Facial Recognition

Agenda

- Introduction – 5 minutes
- How do Humans Recognize Faces – 15 minutes exercise
- How does AI Recognize Faces –15 min lecture, 15 min exercise
- How is Facial Recognition Technology applied? – 15 min lecture, 15 min exercise

Task Definition

Facial Recognition is a form of biometric identification, often split into 2 types

Face Matching: Are these two faces the same person?

Face Identification: Who is this?

Task Corollaries

Other questions we can ask given a picture of someone's face:

Sentiment Analysis – What emotion is this person feeling (happy, sad, angry…)?

Classification – Is this person young or old? Boy or girl?

Face Detection – Is there a person/face in this photo?

Worksheet: How do Humans Recognize Faces (15 minutes)

Split into groups and complete Worksheet FRT Section 1 and 2.

Remember to assign one or more people to be dedicated note takers and time keepers, to make sure that your discussion is recorded and you stay on track!

Worksheet: How do Humans Recognize Faces (15 minutes)

1. If you see a friend on the street, at school, at a party, How do you recognize that it’s them?
2. What do you normally do when you recognize a friend out in public?
3. What do you do when you think you recognize a friend, but you aren’t sure?
4. Has it ever happened where you said hi to, or otherwise interacted with, a friend but it turned out to be someone else?
5. Is this person happy or sad? Exercise: https://greatergood.berkeley.edu/quizzes/ei_quiz/

Facial Recognition Technology

Modern Facial Recognition Technology consists of a pipeline which leverages Computer Vision techniques.

- Recall: Computer Vision is the process of detecting certain feature from an image and making classifications based on those features.

Pipeline:

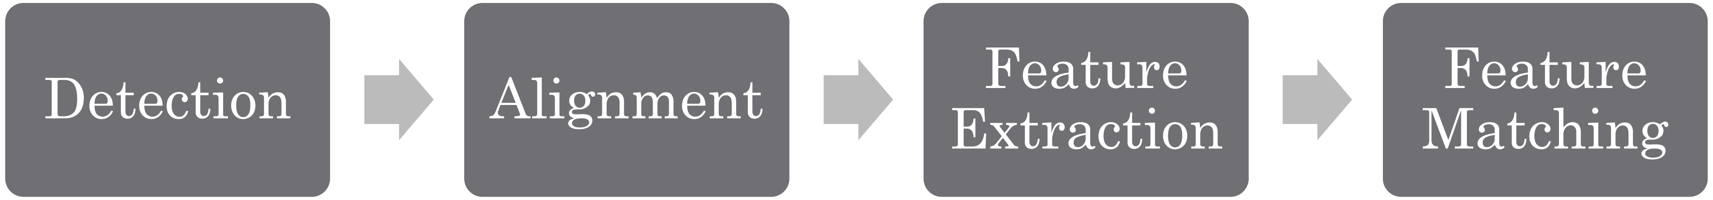

Detection

Alignment

Feature Extraction

Feature Matching

Face Detection

- Is there a face in this picture, and where is it?
- If there are multiple faces, find each individual face and repeat the task on those individuals.

