U.S. Department of Homeland Security

SCIENCE AND TECHNOLOGY DIRECTORATE

Understanding and Mitigating Bias in Human & Machine Face Recognition

Science and Technology

John J. Howard
Principal Scientist
Identity and Data Sciences Laboratory at
the Maryland Test Facility

Arun Vemury
Lead
Biometric and Identity Technology Center
DHS Science & Technology Directorate

April 2023

Disclaimer

- This research was funded by the U.S. Department of Homeland Security, Science and Technology Directorate on contract number 70RSAT18CB0000034.

- This work was performed by the Identity and Data Sciences Laboratory team at the Maryland Test Facility.

- The views presented here are those of the authors and do not represent those of the Department of Homeland Security, the U.S. Government, or their employers.

- The data used in this research was acquired under IRB protocol or is publicly available non-PII data.

Science and Technology

Agenda

- The Maryland Test Facility

- Demographic differentials or “bias” in Face Recognition:
    - What is it?
    - Where does it come from?
    - Why are they bad?
    - How do we measure it?
    - How do we fix it?

Biometric & Identity Technology Center

S&T conducts foundational research to ensure advancements in science and technology are harnessed for cutting-edge solutions to new and emerging operational challenges.

- Drive biometric and identity innovation at DHS through RDT&E capabilities
- Facilitate and accelerate understanding of biometrics and identity technologies for new DHS use cases
- Drive efficiencies by supporting cross cutting methods, best practices, and solutions across programs
- Deliver Subject Matter Expertise across the DHS enterprise
- Engage Industry and provide feedback
- Encourage Innovation with Industry and Academia

The Maryland Test Facility (MdTF)

- Founded in 2014 by the Department of Homeland Security, Science and Technology Directorate.
- 20,000 ft^2 of office and reconfigurable laboratory space
- Fully instrumented and designed for human subject testing
  - Data collection infrastructure: Cameras, ambient light, noise, humidity, real time control center and monitoring capability, informed consent collection facilities, etc.
- Since its founding over 2500 subjects have participated in biometric testing at the MdTF
  - Ages 18-72
  - 114 countries of origin

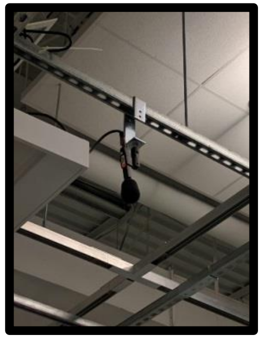



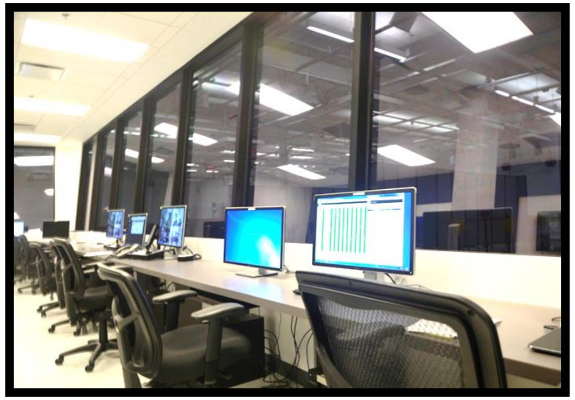

Science and Technology

DHS S&T Biometric Technology Rallies

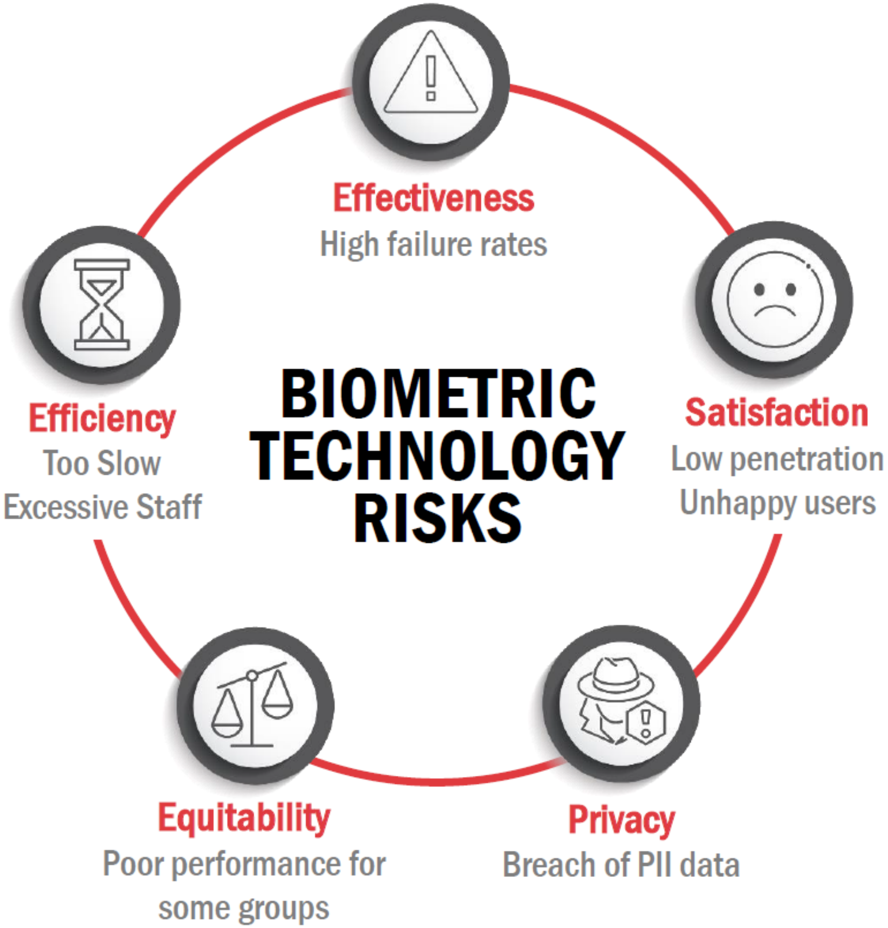

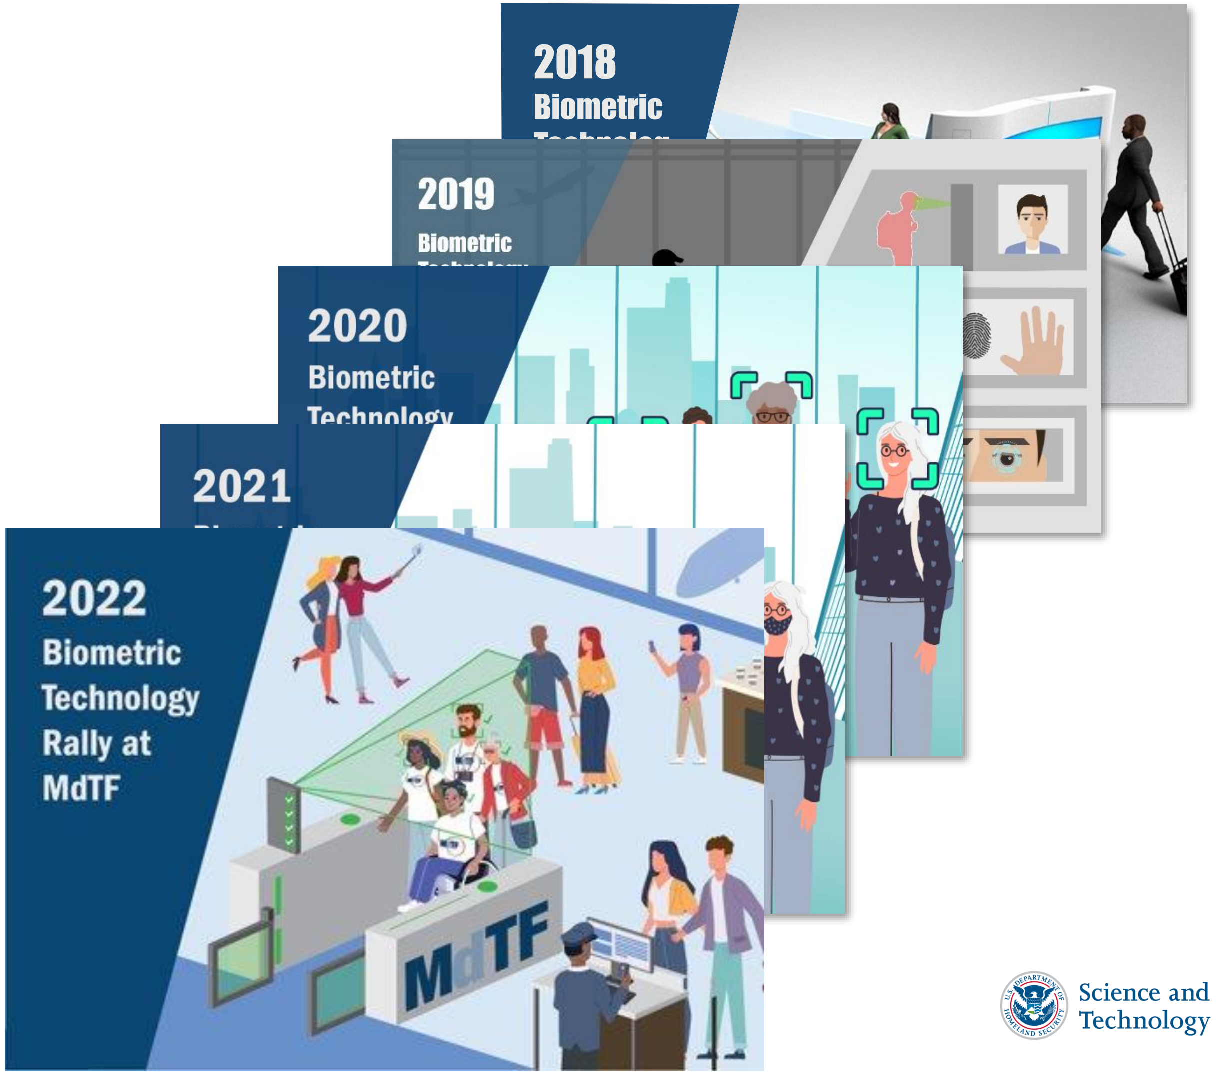

What is demographic "bias" in FR

nature

Explore content ˅ About the journal ˅ Publish with us ˅ Subscribe

nature > news feature > article

NEWS FEATURE | 18 November 2020

Is facial recognition too biased to be let loose?

The technology is improving – but the bigger issue is how it's used.

The Alan Turing Institute

Understanding bias in facial recognition technologies

ACLU

NEWS & COMMENTARY

How is Face Recognition Surveillance Technology Racist?

SITN science in the news celebrating 20 years

HARVARD UNIVERSITY
The Graduate School of Arts and Sciences

OCTOBER 24, 2020

BLOG, SCIENCE POLICY, SPECIAL EDITION: SCIENCE POLICY AND SOCIAL JUSTICE

Racial Discrimination in Face Recognition Technology

MIT Schwarzman College of Computing

HE CASES AUTHOR RESOURCES

Winter 2021 ˅ Published on Feb 05, 2021 DOI 10.21428/2c646de5.62272586

The Bias in the Machine: Facial Recognition Technology and Racial Disparities

What is demographic “bias” in FR

- Despite all the attention, the term “bias” is not well defined
- Overloaded term (computer science, statistics, psychology, public discourse)
- Not specific enough (How is it biased? Does it have an impact?)
- Howard, Sirotin, Vemury. *The Effect of Broad and Specific Demographic Homogeneity on the Imposter Distributions and False Match Rates in Face Recognition Algorithm Performance (2019).*

What is demographic “bias” in FR

- **False negative differential** – tendency for a group not to match
- **False positive differential** – tendency for a group to false match

Images available under Creative Commons and Fair Use criteria

Algorithm: No Match


FND(τ) =

If the rate that this happens

> or <

Algorithm: No Match


the rate that this happens

Algorithm: Match


FPD(τ) =

If the rate that this happens

> or <

Algorithm: Match


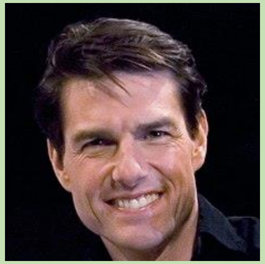

the rate that this happens

Science and Technology

Where does “bias” in FR come from

- Many sources:
  - Most people will highlight **data**

  - Far fewer people bring up:
    - Loss function
    - **Evaluation bias & historical anchoring**
    - **Our own brains**
      - **Projection bias** (we think machine ought to behave like us)
      - **Confirmation bias** (we like it when the machine confirms our beliefs)
      - **Automation bias** (we do what the machine tells us)

Evaluation Bias & Historical Anchoring

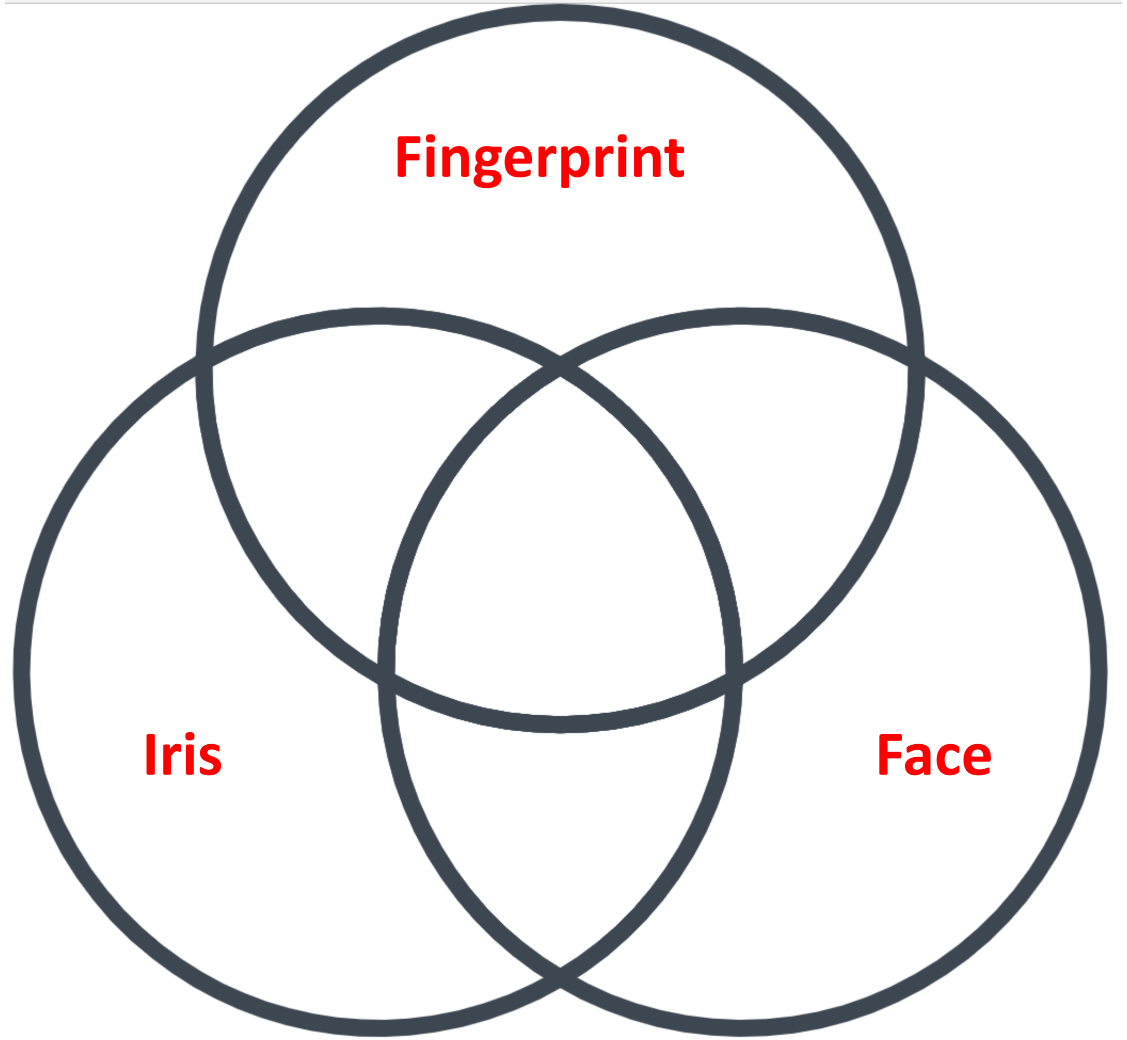

The **means** by which we evaluate fairness **impacts the outcome** of a fairness evaluation

