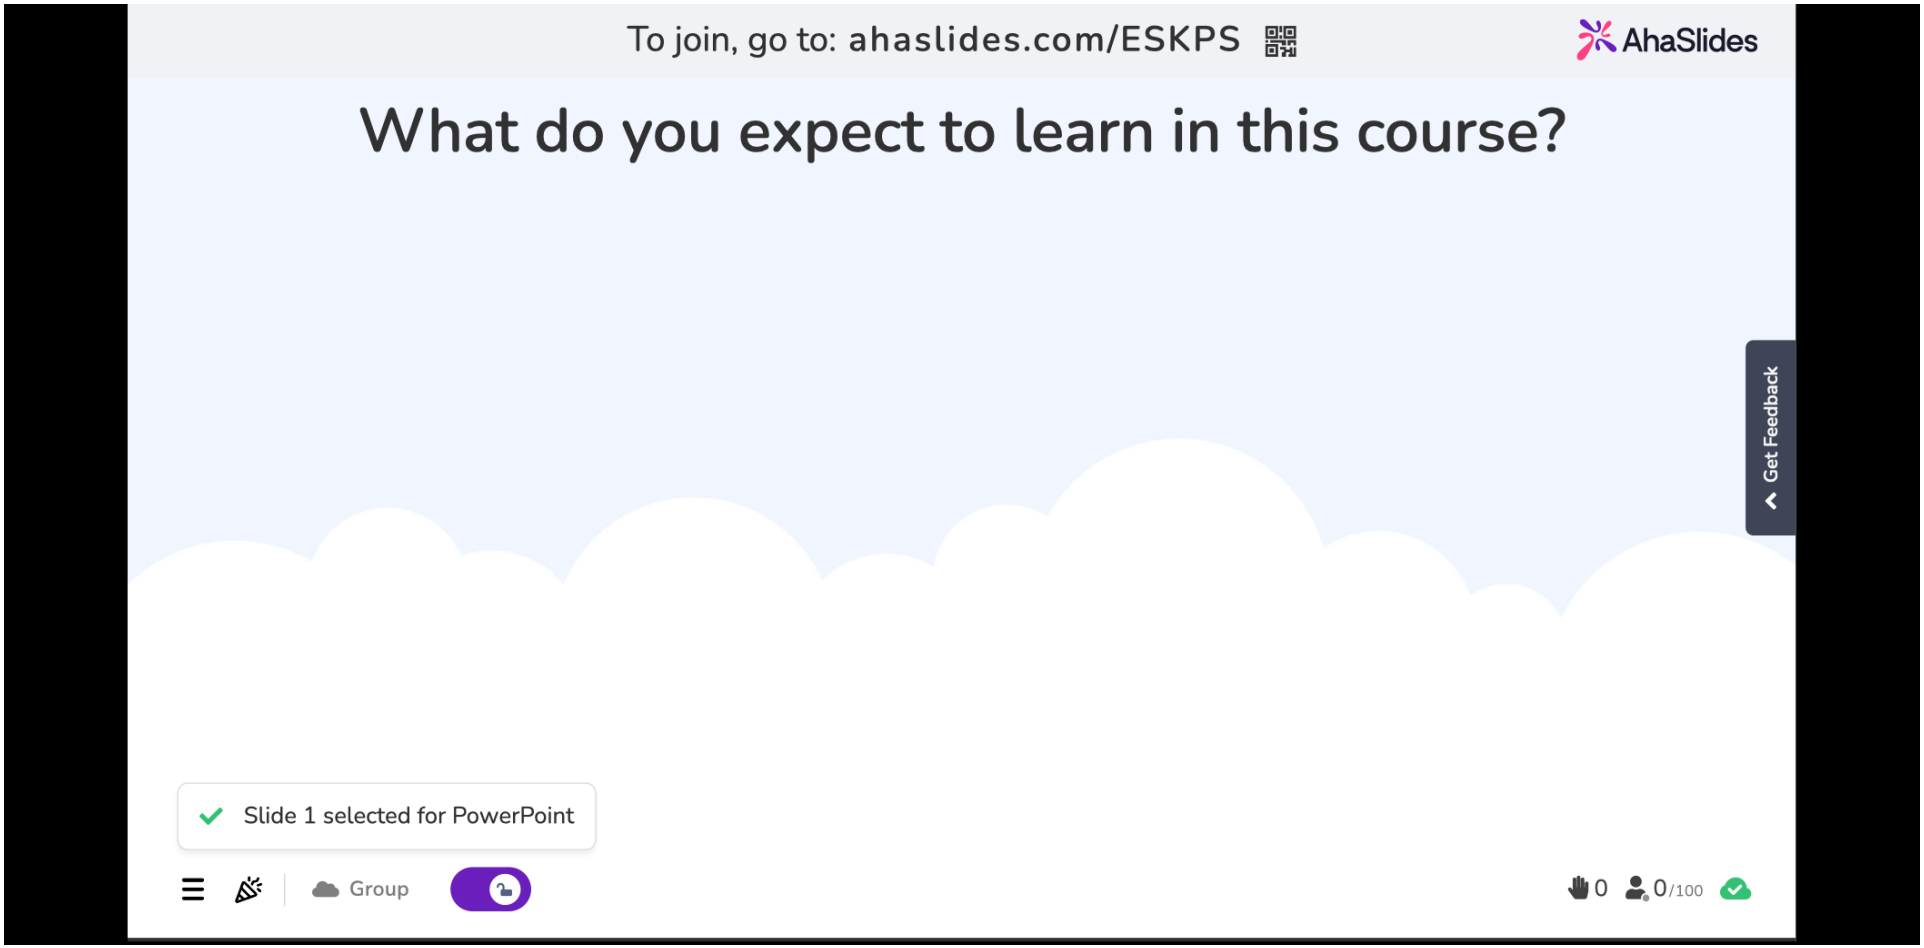
To join, go to: ahaslides.com/ESKPS
AhaSlides
What do you expect to learn in this course?
Get Feedback
Slide 1 selected for PowerPoint
Group
0
0/100