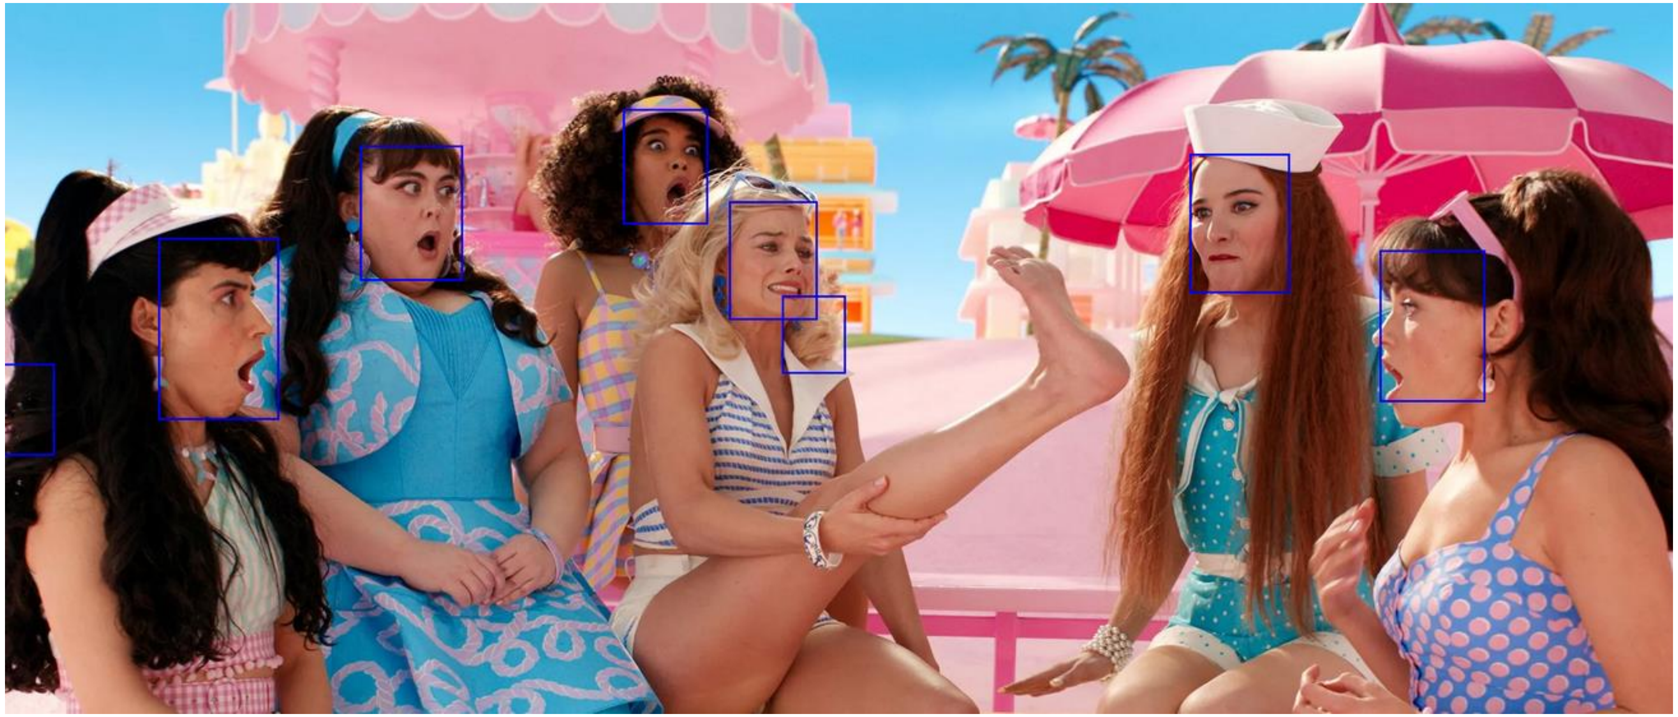

Image Still from Barbie(2023)

Face Alignment

- Extract the desired face from the photo. Align it such that it is centered within the frame and remove undesired features (backgrounds, other faces.)

