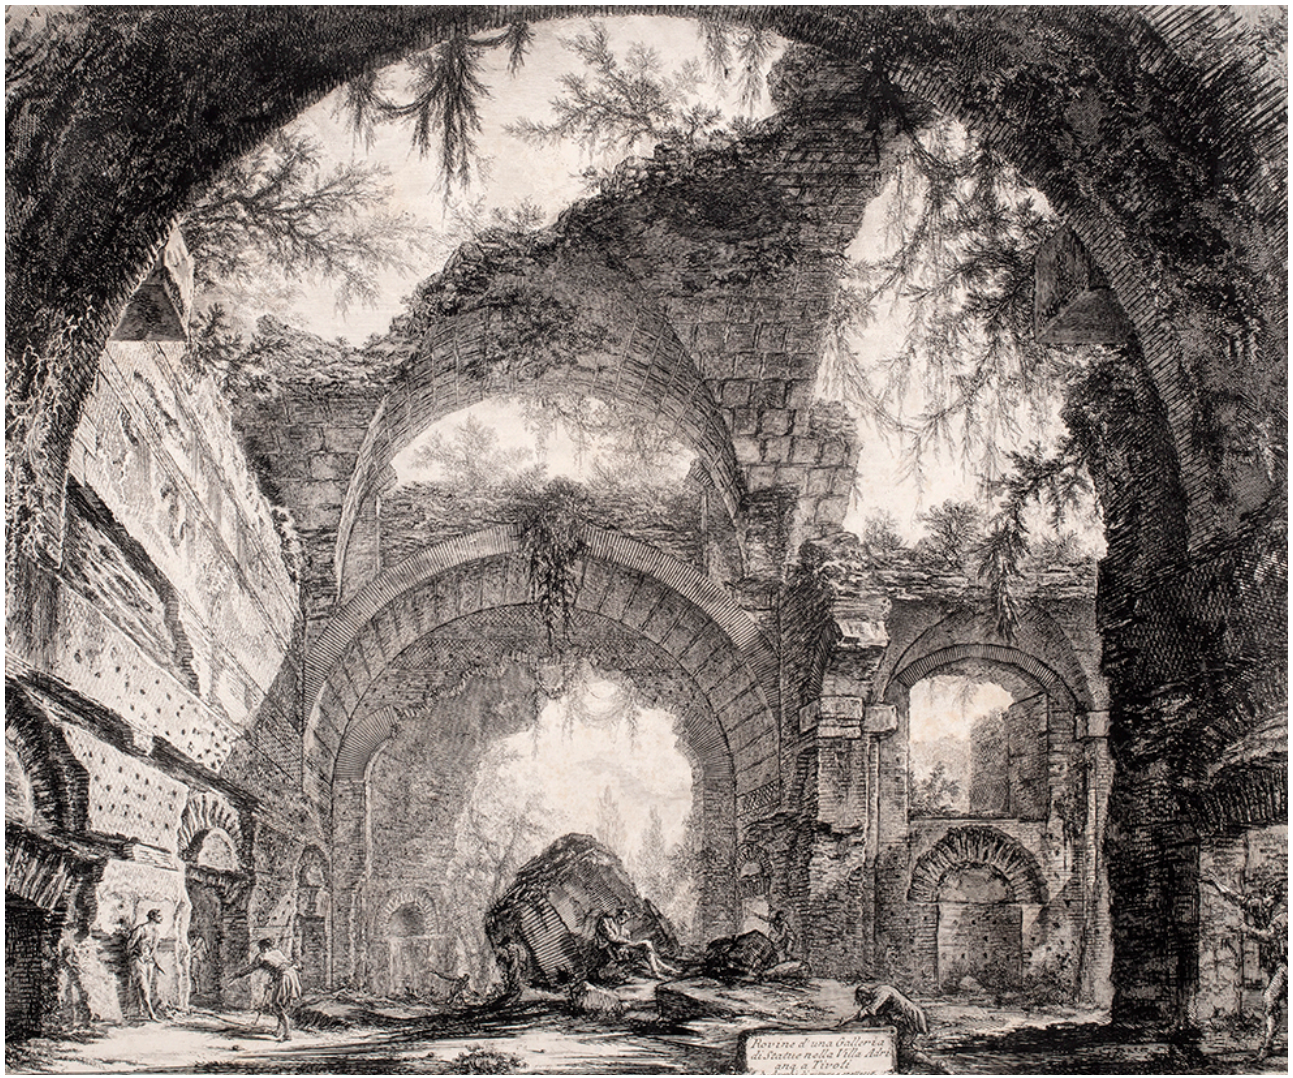

Fall 2023
CS 598 LAZ

Computer vision:
What will stand
the test of time?

Wrap-Up

Final project presentations: Next Wednesday

- 8 minutes
- Structure:
  - Problem definition – with pictures!
  - Method outline
  - Results to date
  - Take-aways

Final project reports: Due Monday, December 11

- See class webpage for instructions

Student presentations: Reflections

Student presentations: Retrospective arrangement

- **From human to artificial intelligence** (5)
- **Technical overviews** (3)
- **Application domains** (2)
- **Academic impacts** (2)
- **Societal impacts** (4)

Paul Klee, Angel of History (Angelus Novus), 1920

Student presentations: Retrospective arrangement

- **From human to artificial intelligence**
  - **Gestalt perception** (Zixuan, Xiyan, Xu)
  - **Memory-inspired neural methods** (Zhen, Savya, Yifei)
  - **Neural visual reasoning** (Volodymyr, Ayush)
  - **Is LLM the answer to AGI?** (Wei-Yu, Mu-Chun)
  - **Can computers create art?** (Pulkit, Utkarsh, Daniel)

Figure from Zhen, Savya, Yifei's presentation