Evaluation Bias & Historical Anchoring

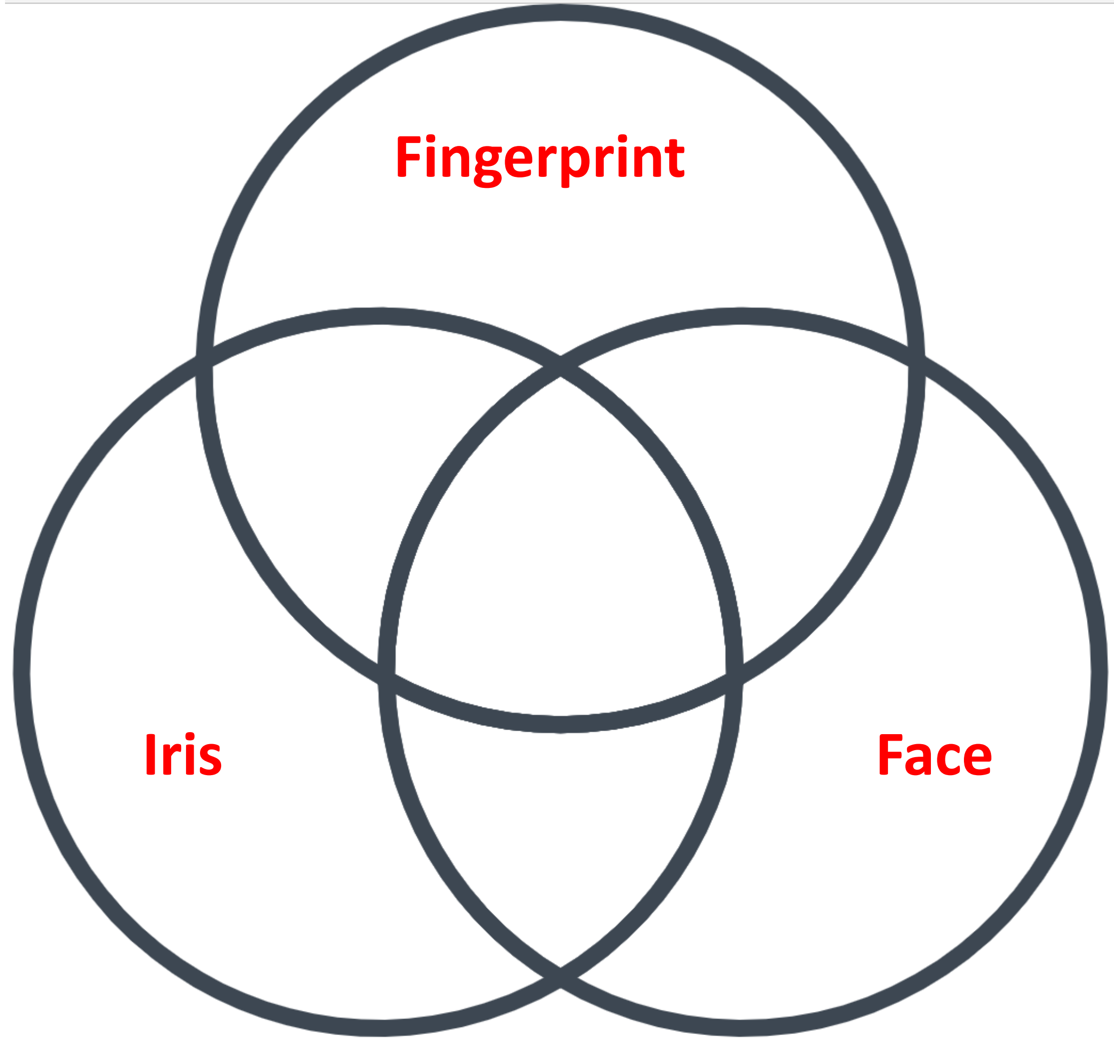

The **means** by which we evaluate fairness **impacts the outcome** of a fairness evaluation

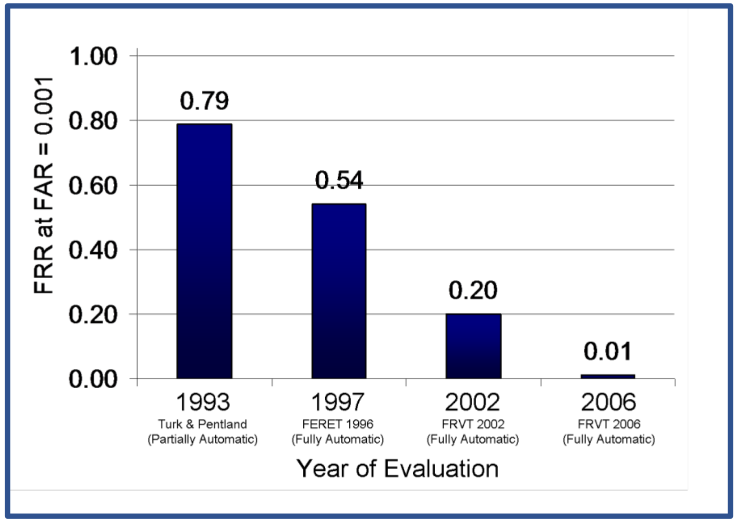

Evaluation Bias & Historical Anchoring

Are we taking lessons from here?

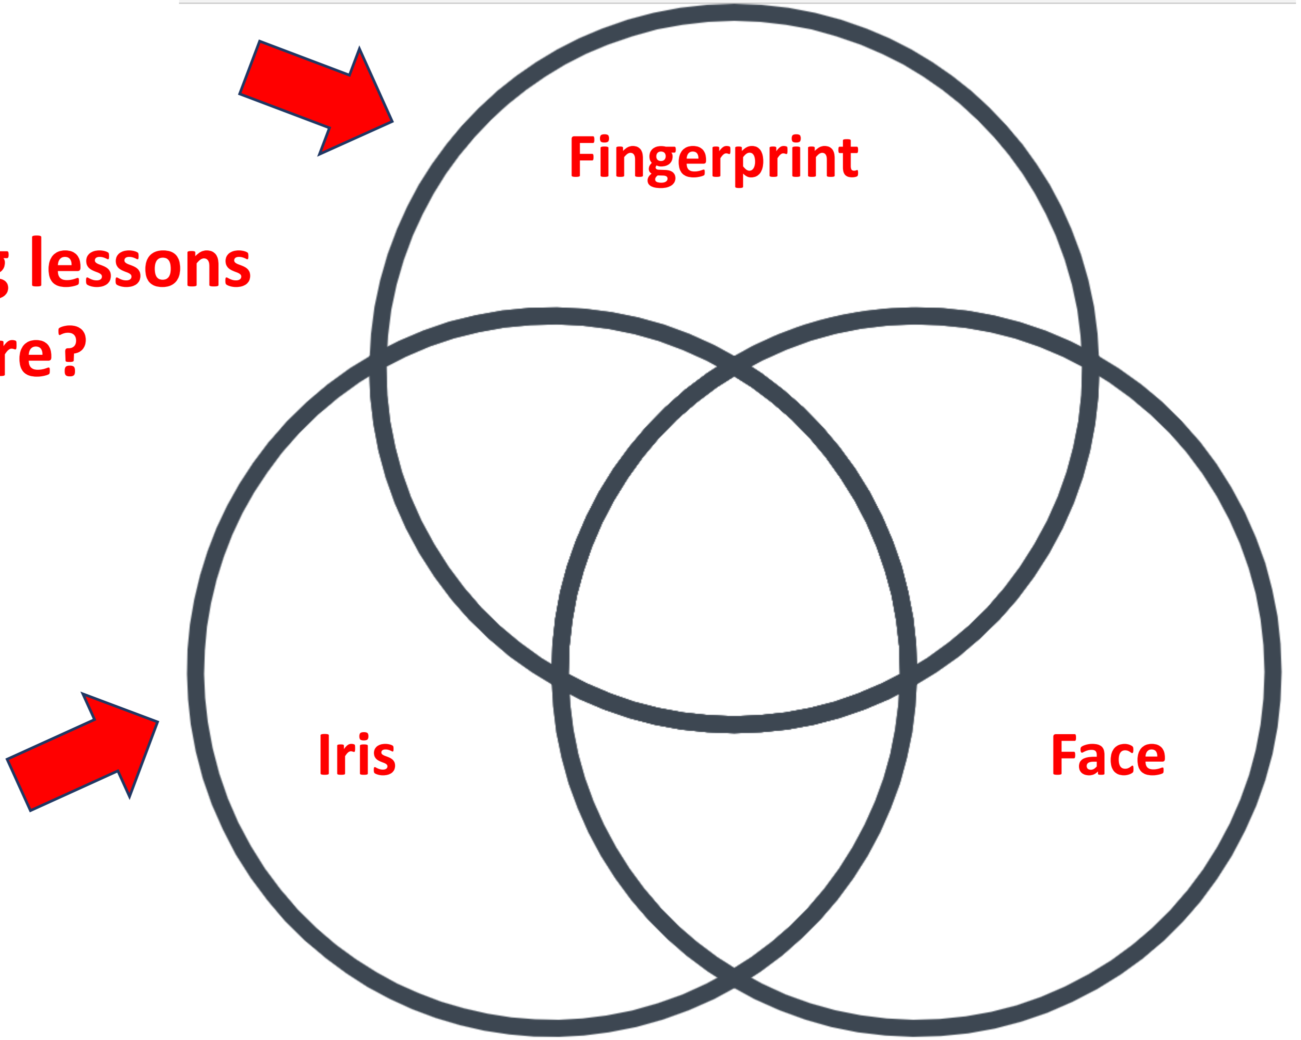

Evaluation Bias & Historical Anchoring

Are we taking lessons from here?

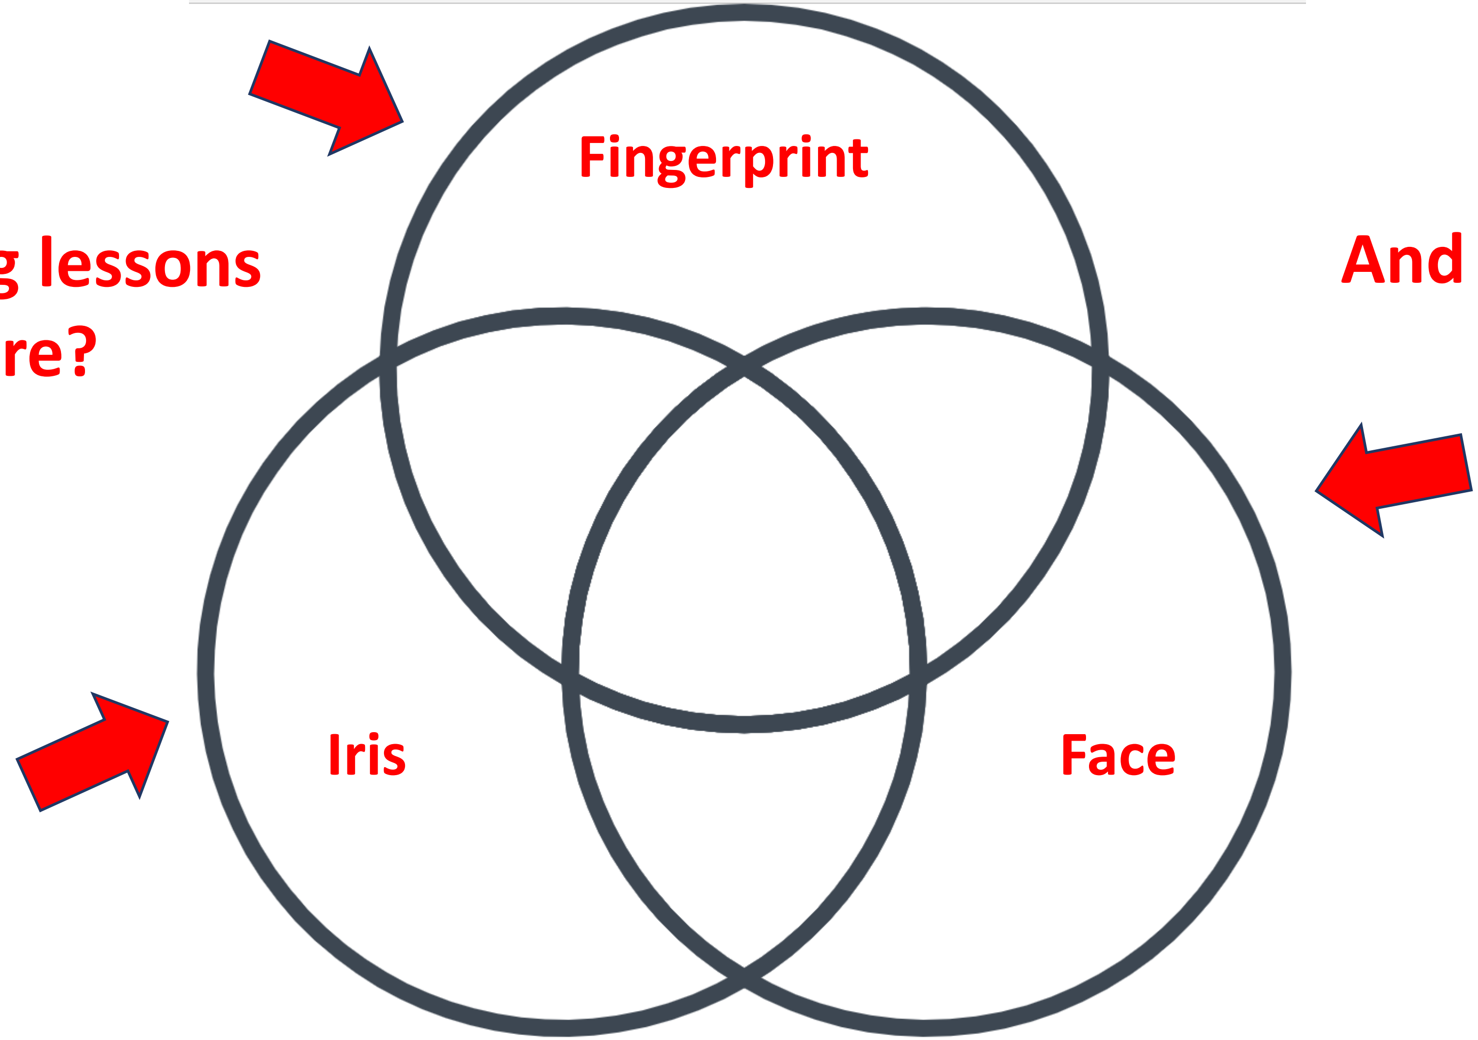

And applying them here?