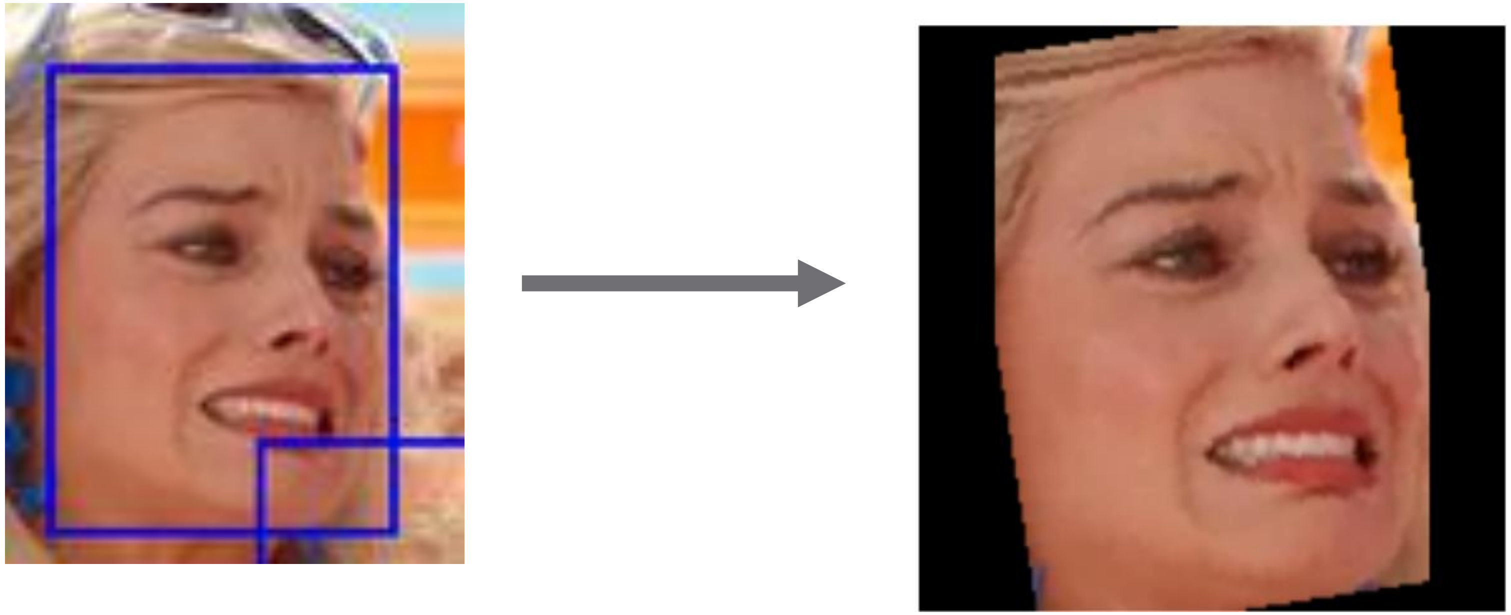

Feature Extraction

- Using a Machine Learning Model, create a mathematical representation of the face.
- Some underlying features that _could_ be part of this representation:
  - Distance between the eyes
  - Contour of the lips
  - Size of cheekbones
  - Length of Nose
  - Much, much more.

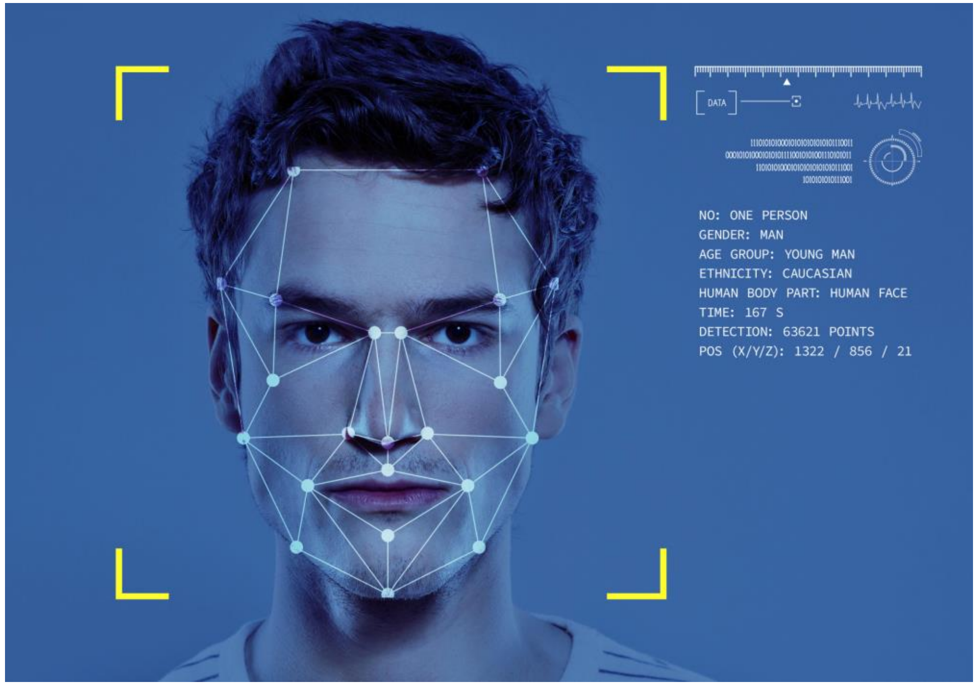

Feature Matching

- Compare the extracted features against a database(identification) or other photo(matching).
- Comparison is performed using a confidence level classification.
- Recall: Classifiers answer a question based on an input. In this case, are these two feature extracted faces the same.
- Confidence level: How likely does the model think the two faces are the same? Or, what is the face which has the highest likelihood of being a match?