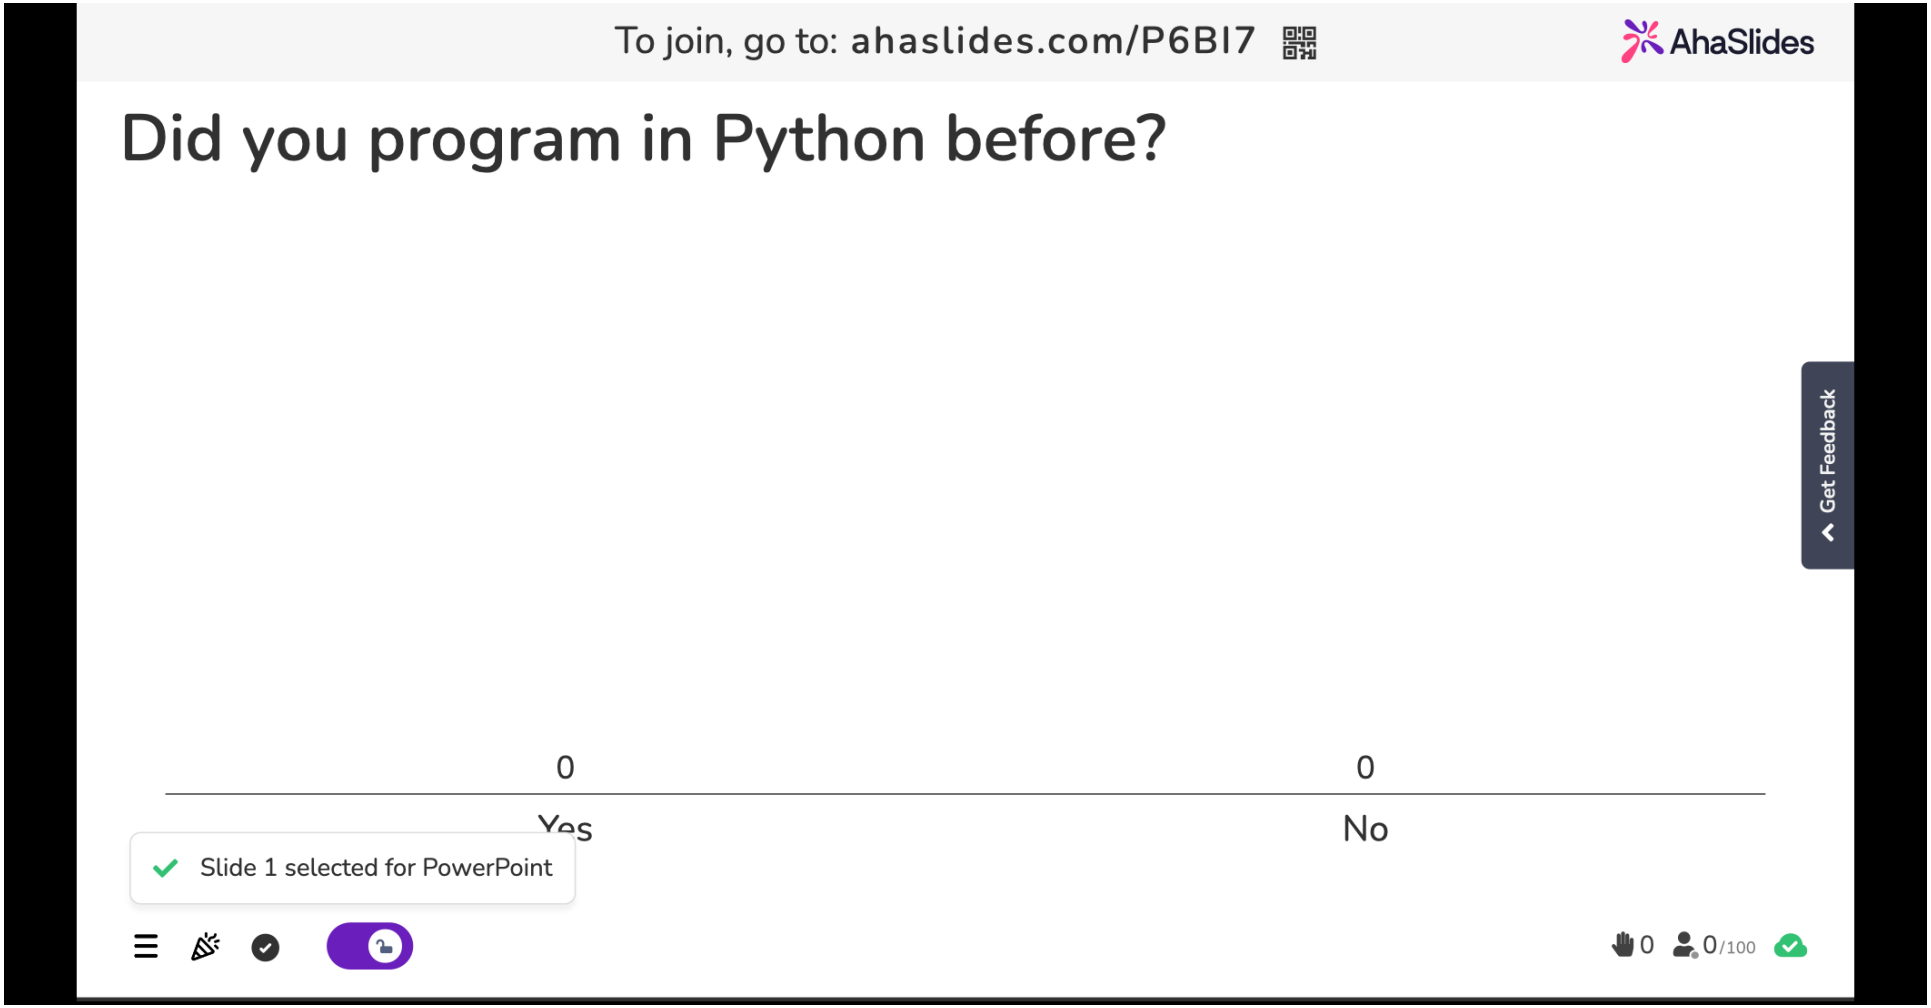
To join, go to: ahaslides.com/P6BI7
AhaSlides
Did you program in Python before?
0
Yes
0
No
Get Feedback
Slide 1 selected for PowerPoint
0
0/100