Evaluation Bias

May be appropriate because this space exists

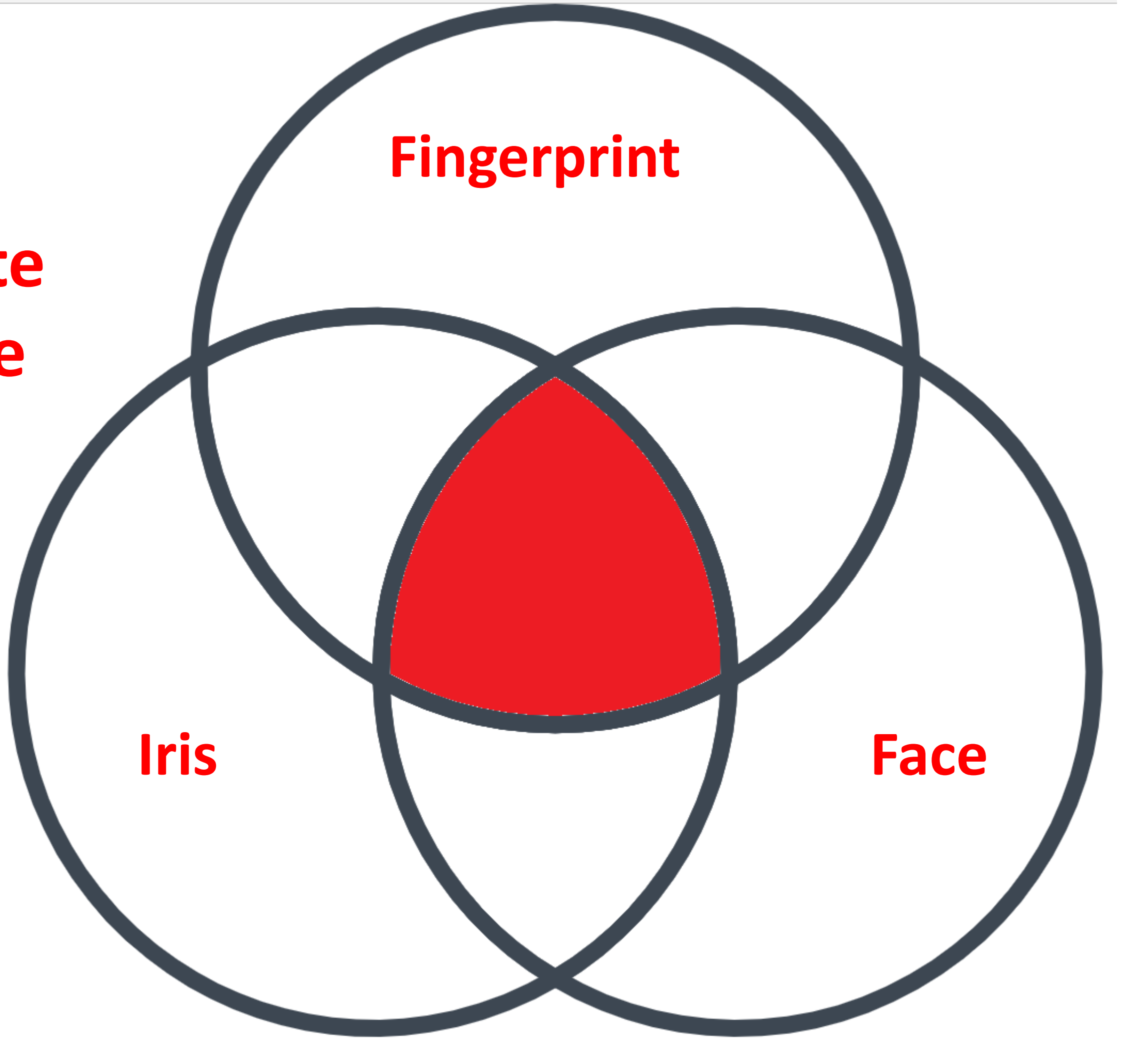

Evaluation Bias

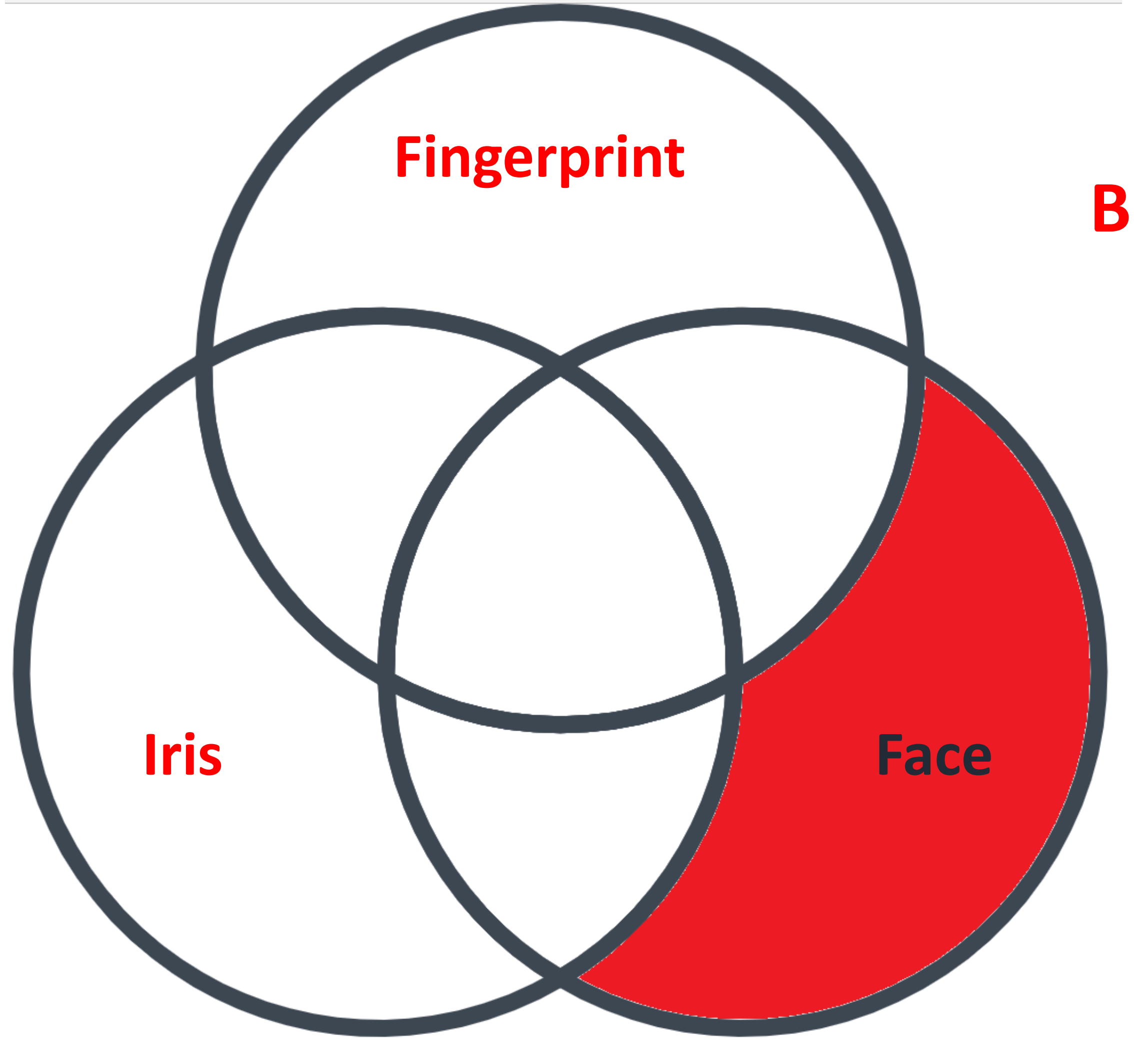

But we need to keep in mind that this space exists as well

Faces are Different for (at least) Two Reasons

- Faces are **genetic**, iris and fingerprint characteristics are determined during development.
  - Face are more alike for siblings, those with common ancestry, and those of the same sex

- Humans have an **innate ability** to perform face recognition tasks, not so with iris and fingerprints.
  - Humans have dedicated brain areas that process faces quickly
  - This was an important function for human evolution
    - Mates, Friends, Foes, Family members
    - Other primates have a similar capability
  - Intuitively perceive same-gender and same-race faces as more similar
  - We even know the exact part of the human brain dedicated to face processing.
    - Evolved to recognize familiar individuals within small social groups (25-100)
  - Prosopagnosia – “face blindness”

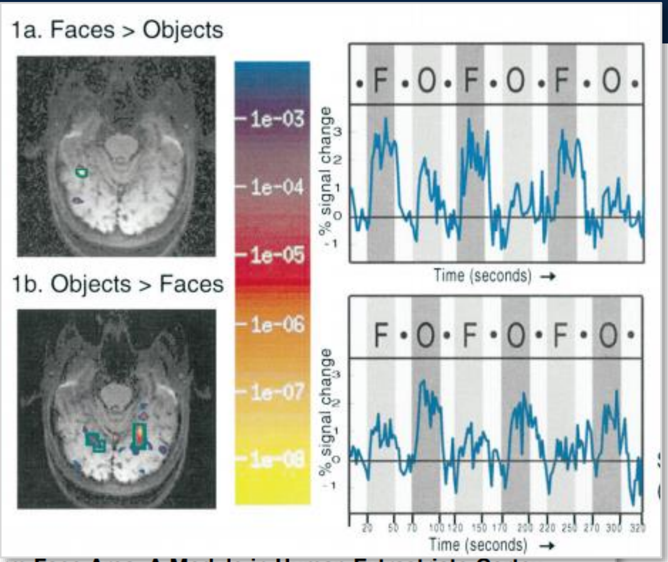

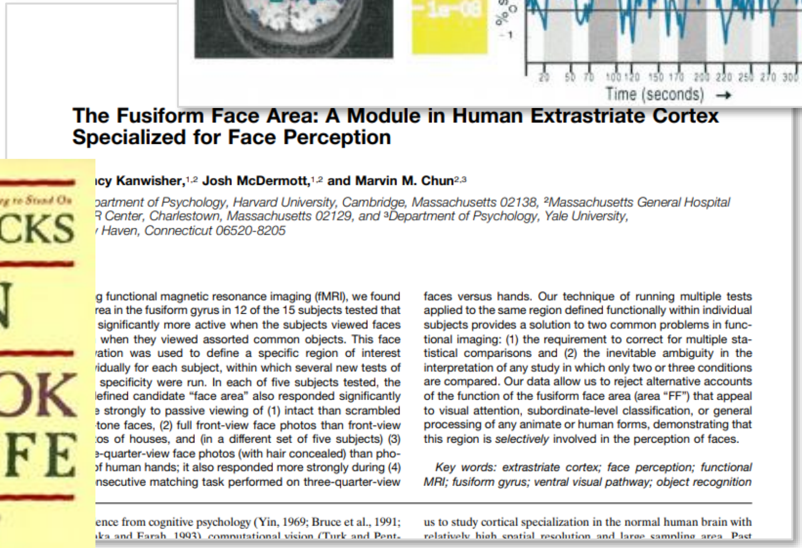
The Fusiform Face Area: A Module in Human Extrastriate Cortex Specialized for Face Perception

...cy Kanwisher,[1,2] Josh McDermott,[1,2] and Marvin M. Chun[2,3]

...partment of Psychology, Harvard University, Cambridge, Massachusetts 02138, [2]Massachusetts General Hospital ...R Center, Charlestown, Massachusetts 02129, and [3]Department of Psychology, Yale University, ... Haven, Connecticut 06520-8205

...faces versus hands. Our technique of running multiple tests applied to the same region defined functionally within individual subjects provides a solution to two common problems in functional imaging: (1) the requirement to correct for multiple statistical comparisons and (2) the inevitable ambiguity in the interpretation of any study in which only two or three conditions are compared. Our data allow us to reject alternative accounts of the function of the fusiform face area (area “FF”) that appeal to visual attention, subordinate-level classification, or general processing of any animate or human forms, demonstrating that this region is *selectively* involved in the perception of faces.

Key words: extrastriate cortex; face perception; functional MRI; fusiform gyrus; ventral visual pathway; object recognition

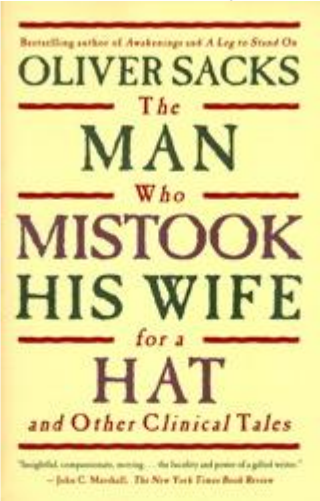

Demographic Effects Exist, Our Understanding of Them may be Clouded.

> It may seem natural to us that FR “clusters” people based on race and gender (projection bias) <

Iris recognition

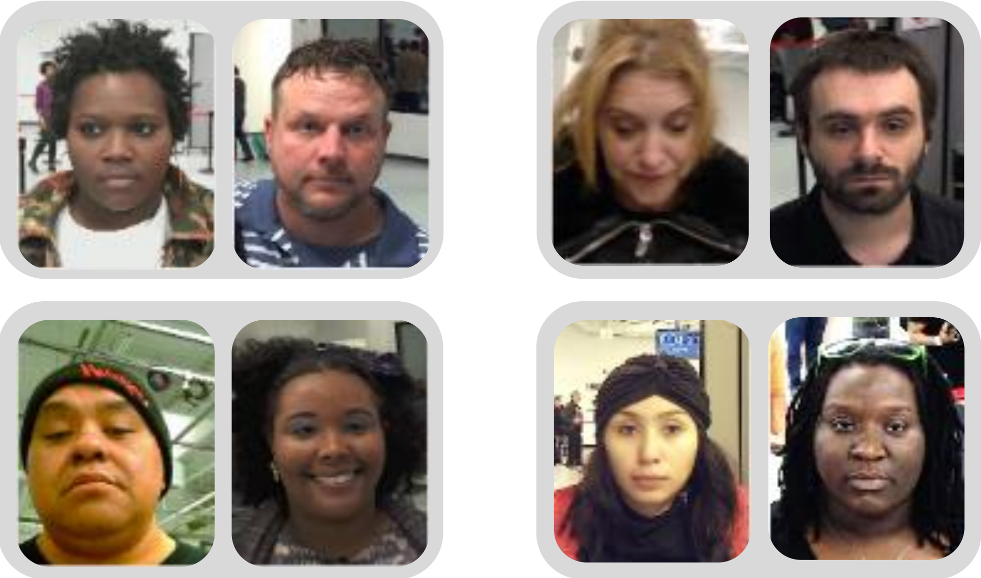

Iris recognition false positives were random relative to race and gender

Face recognition

80% of face recognition false positives were between people of the same race and gender

Subjects consent for use of their image in publications was obtained

Apples and Apples or Apples and Oranges?

> All of these “errors” are called “false matches”, but those on the right are different than those on the left <

Iris recognition

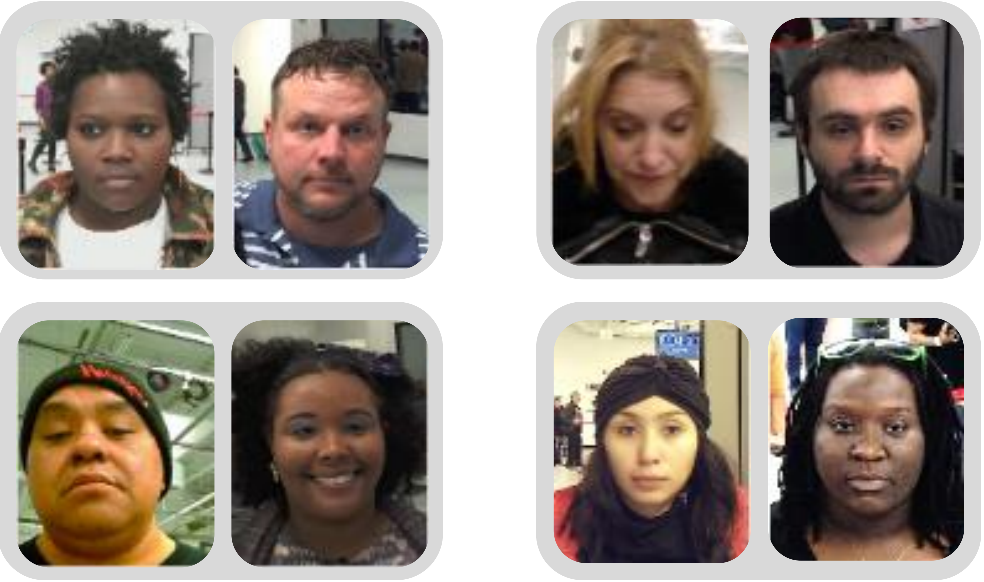

Iris recognition false positives were random relative to race and gender

Face recognition

80% of face recognition false positives were between people of the same race and gender

ience and
Technology

Subjects consent for use of their image in publications was obtained

Problem – When an algorithm errors in this way, it makes the human's job harder & slower