DALL-E 3 Prompt: Collage of moments of the same person—child riding a bike, student studying, elderly couple reminiscing

Student presentations: Retrospective arrangement

- **From human to artificial intelligence**
  - **Gestalt perception** (Zixuan, Xiyan, Xu)
  - **Memory-inspired neural methods** (Zhen, Savya, Yifei)
  - **Neural visual reasoning** (Volodymyr, Ayush)
  - **Is LLM the answer to AGI?** (Wei-Yu, Mu-Chun)
  - **Can computers create art?** (Pulkit, Utkarsh, Daniel)

Which one forms the cube?

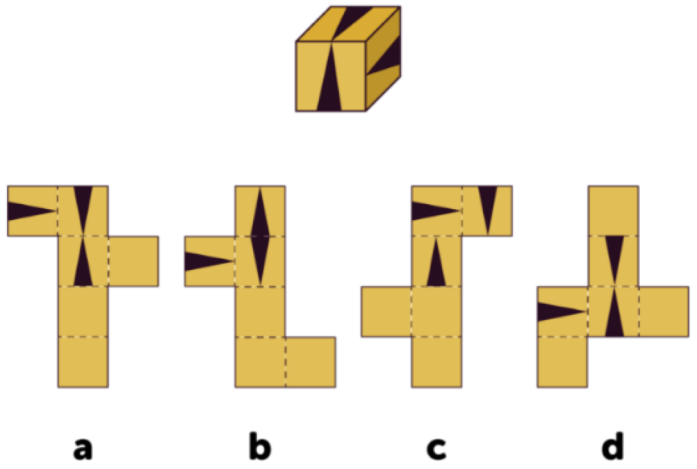

Figure from Volodymyr and Ayush's presentation

Student presentations: Retrospective arrangement

- **Technical overviews**
  - **NeRFs** (Albert, Yilin, Tianhang)
  - **SLAM** (Jose, Shreya, Ahmed)
  - **Neural generative modeling** (Steven, Mingtong, Seemandhar)

Figure from Albert, Yilin, Tianhang's presentation: Photosculpture

Student presentations: Retrospective arrangement

- **Application domains**
  - **CV in healthcare**
    (Josh, Mario, Ansh)
  - **HCI before and after CV**
    (Nancy, Heth, Jenish)