Course intro: Syllabus

- Contact Information
  - Prof. Nils Murrugarra
  - nem177@pitt.edu
  - Please, add prefix “[CS 2770]” in all emails.
  - Website: https://nineil.github.io/courses/spring26_cs2770/

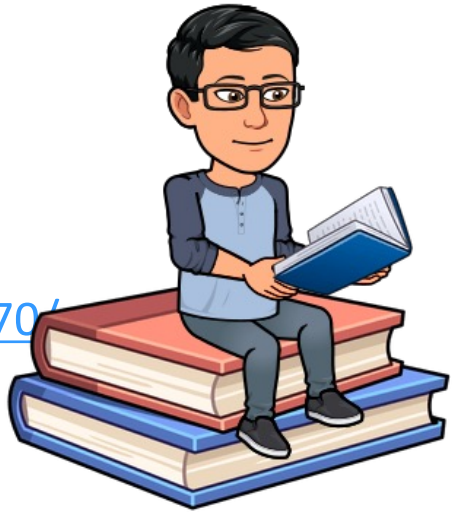

- Lectures:
  - Mon/Wed: 9:30am - 10:45am @ SENSQ 5313

- Office hours:
  - TBD (Please, fill this form). Inputs will be considered with my other courses, and my own schedule

Course intro: Textbook

Computer Vision algorithms and applications

Edition: 2nd

By Richard Szeliski

ISBN: 978-3030343712

Year: 2022

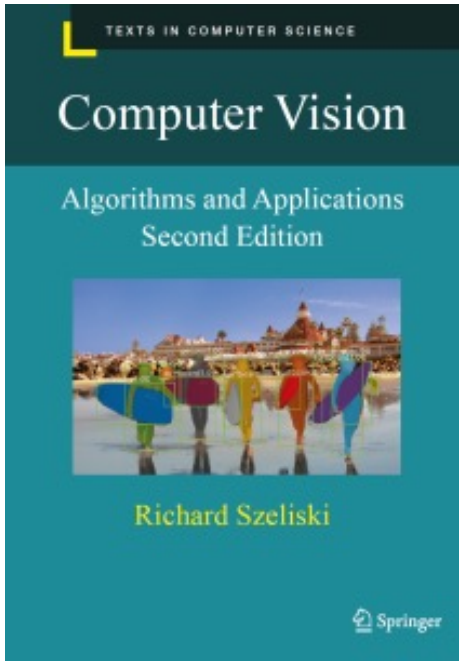

Course intro: Textbook

Foundations of Computer Vision

Edition: 1st

By Antonio Torralba, Phillip Isola, and William Freeman

ISBN: 978-0262048972

Year: 2024

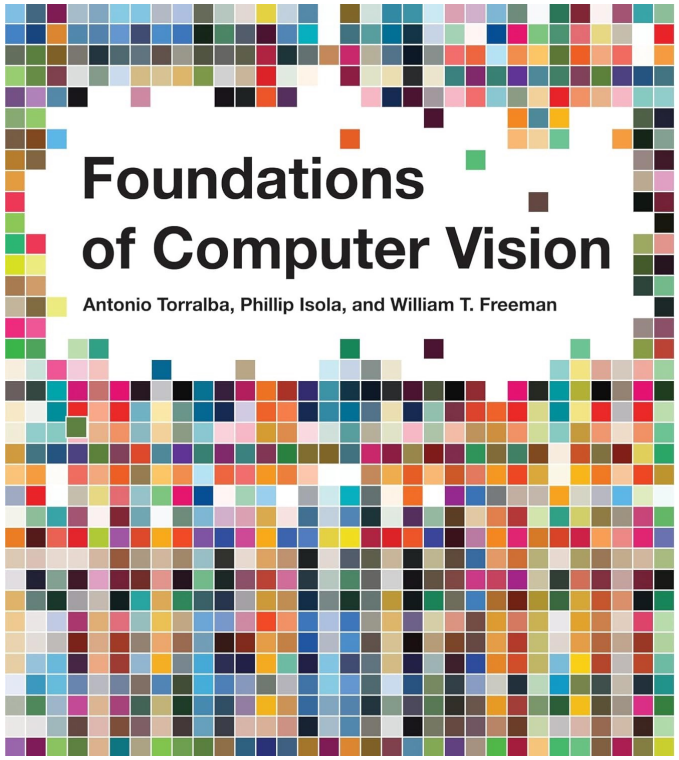

Course intro: What to expect?

- Material is based on previous iterations of my Computer Vision courses and material from US well-recognized Universities
- Exams mainly cover this material
- We will do around 4 to 6 programming assignments

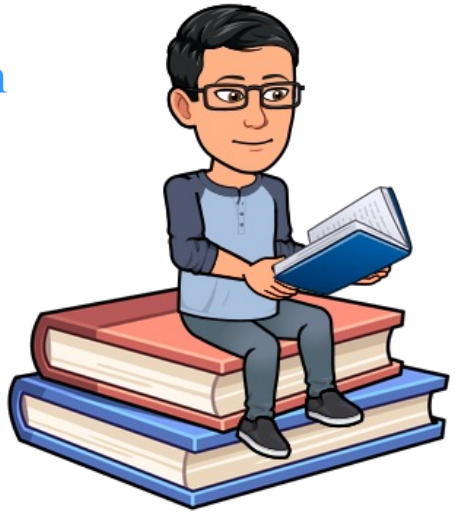

Course intro: What to expect?

- 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.
- Quizzes and projects cannot be made up unless arrangements are made to take/submit them ahead of time.
- Late assignments will be accepted with a 10% penalty per day up to 2 days to provide for unforeseen circumstances.

Review Syllabus

Canvas Link:

https://canvas.pitt.edu/courses/350892/files/231502l8?module_item_id=6256059

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

https://www.mathworks.com/matlabcentral/fileexchange/18169-optical-character-recognition-ocr

Motivation: Accurate object detection

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

Motivation: Exploring photo collections

Snavely et al.