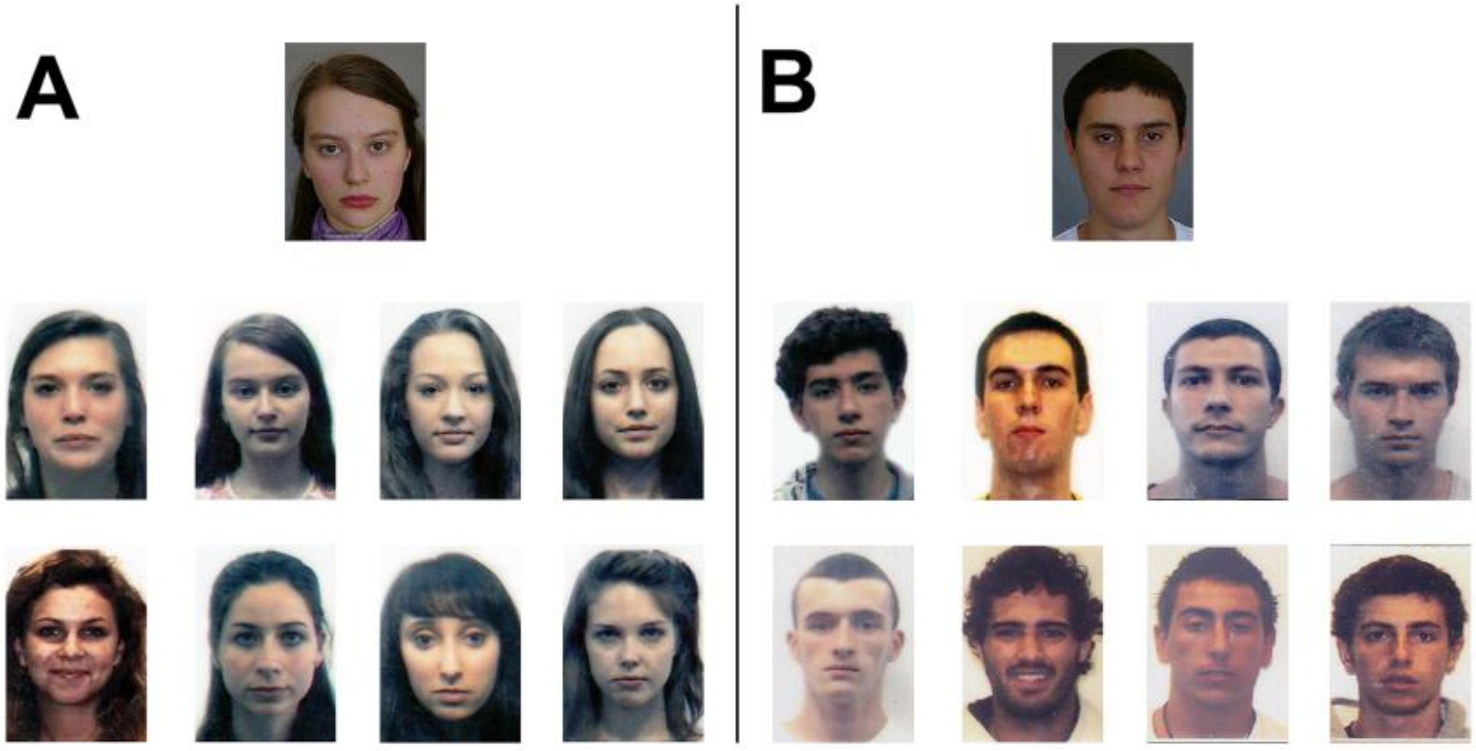

- White et. al "Error Rates in Users of Automatic Face Recognition Software"
- **50% - 60%** errors rates
- If ability of the human to correct the error is the distinguishing factor, **within group false match is not the same as an out group false match**

Problem – Hard tasks are more susceptible to automation bias

- Howard, Rabbitt, Sirotin, *Human-algorithm teaming in face recognition: How algorithm outcomes cognitively bias human decision-making.* PLoS 2020

- 343 volunteers performed face matching task (12 face pairs)
  - Glasglow Face Matching Test (8 pairs)
  - Select stimuli from MEDS for diversity in pairs (4 face pairs)

- Asked to rate similarity on a 7-point scale:

-3	**I am absolutely certain these are different people**
-2	**I am mostly certain these are different people**
-1	**I am somewhat certain this is the different person**
0	**I am not sure**
1	**I am somewhat certain these are same people**
2	**I am mostly certain this is the same person**
3	**I am absolutely certain this is the same person**

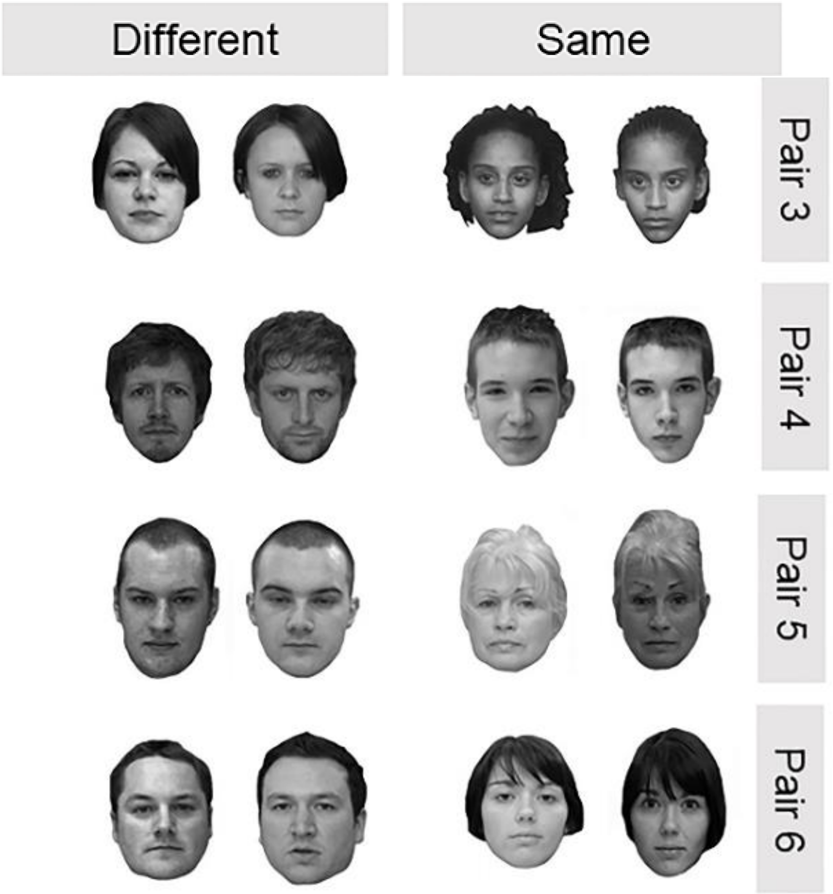

Automation Bias in FR

- Subjects were given face pairs under two conditions:

Control:

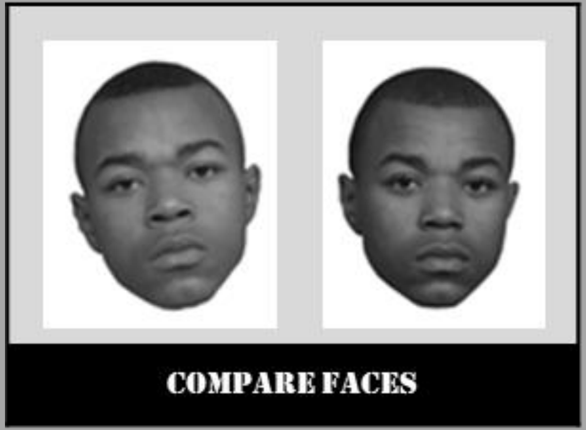

With a-priori identity information:

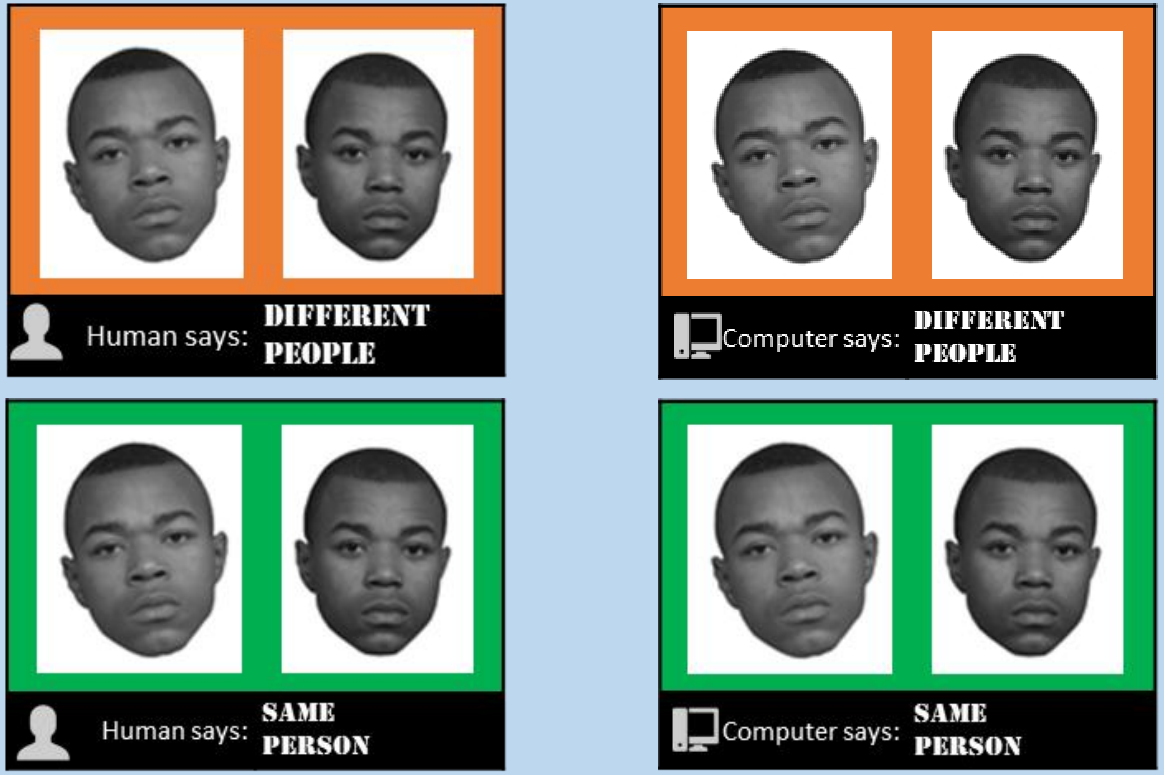

Automation Bias in FR

- At a threshold of 0.5:

Source	N	Accuracy	FPR	TPR
Control	120	0.75	0.19	0.70
Same	223	0.73	0.25	0.72
Different	223	0.75	0.17	0.66

No Match	-3	I am absolutely certain these are different people
	-2	I am mostly certain these are different people
	-1	I am somewhat certain this is the different person
	0	I am not sure
Match	1	I am somewhat certain these are same people
	2	I am mostly certain this is the same person
	3	I am absolutely certain this is the same person

Automation Bias in FR

- Across thresholds:

-3	I am absolutely certain these are different people
-2	I am mostly certain these are different people
-1	I am somewhat certain this is the different person
0	I am not sure
1	I am somewhat certain these are same people
2	I am mostly certain this is the same person
3	I am absolutely certain this is the same person

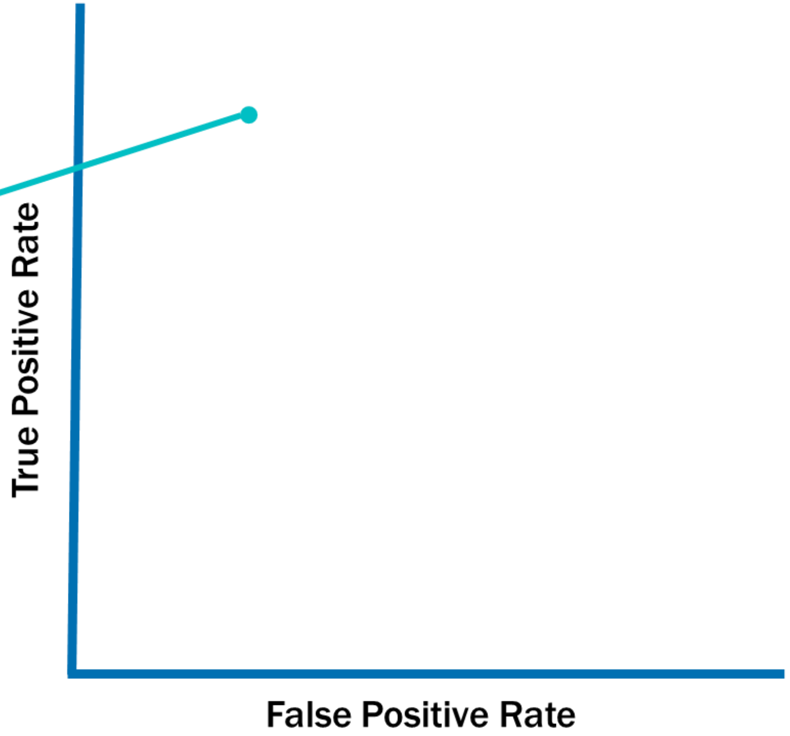

Source	FPR	TPR
Control	0.19	0.70
Same	0.25	0.72
Different	0.17	0.66

Automation Bias in FR

- Across thresholds:

-3	I am absolutely certain these are different people
-2	I am mostly certain these are different people
-1	I am somewhat certain this is the different person
0	I am not sure
1	I am somewhat certain these are same people
2	I am mostly certain this is the same person
3	I am absolutely certain this is the same person

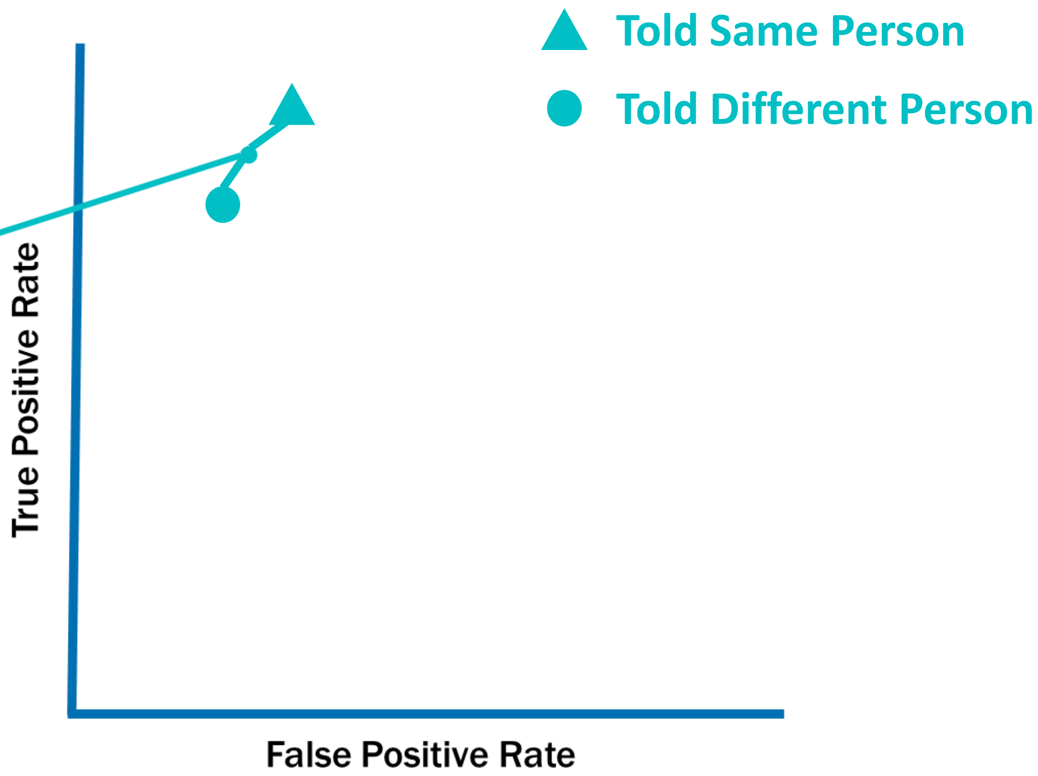

	Source	FPR	TPR
	Control	0.19	0.70
▲	Same	0.25	0.72
●	Different	0.17	0.66

Automation Bias in FR

- Across thresholds:
- The overlap in middling threshold indicates prior identity information can shift responses by a whole step
  - I am not sure → I am somewhat sure
- But only for **challenging** face pairs (I am not sure)
- Prior identity information effect was present but modest
- Humans mostly trusted their own abilities (**under ideal conditions**)