Figure from Nancy, Heth, Jenish's presentation

Student presentations: Retrospective arrangement

- **Academic impacts**
  - **Should peer review be abandoned?** (Aditya, Arjun)
  - **Academic vs. industry AI research** (Sharath, Pradhyumna, Chandan)

Figure from Sharath, Pradhyumna, Chandan's presentation.
DALLE - 3 Prompt: display a split screen. On one side, show a traditional academic setting: perhaps an old university lecture hall, with students pouring over textbooks. On the other side, depict a bustling tech company, with sleek computers, buzzing servers, and teams in agile huddles.

Student presentations: Retrospective arrangement

- **Academic impacts**
  - **Should peer review be abandoned?** (Aditya, Arjun)
  - **Academic vs. industry AI research** (Sharath, Pradhyumna, Chandan)

Figure from Sharath, Pradhyumna, Chandan's presentation.
DALLE-3 Prompt: Image of a strong and vibrant bridge connecting the academic and industrial realms. Images of researchers from both academia and industry collaborating and sharing knowledge.

Student presentations: Retrospective arrangement

- **Societal impacts**
  - **Environmental impact**
    (Genglin, Chandni, Ritik)
  - **Privacy concerns**
    (Soorya, Nidhish, Roubo)
  - **Generative AI: Societal issues and ethics**
    (Asher, Matthew, Sid)
  - **Will AI cause mass unemployment?**
    (Sainath, Sindhu, Sristi)

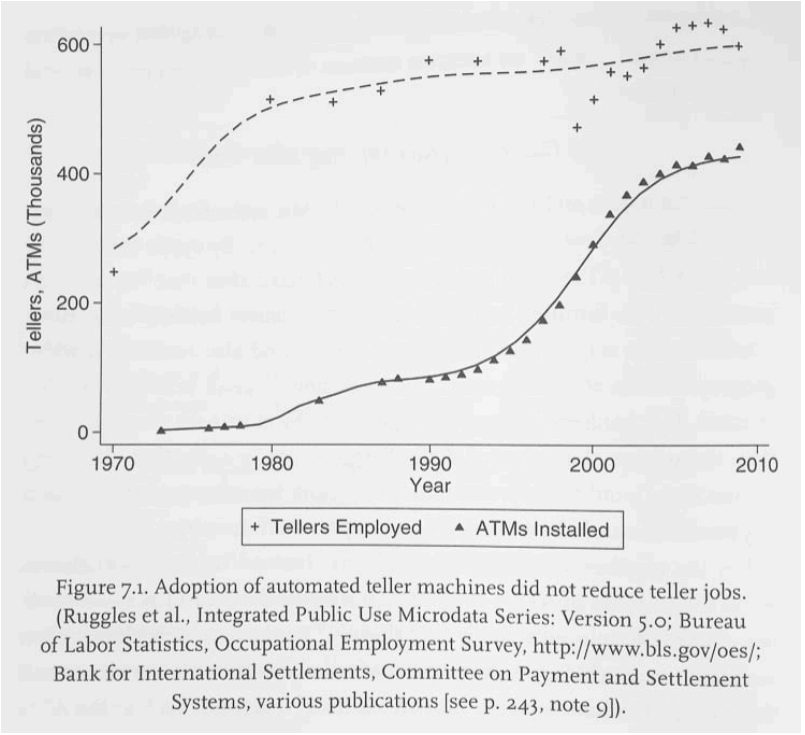

Figure 7.1. Adoption of automated teller machines did not reduce teller jobs. (Ruggles et al., Integrated Public Use Microdata Series: Version 5.0; Bureau of Labor Statistics, Occupational Employment Survey, http://www.bls.gov/oes/; Bank for International Settlements, Committee on Payment and Settlement Systems, various publications [see p. 243, note 9]).

Figure from Sindhu and Sristi's presentation

Quo Vadis, AI/CV?