Motivation: Linking info with a mobile device

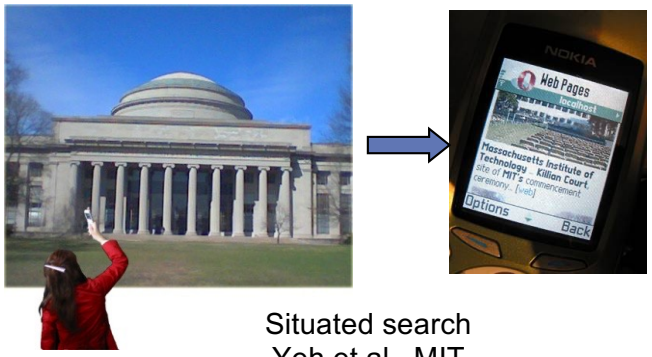

Situated search
Yeh et al., MIT

kooaba

MSR Lincoln

Motivation: Transferring art styles

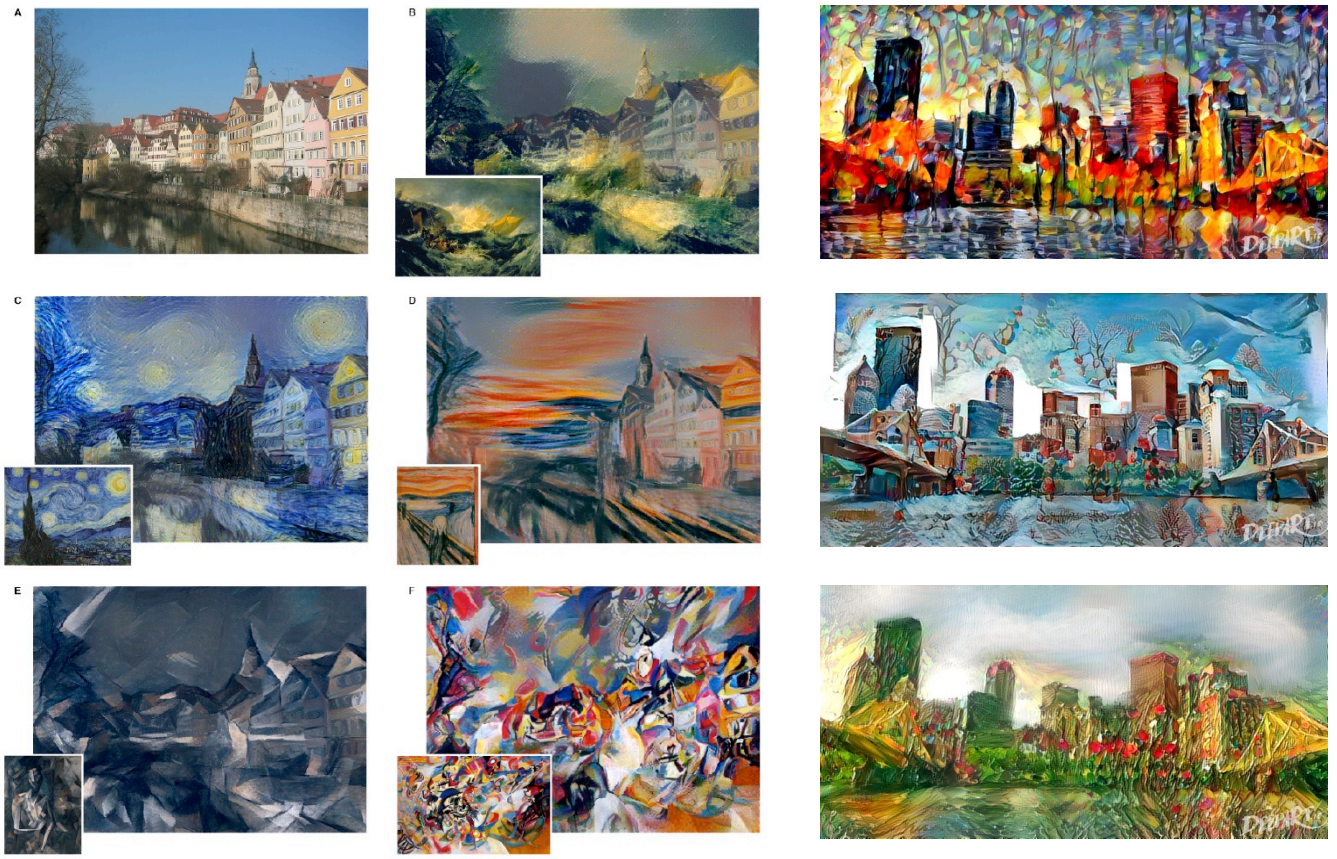

Gatys et al., CVPR 2016

DeepArt.io – try it for yourself!

Motivation: Image Generation (faces)