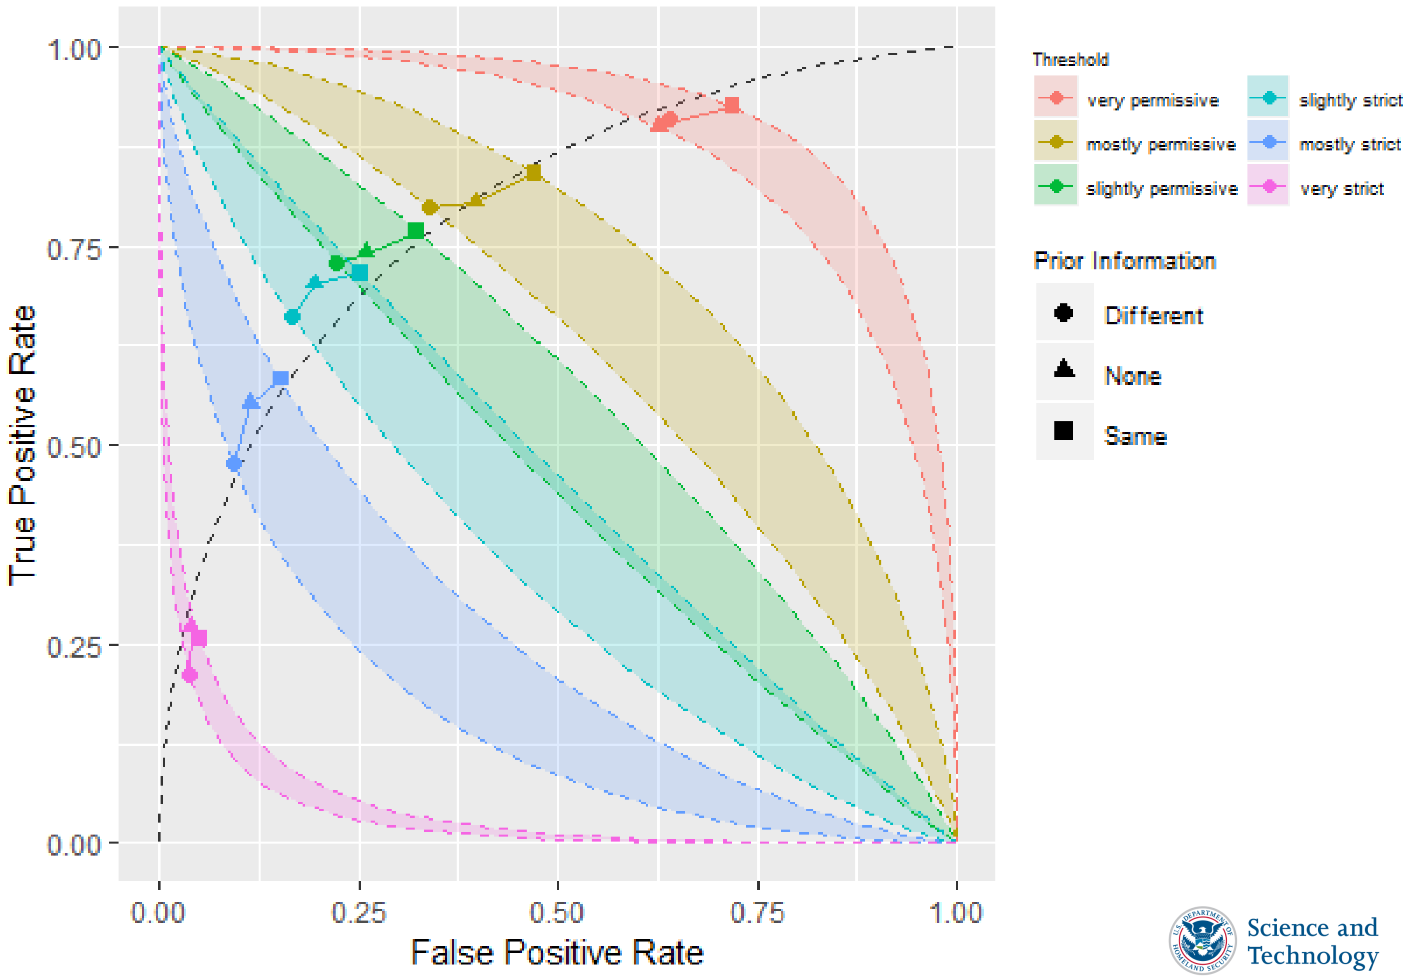

Automation Bias in FR (when it’s hard)

- Barragan, Howard, Rabbitt, Sirotin. *COVID-19 Masks Increase The Influence of Face Recognition Algorithm Decisions on Human Decisions in Unfamiliar Face Matching.* PLoS 2022

Automation Bias in FR (when it's hard)

- Barragan, Howard, Rabbitt, Sirotin. *COVID-19 Masks Increase The Influence of Face Recognition Algorithm Decisions on Human Decisions in Unfamiliar Face Matching.* PLoS 2022

Control

Computer-No Mask

Computer-Mask

Automation Bias in FR (when it's hard)

- 150 test subjects
- Largely replicated 2020 "No Mask" study

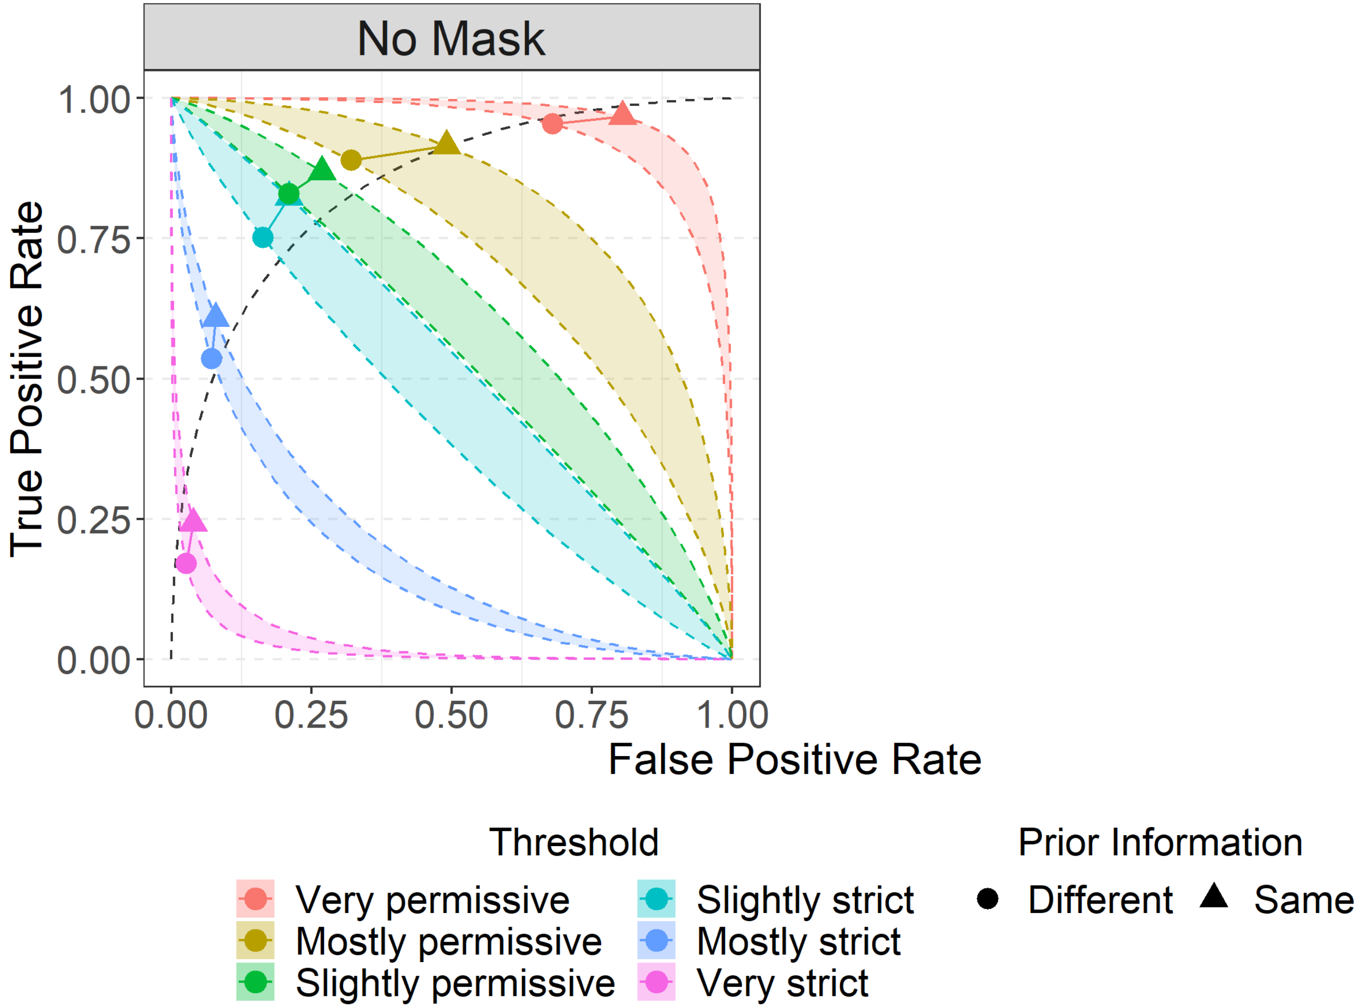

Automation Bias in FR (when it's hard)

- 150 test subjects
- Largely replicated 2020 "No Mask" study
- However, the presence of masks greatly increased the influence of prior algorithm information
- It also reduced accuracy 10-20% points.

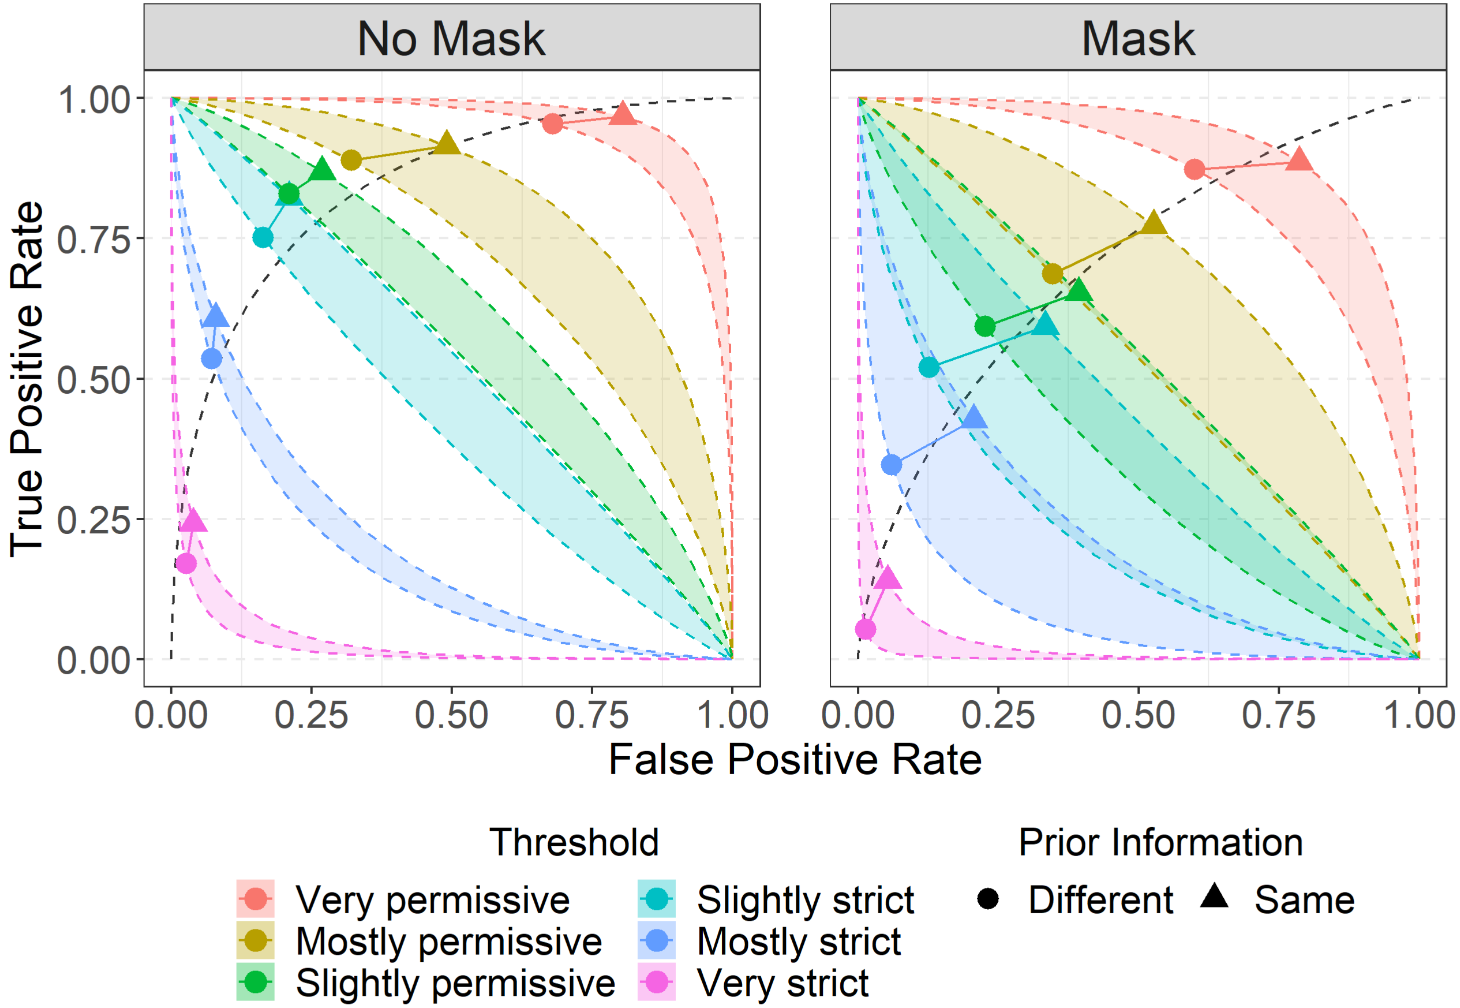

Automation Bias in FR (when it's hard)

- Our results showed that masks increased human reliance on algorithm determinations (if presented)

- Its likely (in our minds) that this is true for many factors that increase difficulty in face recognition tasks:
    - True across many categories of socio-technical systems (Google maps effect)
    - Lack of information in the image due to pose, blur, lighting etc.
    - Human perceived similarity **demographic homogeneity**

Agenda

- ~~The Maryland Test Facility~~

- Demographic differentials or “bias” in Face Recognition:
    - ~~What is it?~~
    - ~~Where does it come from?~~
    - ~~Why are they bad?~~
    - How do we measure it?
    - How do we fix it?

How do we measure demographic differentials

- Remember, these two things are **both** called a “false match error” in biometrics parlance:

Two people who share a similar **iris pattern** (according to an algorithm)

Two people who share a similar **face pattern** (according to an algorithm)

- Demographic **sameness, i.e. homogeneity** makes one of these much harder for a human to adjudicate

Science and Technology

Broad Homogeneity

- We coined the term “broad homogeneity” to describe this “sameness” effect in 2019

Different Demographics ⟷ Same Demographics

More similar

Algorithm score

Less similar

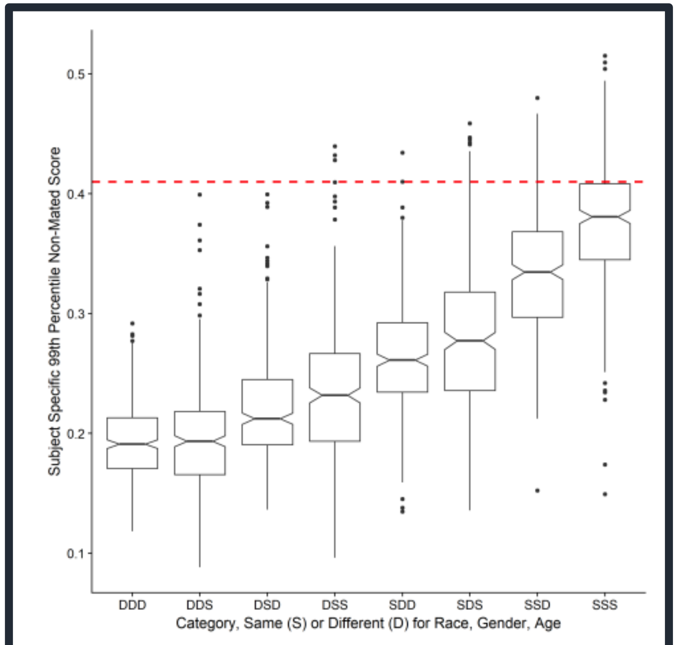

Figure 4. Distributions of the 99th percentile subject-specific non-mated scores across broad homogeneous versus heterogeneous race, gender, and age categories.