CVPR 2023 workshop

https://sites.google.com/view/academic-cv/

ICCV 2023 workshop

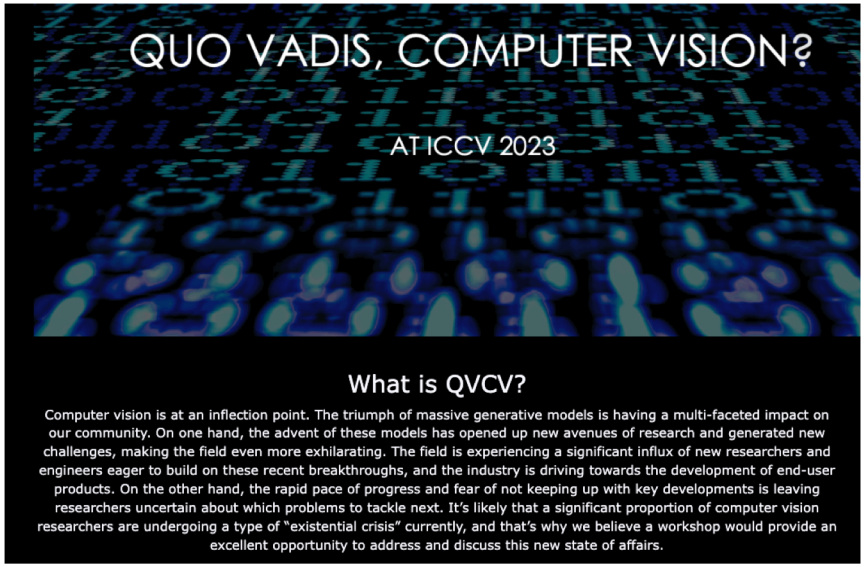

https://gkioxari.github.io/Tutorials/iccv2023/

Parting thoughts

- R. Sutton, The bitter lesson
- H. Simon, Why should machines learn?
- J. Barron, How technology works

The future?

THE NEW YORKER

PROFILES

WHY THE GODFATHER OF A.I. FEARS WHAT HE'S BUILT

Geoffrey Hinton has spent a lifetime teaching computers to learn. Now he worries that artificial brains are better than ours.

By Joshua Rothman

November 13, 2023

https://www.newyorker.com/magazine/2023/11/20/geoffrey-hinton-profile-ai

The future?

THE NEW YORKER

PROFILES

WHY THE GODFATHER OF A.I. FEARS WHAT HE'S BUILT

Geoffrey Hinton has spent a lifetime teaching computers to learn. Now he worries that artificial brains are better than ours.

By Joshua Rothman

November 13, 2023

https://www.newyorker.com/magazine/2023/11/20/geoffrey-hinton-profile-ai

During his last few years at Google, Hinton focussed his efforts on creating more traditionally mindlike artificial intelligence using hardware that more closely emulated the brain. In today's A.I.s, the weights of the connections among the artificial neurons are stored numerically; it's as though the brain keeps records about itself. In your actual, analog brain, however, the weights are built into the physical connections between neurons. Hinton worked to create an artificial version of this system using specialized computer chips.

"If you could do it, it would be amazing," he told me. The chips would be able to learn by varying their "conductances." Because the weights would be integrated into the hardware, it would be impossible to copy them from one machine to another; each artificial intelligence would have to learn on its own. "They would have to go to school," he said. "But you would go from using a megawatt to thirty watts." As he spoke, he leaned forward, his eyes boring into mine; I got a glimpse of Hinton the evangelist. Because the knowledge gained by each A.I. would be lost when it was disassembled, he called the approach "mortal computing." "We'd give up on immortality," he said. "In literature, you give up being a god for the woman you love, right? In this case, we'd get something far more important, which is energy efficiency." Among other things, energy efficiency encourages individuality: because a human brain can run on oatmeal, the world can support billions of brains, all different. And each brain can learn continuously, rather than being trained once, then pushed out into the world.