Model Training

How would we train both of the models in pipeline?

Possible Solutions:

1. Train the extractor and matcher separately, using a dataset of pictures and actual face measurements for the extraction and a dataset of faces labelled as matching or non matching for the matcher.
2. Train them at the same time, using only a dataset of faces labelled as matching or non matching.
3. Other suggestions?

Question: Which of these would be easier to implement? Do you notice any problems?

Common Data Sets

Labelled Faces in the Wild: http://vis-www.cs.umass.edu/lfw/

- Faces are taken from Yahoo News, and have been labelled with many attributes, including race, gender, age, id.
- Demographics on the dataset are not provided by the publishers. Examples provided feature mostly white faces.

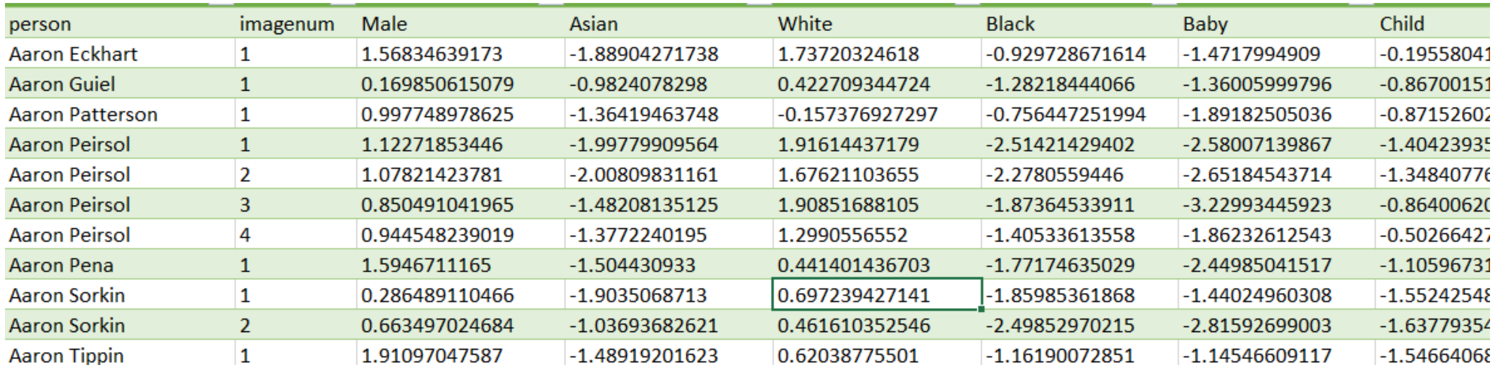

person	imagenum	Male	Asian	White	Black	Baby	Child
Aaron Eckhart	1	1.56834639173	-1.88904271738	1.73720324618	-0.929728671614	-1.4717994909	-0.19558041
Aaron Guiel	1	0.169850615079	-0.9824078298	0.422709344724	-1.28218444066	-1.36005999796	-0.86700151
Aaron Patterson	1	0.997748978625	-1.36419463748	-0.157376927297	-0.756447251994	-1.89182505036	-0.87152602
Aaron Peirsol	1	1.12271853446	-1.99779909564	1.91614437179	-2.51421429402	-2.58007139867	-1.40423935
Aaron Peirsol	2	1.07821423781	-2.00809831161	1.67621103655	-2.2780559446	-2.65184543714	-1.34840776
Aaron Peirsol	3	0.850491041965	-1.48208135125	1.90851688105	-1.87364533911	-3.22993445923	-0.86400620
Aaron Peirsol	4	0.944548239019	-1.3772240195	1.2990556552	-1.40533613558	-1.86232612543	-0.50266427
Aaron Pena	1	1.5946711165	-1.504430933	0.441401436703	-1.77174635029	-2.44985041517	-1.10596731
Aaron Sorkin	1	0.286489110466	-1.9035068713	0.697239427141	-1.85985361868	-1.44024960308	-1.55242548
Aaron Sorkin	2	0.663497024684	-1.03693682621	0.461610352546	-2.49852970215	-2.81592699003	-1.63779354
Aaron Tippin	1	1.91097047587	-1.48919201623	0.62038775501	-1.16190072851	-1.14546609117	-1.54664068

- **Question**: Why would it be important to know the demographics of the dataset?