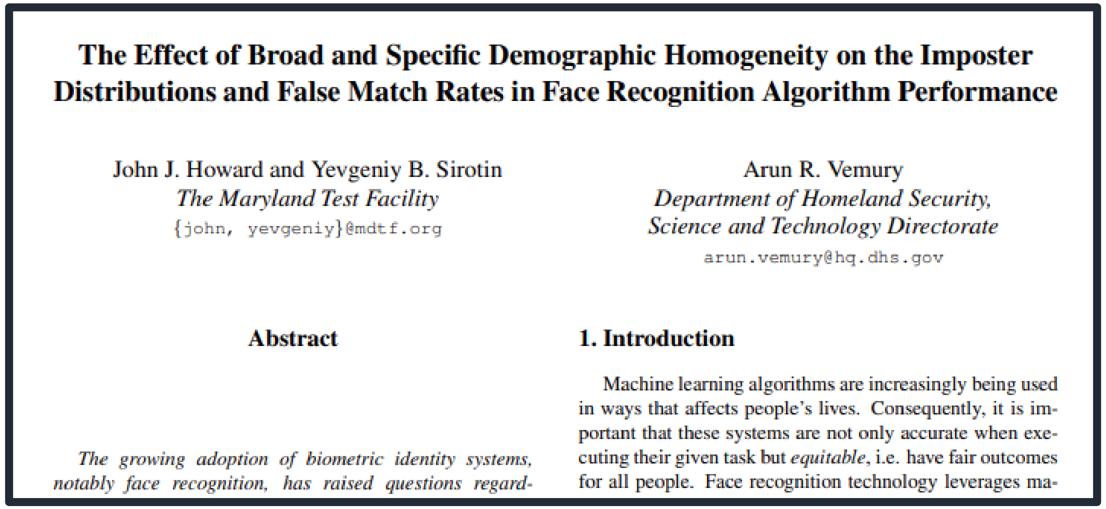

The Effect of Broad and Specific Demographic Homogeneity on the Imposter Distributions and False Match Rates in Face Recognition Algorithm Performance

John J. Howard and Yevgeniy B. Sirotin
The Maryland Test Facility
{john, yevgeniy}@mdtf.org

Arun R. Vemury
Department of Homeland Security, Science and Technology Directorate
arun.vemury@hq.dhs.gov

Abstract

The growing adoption of biometric identity systems, notably face recognition, has raised questions regard-

1. Introduction

Machine learning algorithms are increasingly being used in ways that affects people’s lives. Consequently, it is important that these systems are not only accurate when executing their given task but *equitable*, i.e. have fair outcomes for all people. Face recognition technology leverages ma-

- We showed this effect exists in **one** commercial face recognition algorithm

- Not present in iris or fingerprint biometrics

This is (Likely) (Currently) a Universal Feature of Face Recognition

- NIST subsequently confirmed this exists in **all 138 algorithms** submitted to FRVT in 2019.
    - NIST FRVT Part 3: Demographics – Annex 5.

Higher (non-mate) similarity score

More similar demographics

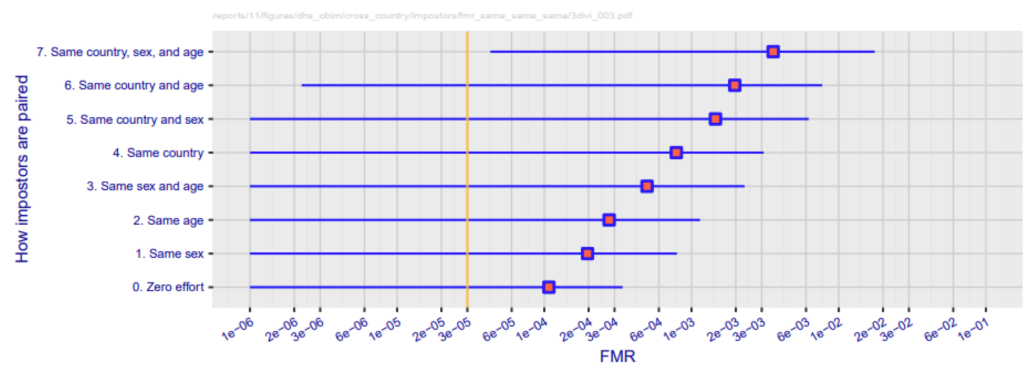

Figure 1: FMR for increasing matched covariates, 3divi-003

The Effect of Broad and Specific Demographic Homogeneity on the Imposter Distributions and False Match Rates in Face Recognition Algorithm Performance

John J. Howard and Yevgeniy B. Sirotin
The Maryland Test Facility
{john, yevgeniy}@mdtf.org

Arun R. Vemury
Department of Homeland Security, Science and Technology Directorate
arun.vemury@hq.dhs.gov

Abstract

stems,
egard-

1. Introduction

Machine learning algorithms are increasingly being used in ways that affects people's lives. Consequently, it is important that these systems are not only accurate when executing their given task but *equitable*, i.e. have fair outcomes for all people. Face recognition technology leverages ma-

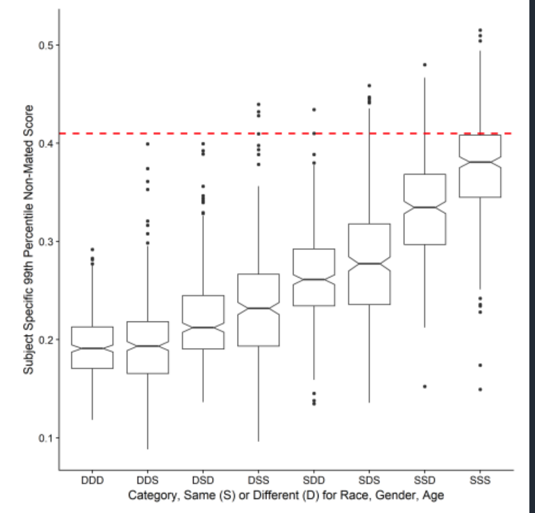

Figure 4. Distributions of the 99th percentile subject-specific non-mated scores across broad homogeneous versus heterogeneous race, gender, and age categories.

But There May be Solutions

- **IF** we recognize this as a problem..
- We may be able to address it
- Estimated **6 – 14%** of **face information content** clustered by race and gender (2021).

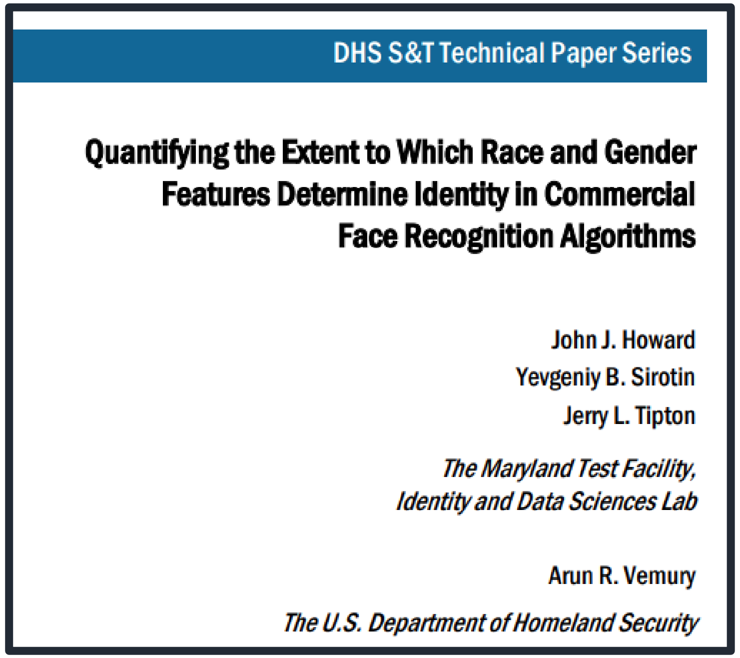
DHS S&T Technical Paper Series

Quantifying the Extent to Which Race and Gender Features Determine Identity in Commercial Face Recognition Algorithms

John J. Howard
Yevgeniy B. Sirotin
Jerry L. Tipton

The Maryland Test Facility, Identity and Data Sciences Lab

Arun R. Vemury

The U.S. Department of Homeland Security

Face Information Content?

- There are many detectable **points** on the human face
- The distances, shapes, and contours formed by those points make up some of the **face information** used by face recognition algorithms
- Some of that information content (but not all) **can cluster** people by ancestry, gender, etc.
- For example, male noses are on average **shorter** and **broader** than female noses

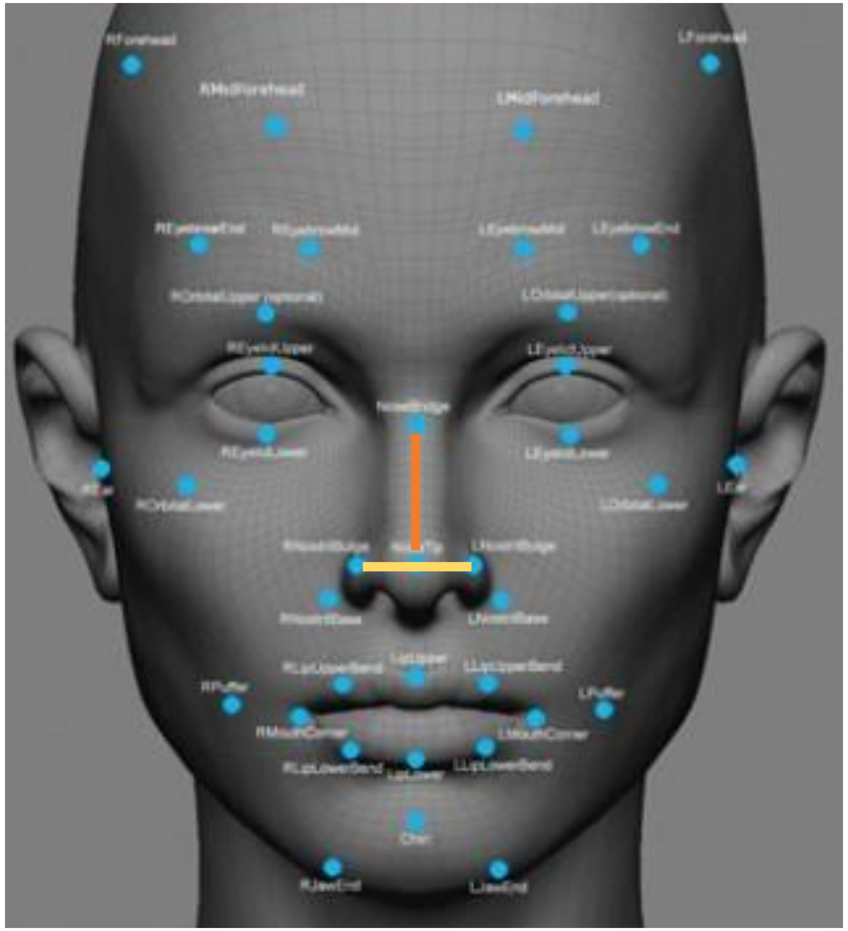

Face Information Content?

- We can visualize this clustering
- And measure it across many types of face information
- To find components that cluster (Comp.1, plot A)*
- And those that don't (Comp.3, plot B)*

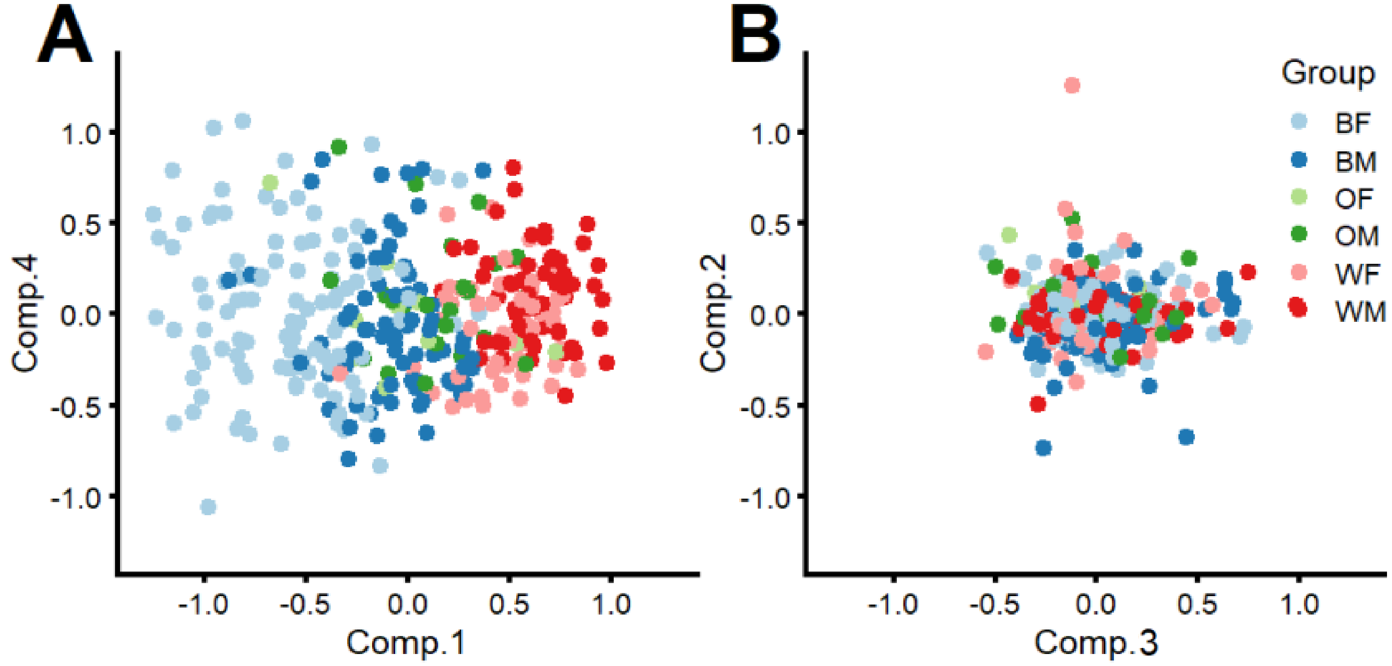

* Howard, Sirotin, Tipton, Vemury. *Quantifying the extent to which race and gender features determine identity in commercial face recognition algorithms*. DHS Technical Paper Series 2020.