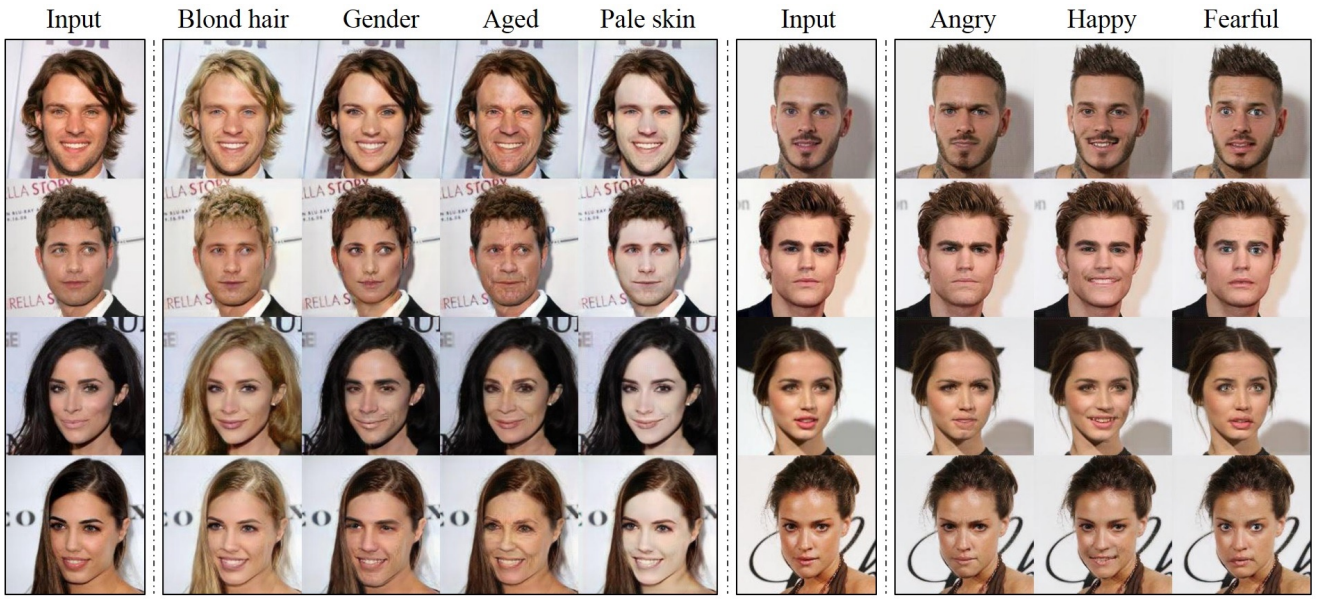

Choi et al., CVPR 2018

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

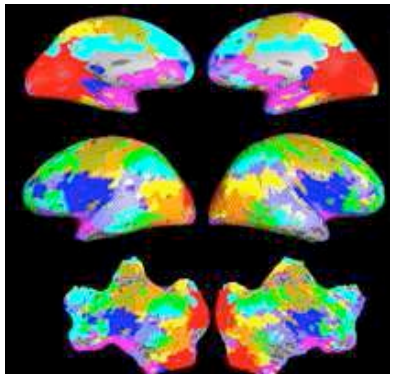

fMRI data
Golland et al.

Motivation: Safety and security

Navigation, driver safety

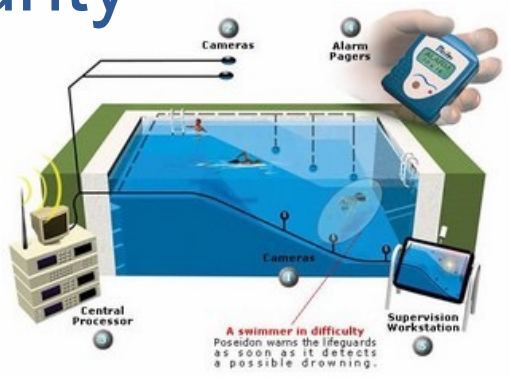

Monitoring pool (Poseidon)

Pedestrian detection
MERL, Viola et al.