Worksheet – Connecting (15 min)

1. In your own words, how does Facial Recognition Technology Work? Is it similar to any other technology we have seen so far?

2. Brainstorm possible beneficial uses of this technology? How can this solve a modern day problem, or be useful in day to day life?

Known Bias

According to NIST study of FRT:

- False Match Rates are 10 to 100 times higher for:
  - Older People
  - Women
  - Non-White People

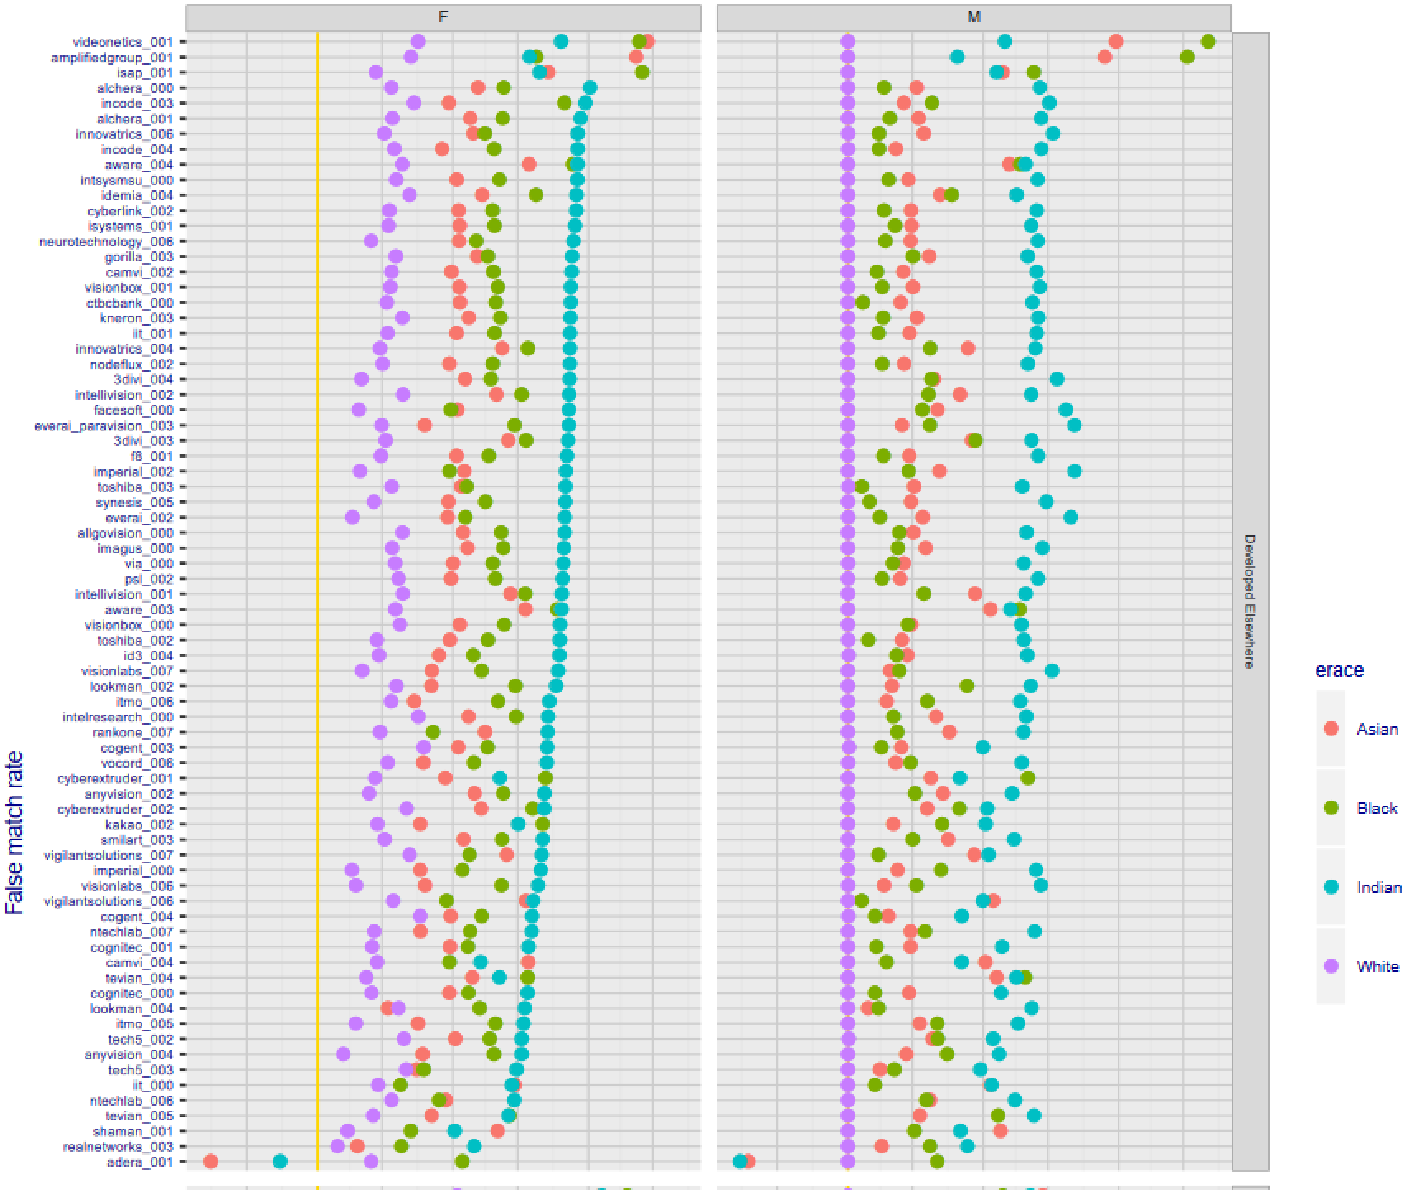

https://doi.org/10.6028/NIST.IR.8280

Applications – Personal Security

- Verify it's really you trying to access your phone, device, home, etc.
- Amazon Astro: Home robot for security and assistance which uses facial recognition and tracking

Applications – Police Surveillance

- Given video footage of a crime scene, compare faces in the video against a database of known faces to identify a potential culprit.
- Databases can include pictures from many sources: mug shots, government ids, publicly available internet photos (social media, news, …)

Figure 2: Selected Federal, State, and Non-government Systems with Facial Recognition Technology Used by Federal Agencies that Employ Law Enforcement Officers, and the Number of Photos in Them