Science and Technology

Selecting Face Information Content

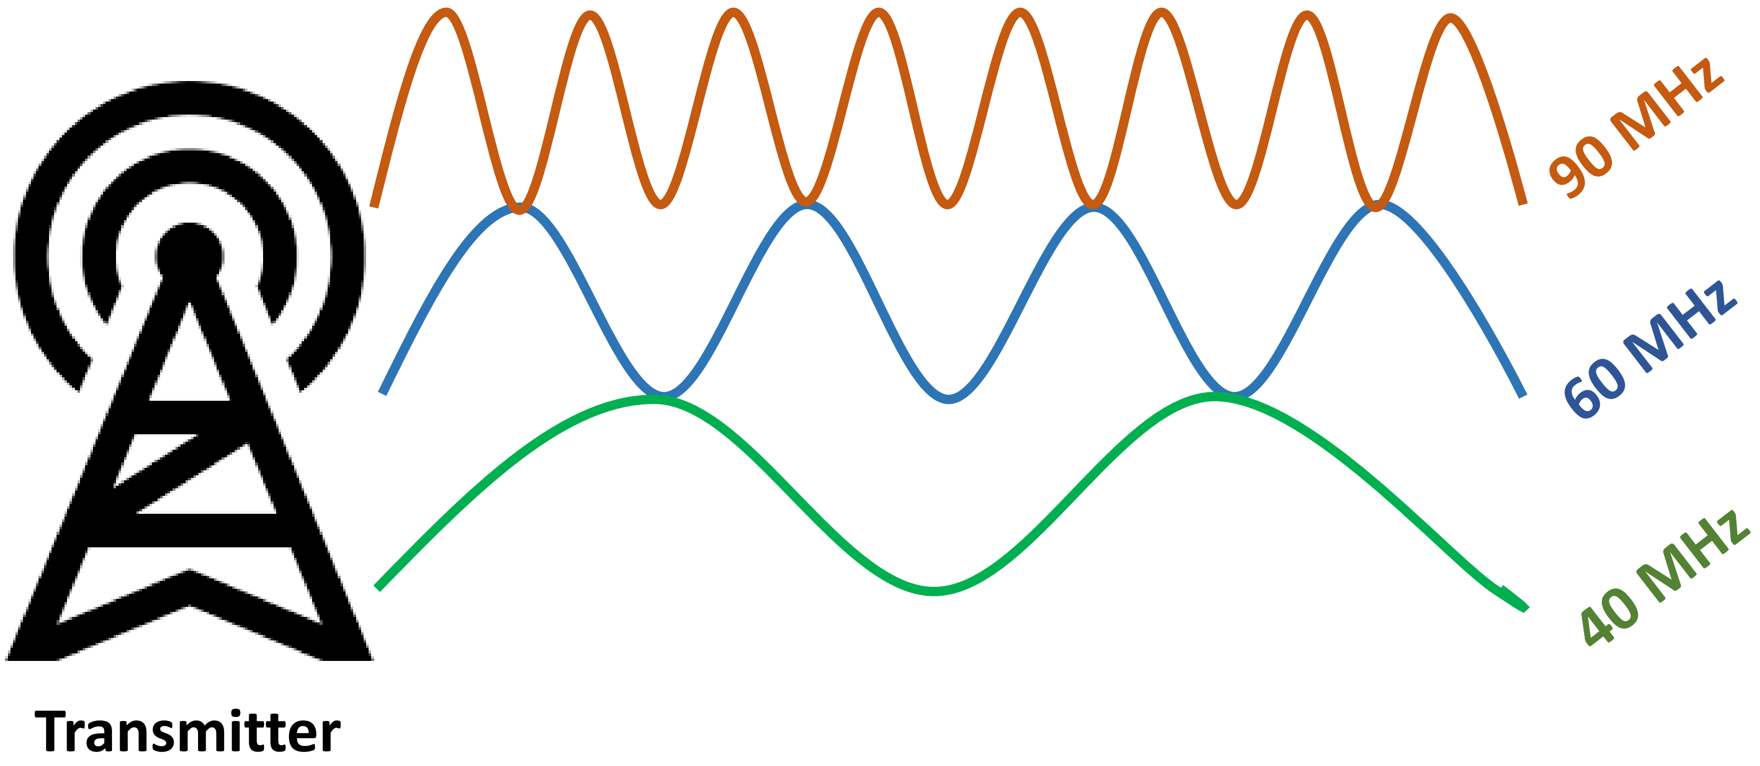

Face Information Content

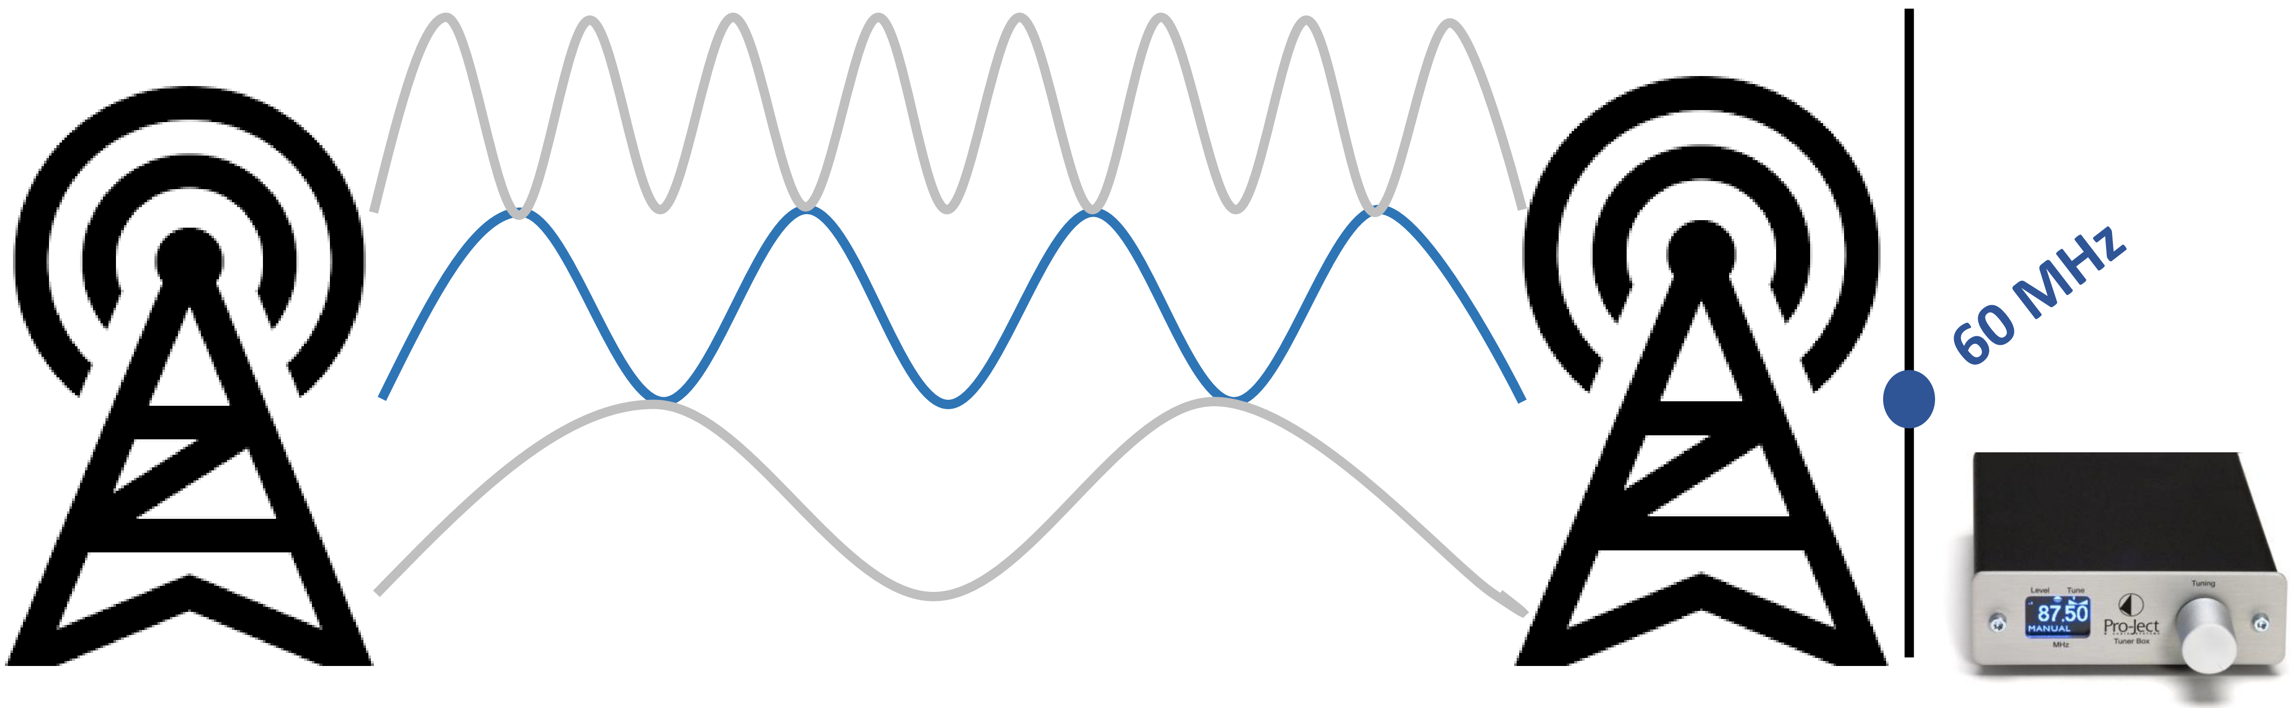

Transmitter

Receiver

Face Information Content

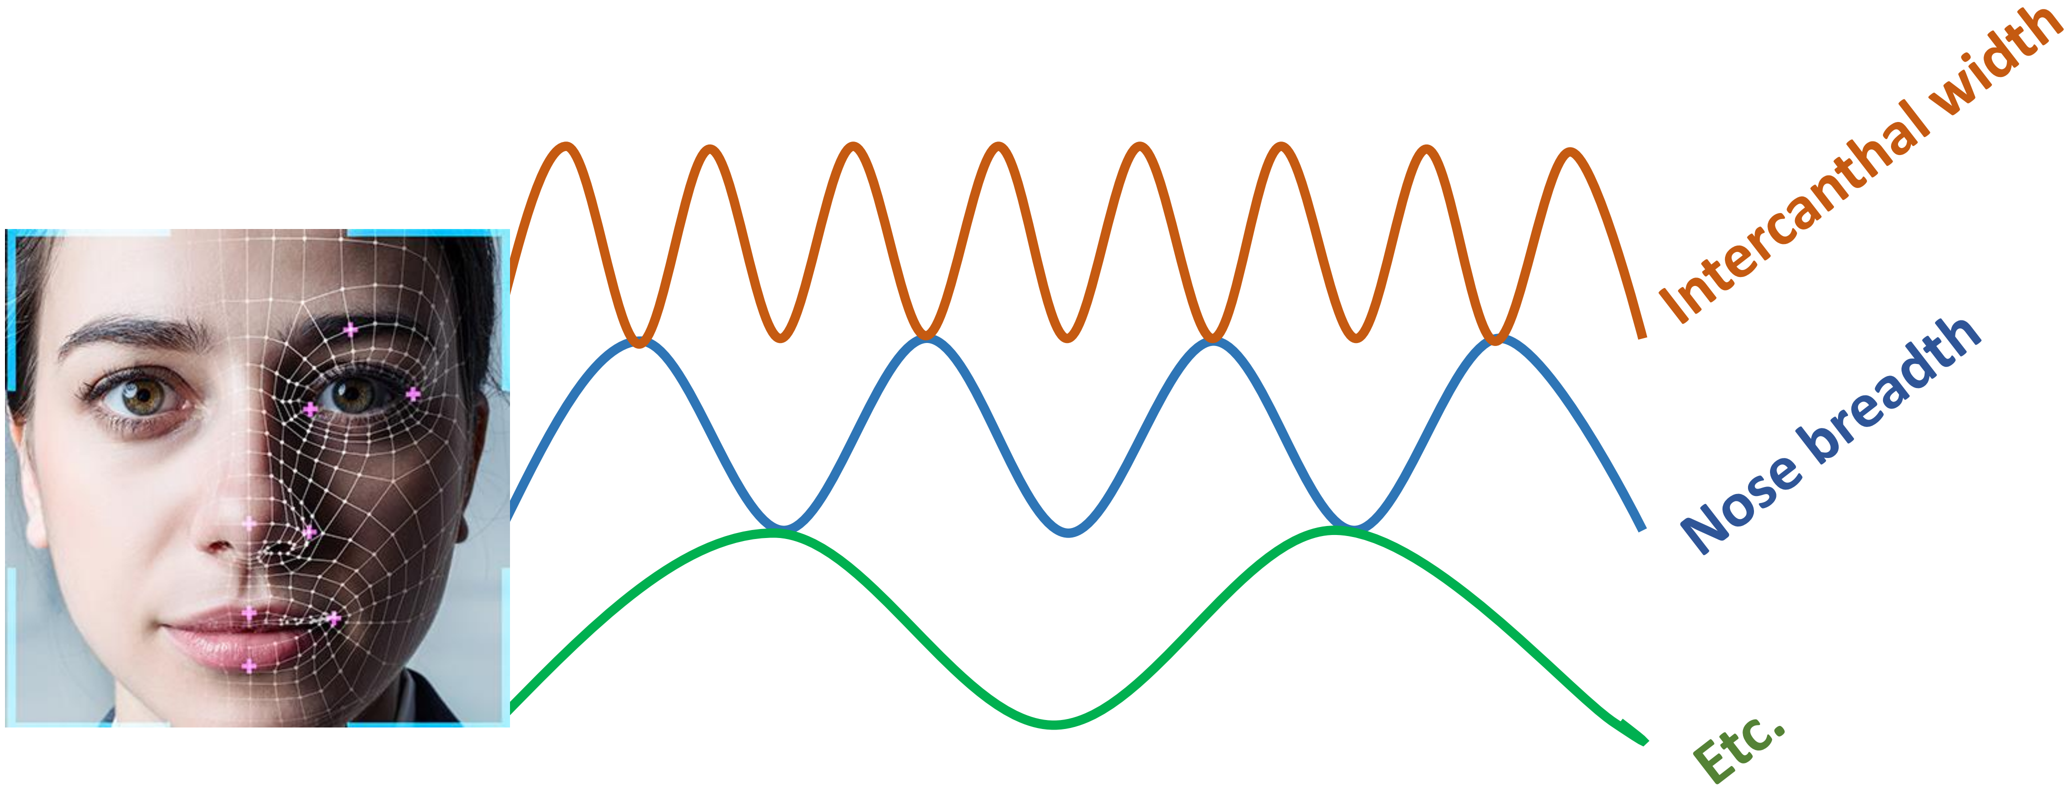

Human Face

Face Information Content

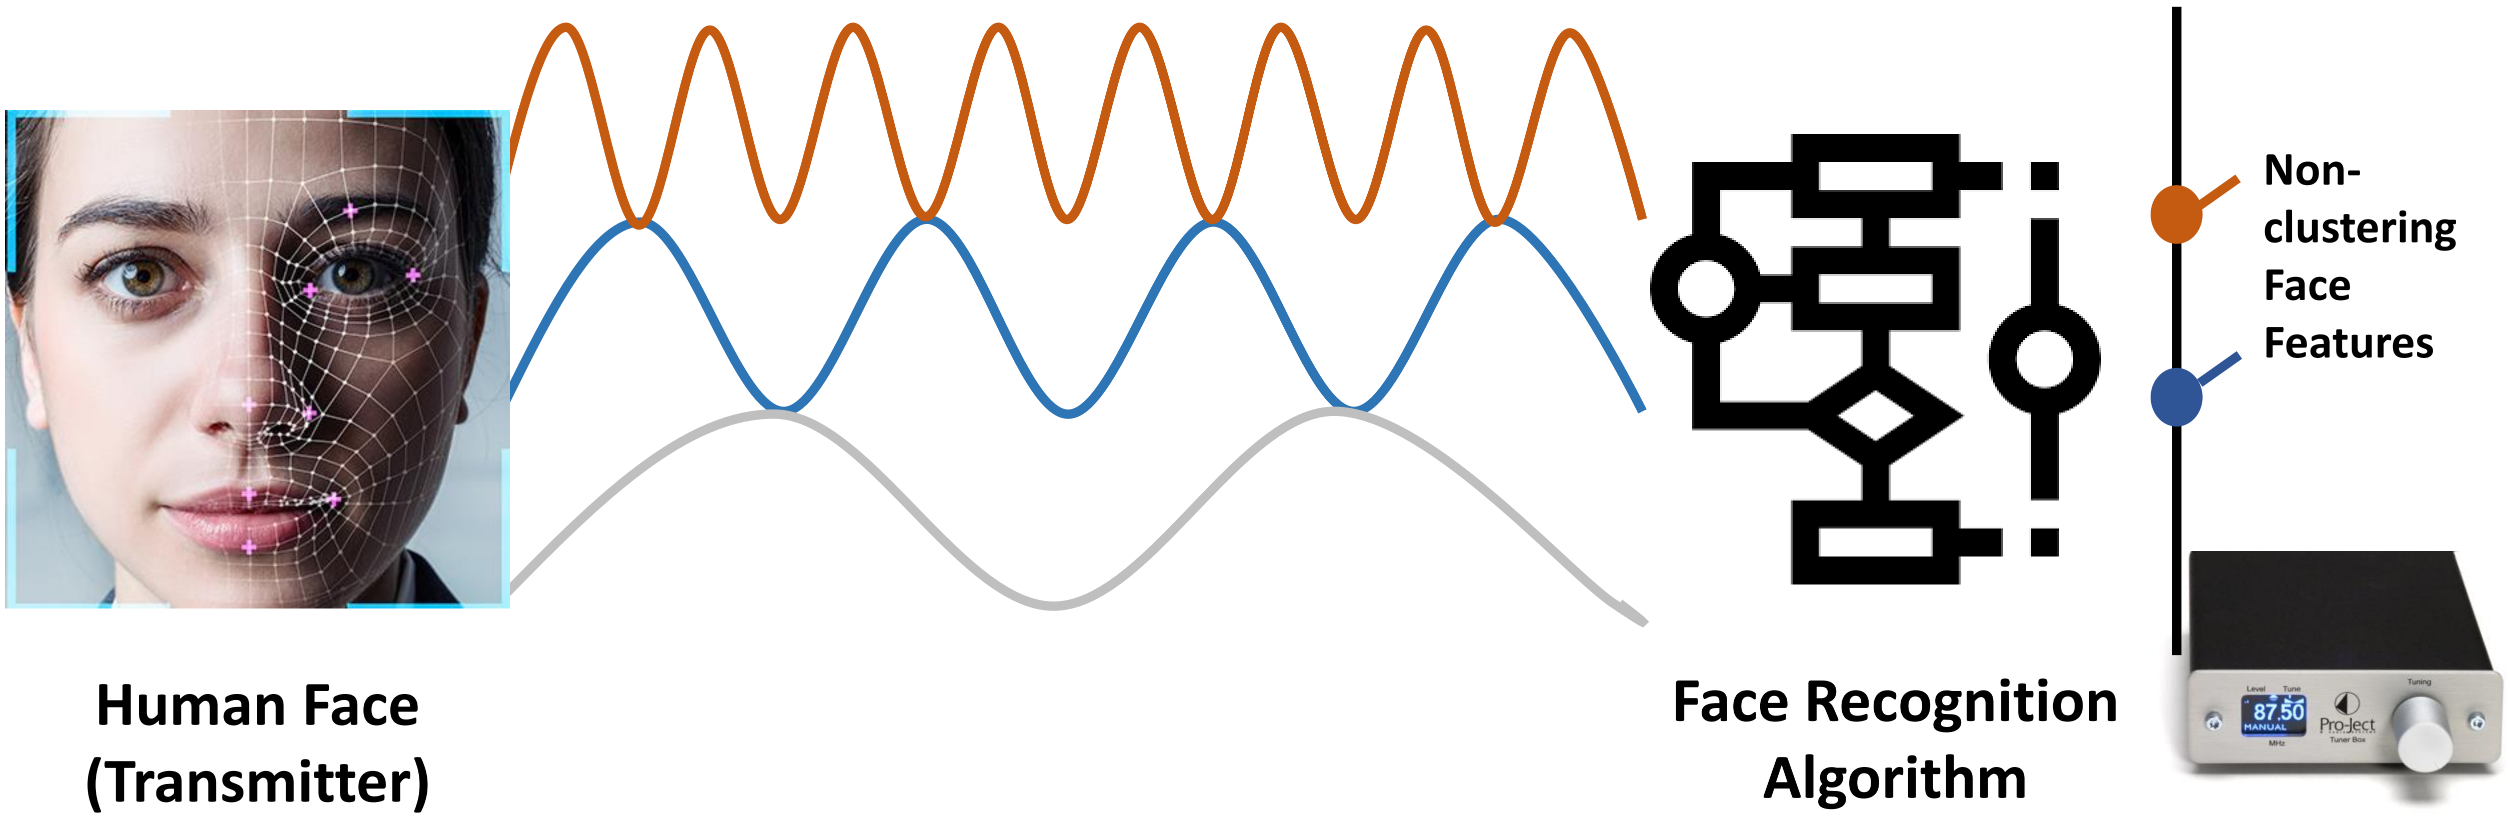

But There May be Solutions

- Estimated **6 – 14%** of face information content clustered by race and gender (2021).

DHS S&T Technical Paper Series

Quantifying the Extent to Which Race and Gender Features Determine Identity in Commercial Face Recognition Algorithms

John J. Howard
Yevgeniy B. Sirotin
Jerry L. Tipton

The Maryland Test Facility, Identity and Data Sciences Lab

Arun R. Vemury

The U.S. Department of Homeland Security

But There May be Solutions

- Estimated **6 – 14%** of face information content clustered by race and gender (2021).

- Showed a method to **remove this clustering** improved “fairness” across five different fairness measures (2022).

DHS S&T Technical Paper Series

Quantifying the Extent to Which Race and Gender Features Determine Identity in Commercial Face Recognition Algorithms

John J. Howard
Yevgeniy B. Sirotin
Jerry L. Tipton

The Maryland Test Facility,
...entity and Data Sciences Lab

Arun R. Vemury

...artment of Homeland Security

Appeared in 26th International Conference on Pattern Recognition (ICPR 2022), Fairness in Biometrics Workshop, Montreal, Quebec, August 2022.

Disparate Impact in Facial Recognition Stems from the Broad Homogeneity Effect: A Case Study and Method to Resolve

John J. Howard*[1], Eli J. Laird*†[1], and Yevgeniy B. Sirotin*[1]

The Identity and Data Sciences Lab at The Maryland Test Facility, Maryland, USA
{elaird, jhoward, ysirotin}@idslabs.org

Abstract. Automated face recognition algorithms generate encodings of face images that are compared to other encodings to compute a similarity score between the two originating face images. These face encodings, also known as feature vectors, contain representations of various facial features. Some of these facial features, but not all, have been shown to resemble each other across different subjects that happen to share a de-

What data did we use?

- Data
  - Three of face samples collected from the 2018-200 Biometric Technology Rallies:
    - S1 – demographically balanced training set
    - S2 – disjoint test set
    - S3 – mated pairs to subjects in S1

Dataset	Subjects (Samples)			
	Black Female	**Black Male**	**White Female**	**White Male**
S1	150 (150)	150 (150)	150 (150)	150 (150)
S2	50 (50)	50 (50)	49 (49)	43 (43)
S3	106 (300)	117 (339)	126 (321)	117 (278)

  - Two algorithms
    - ArcFace pre-trained on MS-Celeb-1M
    - ArcFace pre-trained on Glint 360k
  - Requirement for white box template structures

What did we do?

- **Goal:** Given a matrix V of face recognition **feature vectors**, identify components of those vectors that exhibit demographic clustering.
- **Process:**
  - SVD on normalized feature vector matrix, creates subject specific space (U) and a feature space (W^T)

$$\hat{V} = U\Sigma W^T, \text{where U} \in \mathbb{R}^{n\times n}, \Sigma \in \mathbb{R}^{n\times p}, W^T \in \mathbb{R}^{p\times p}$$

  - Calculate clustering index (C_k)

$$C_k = 1 - \frac{\sum_D \sum_{i\in D}(u_i - \bar{u}_D)^2}{\sum_i (u_i - \bar{u})^2}, \quad k, i \in \{1, \ldots, n\}$$

  - Identify components in U with $C_k > 99^{th}\ percentile$ of the bootstrapped C_k distribution

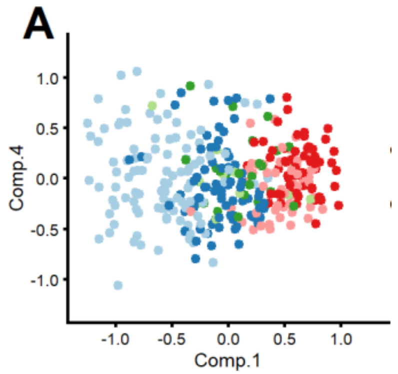

Science and Technology

What did we do?

- Given we found r components in the U matrix with statistically significant clustering

- Remove r columns from W which correspond to the r clustered components in U,
  - *Leaving* $\widehat{W} \in \mathbb{R}^{p \times m}$, *where* $m = p - r$