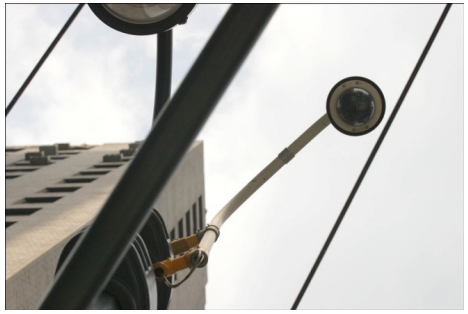

Surveillance

Motivation: Generative AI

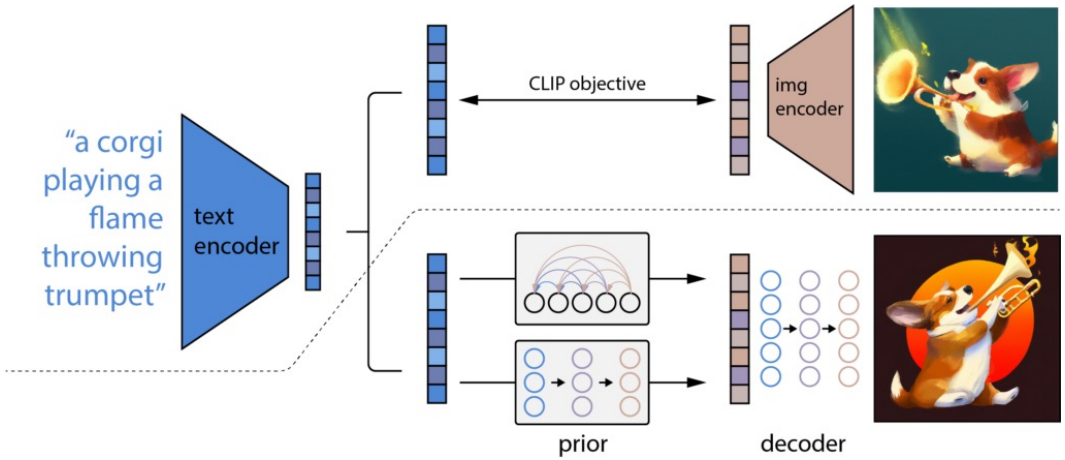

Dall.e 2: https://learnopencv.com/mastering-dall-e-2/

Stable Diffusion: "Triceratops programming on a MacBook in a startup office"

"a man with red hair"

"a vase of flowers"

Text-conditional image-inpainting [ref]

Motivation: Multimodal Large Language Models

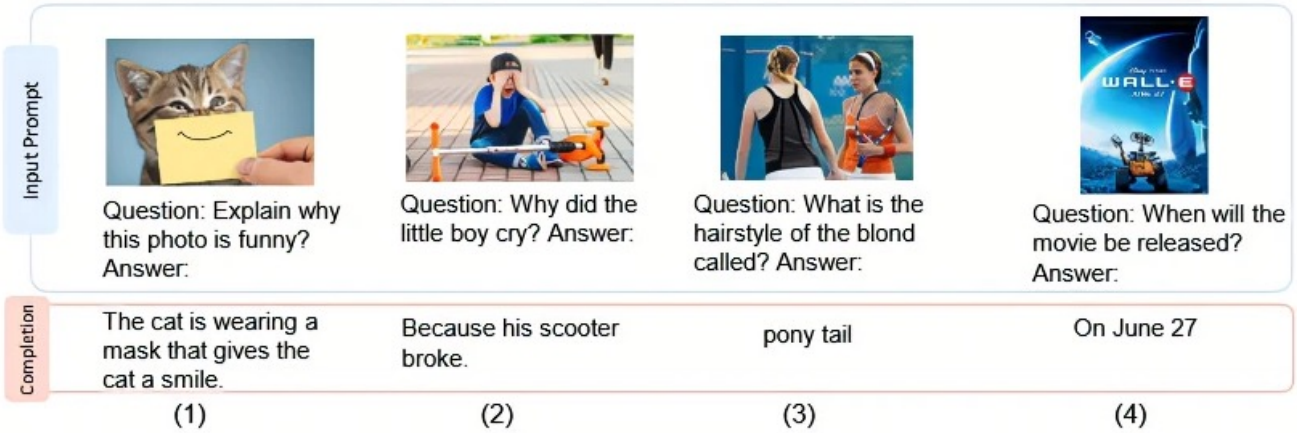

Gpt-4: https://medium.com/@amol-wagh/whats-new-in-gpt-4-an-overview-of-the-gpt-4-architecture-and-capabilities-of-next-generation-ai-900c445d5ffe

Ask an Image Question to https://gemini.google.com/

Motivation: NVIDIA Applications

https://www.youtube.com/watch?v=OnTgbN3uXvw&ab_channel=NVIDIA

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/cs2770_spri26