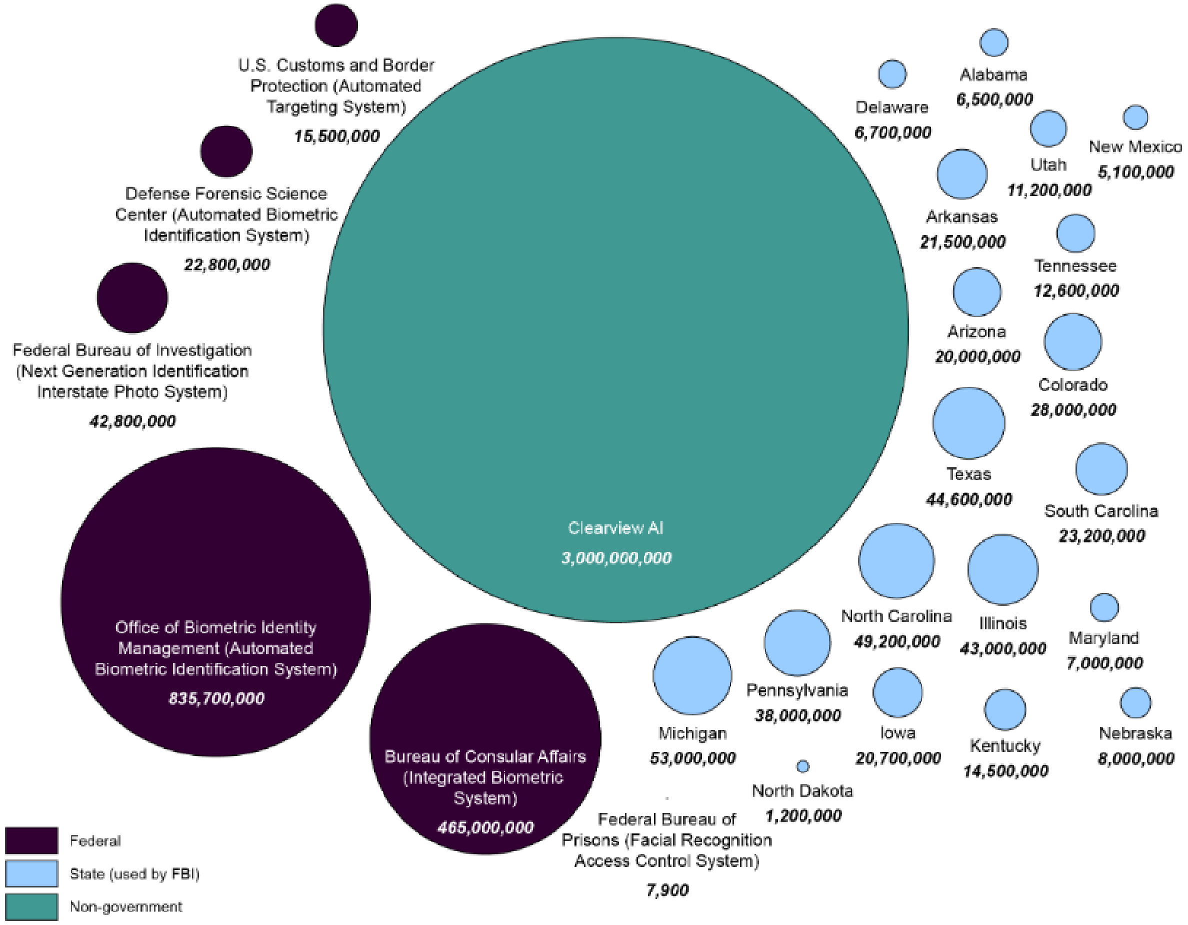

Source: GAO analysis of information provided by system users or owners. | GAO-21-105309

Applications – Police Surveillance

- Identification models are not perfectly accurate. They have been shown to work less well for people with darker skin.
- What happens if we trust AI completely in this case?
  - Wrongful arrests/incarceration
  - Emotional strain on the victim (embarrassment, loss of dignity, powerlessness)
  - Financial costs to fight the wrongful arrest (which might be unaffordable to some)
- This has happened.

Robert Williams (pictured here with his daughter) was wrongfully arrested due to FRT use by Michigan Police in 2019.

More on his story: https://www.aclu.org/news/privacy-technology/i-did-nothing-wrong-i-was-arrested-anyway

Applications – Eye Tracking

- Track where someone's focus is based on eye movement.
  - Interact with websites using eyes (accessibility tech)
  - Allow a professional gamer/athlete to hone their skills by evaluating their focus and improving their focus.
  - Track cheating for online exams, by checking if students are looking away from the screen.

Applications – Eye Tracking

- Cheating detection was done during the pandemic by multiple institutions.
- An eyetracker does not know why someone looked away from the screen:
  - Distractions taking an exam in a busy household.
  - Needing to look away from the screen to rest your eyes/think.
  - Inability to stay focussed on a screen for long periods of time.
- What happens if we fully trust the AI in this case?
  - Wrongful accusations leading to students suspended/expelled
  - Emotional/Mental strain fighting the accusations
  - Feeling powerless
- These detection systems disproportionately labelled students of colour, students who wear religious facial coverings, students with disabilities as cheaters.

Worksheet – Reflecting (15min)

- Does this technology work the same for everyone? What groups does it work better on?

- What does a “mistake” look like for this technology?

- What are some potentially harmful outcomes from these mistakes (or from use of this tech in general) Is this exacerbated by bias?

Recommended Reading

- Chapter 3 – More than a Glitch: Confronting Race, Gender and Ability Bias in Tech by Meredith Broussard
- Robert Williams' Story - https://www.aclu.org/news/privacy-technology/i-did-nothing-wrong-i-was-arrested-anyway

References

https://www.nytimes.com/2020/06/24/technology/facial-recognition-arrest.html

https://centerforhealthjournalism.org/our-work/insights/when-world-went-remote-communities-wrong-side-digital-divide-got-shut-out

https://medium.com/backprop-labs/face-recognition-pipeline-clearly-explained-f57fc0082750

https://www.technologyreview.com/2020/08/07/1006132/software-algorithms-proctoring-online-tests-ai-ethics/

Face Recognition Vendor Test (FRVT), Part 3: Demographic Effects. Patrick Grother, Mei Ngan, Kayee Hanaoka. 2019. https://doi.org/10.6028/NIST.IR.8280