- Define **de-clustering transform** $\widehat{\mathrm{W}}\widehat{\mathrm{W}}^{\mathrm{T}}$

What did we do?

- Can apply $\widehat{\mathrm{W}}\widehat{\mathrm{W}}^{\mathrm{T}}$ to the set of feature vectors it was learned on
  - $\dot{V} = V\widehat{W}\widehat{W}^T$
  - **Q1**: How demographically “fair” are comparison scores generated from $\dot{V}$ versus V ?

- Can apply $\widehat{\mathrm{W}}\widehat{\mathrm{W}}^{\mathrm{T}}$ to any arbitrary face feature vector v (from the same algorithm)
  - $\dot{v} = v\widehat{W}\widehat{W}^T$
  - **Q2**: If we learn features that exhibit demographic clustering on one set of subjects, do those same featured cluster on other subjects?

What did we do?

- **Experiment 1** - De-clustering Learned and Applied to the Same Dataset (S1)
  - Performed *n x n* comparisons for S1 (360,000 comparisons)
  - Learned & Applied de-clustering transform to S1 feature vectors
  - Evaluated false match rate (FMR) differentials pre- and post-applying transformation

- **Experiment 2** - De-clustering Learned on One Dataset and Applied to a Disjoint Dataset (S2)
  - Performed *n x n* comparisons for S2 (36,864 comparisons)
  - Applied de-clustering transform learned on S1 to S2 feature vectors
  - Evaluated false match rate differentials (FMR) pre- and post-applying transformation

Dataset	Subjects (Samples)			
	Black Female	**Black Male**	**White Female**	**White Male**
S1	150 (150)	150 (150)	150 (150)	150 (150)
S2	50 (50)	50 (50)	49 (49)	43 (43)
S3	106 (300)	117 (339)	126 (321)	117 (278)

Science and Technology

How did we measure success?

- Five face recognition fairness measures:
  - Net Clustering [1]
  - Gini Aggregation Rate for Biometric Equitability (GARBE) [2]
  - Fairness Discrepancy Rate (FDR) [3]
  - NIST Inequity Ratio* – all ratios
  - NIST Inequity Ratio [4] – along the diagonal

- Investigated these measures at a threshold that gives a global FMR of 1e-3

- Broad homogeneity is a non-mated effect (alpha = 1, Beta = 0)

[1] Howard, J.J., Sirotin, Y.B., Tipton, J.L., Vemury, A.R.: Quantifying the extent to which race and gender features determine identity in commercial face recognition algorithms (2020)
[2] Howard, J., Laird, E., Sirotin, Y., Rubin, R., Tipton, J., and Vemury, A.. (2022). Evaluating Proposed Fairness Models for Face Recognition Algorithms.
[3] Pereira, T.d.F., Marcel, S.: Fairness in biometrics: a figure of merit to assess biometric verification systems. IEEE Transactions on Biometrics, Behavior, and Identity Science pp. 11 (2021). https://doi.org/10.1109/TBIOM.2021.3102862
[4] Grother, P.: Face recognition vendor test (frvt) part 8: Summarizing demographic differentials (2022)

What we found

- Most “fair” values are in **bold** (higher for FDR, lower for all others)
- Applying this demographic de-clustering **universally improved “fairness”**
- Across **two face recognition algorithms**
- Even when applied to an **“unknown” set of subjects** (S2)

Algorithm	Fairness Metric	Experiment 1		Experiment 2	
		S1 Original	S1 Transformed	S2 Original	S2 Transformed
ArcFace-MS1MV2	Net Clustering	0.0163	**0.00549**	0.0252	**0.0207**
	GARBE	0.8540	**0.65000**	0.922	**0.909**
	FDR	0.9900	**0.99900**	0.991	**0.993**
	INEQ	219.00	**30.2000**	22.00	**18.00**
	INEQ*	15.58	**3.74**	10.56	**6.62**
ArcFace-Glint360k	Net Clustering	0.0150	**0.00497**	0.0250	**0.0197**
	GARBE	0.8350	**0.67100**	0.955	**0.881**
	FDR	0.9910	**0.99900**	0.990	**0.996**
	INEQ	199.00	**22.1000**	12.5	**10.20**
	INEQ*	16.23	**3.67**	12.47	**3.68**

What does this do to human review?

- Pulled two rank 4 probe and candidate lists:

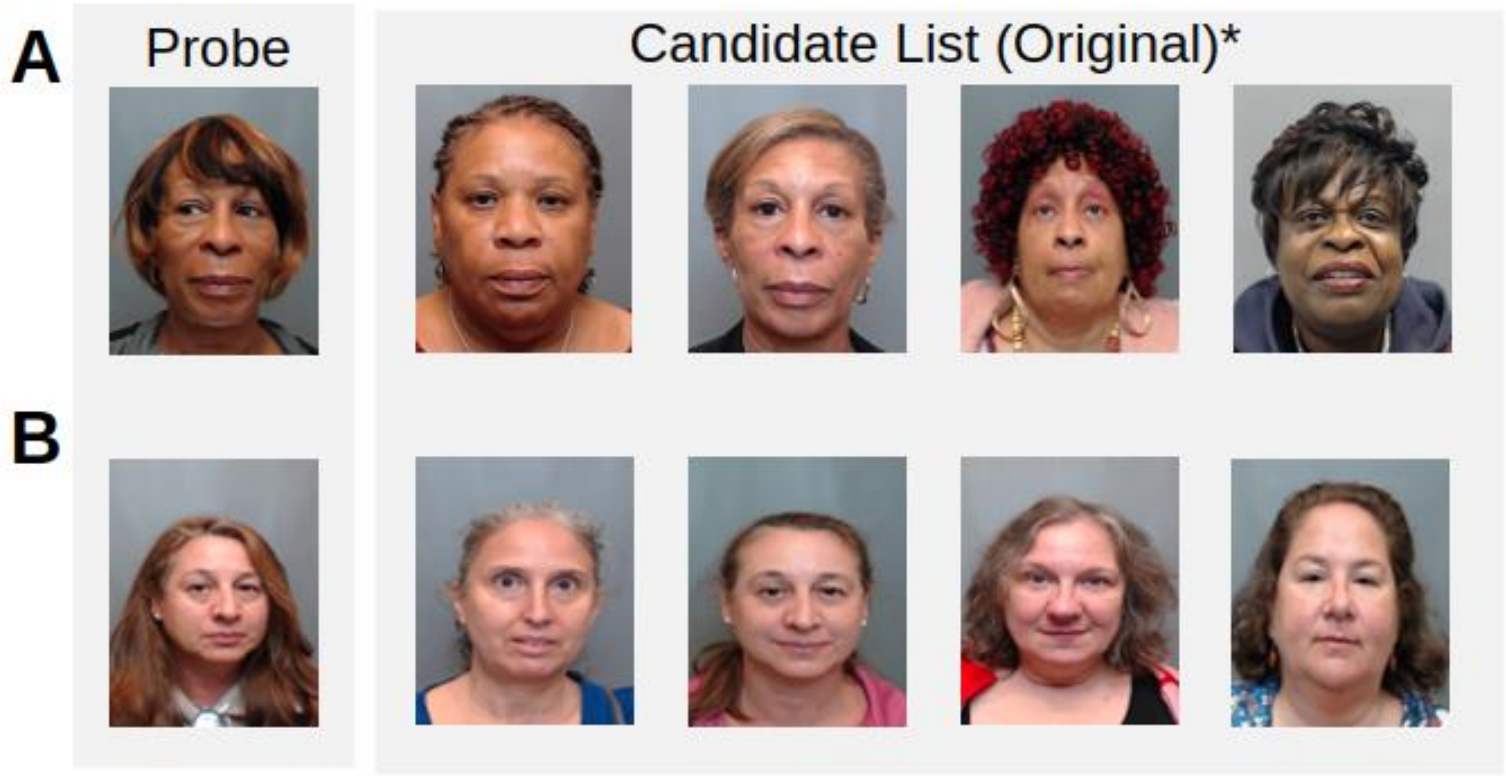

What does this do to human review?

Before De-clustering

After De-clustering

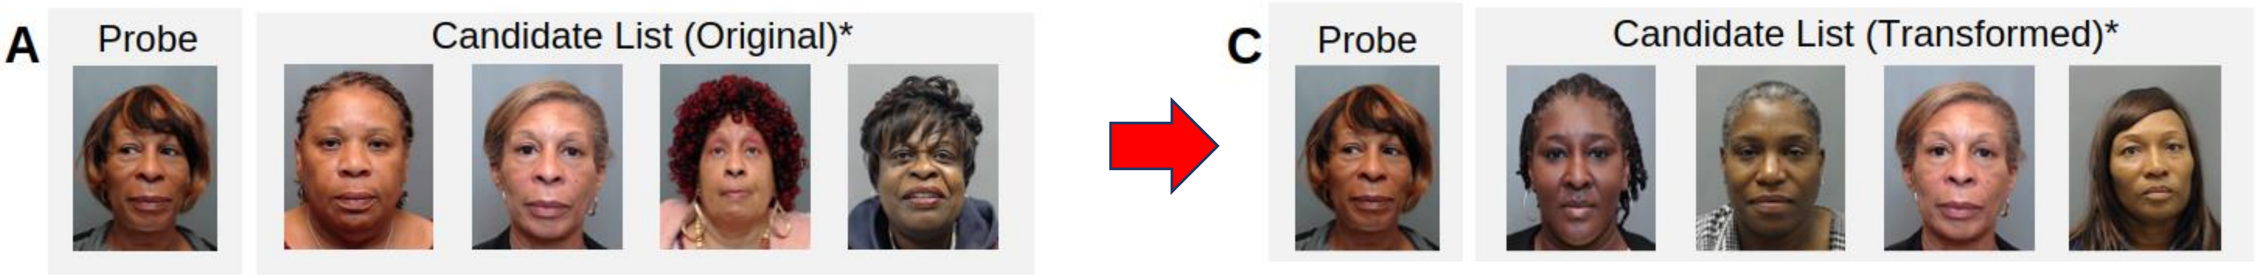

For some subjects, one broadly homogenous candidate set was replaced with another

What does this do to human review?

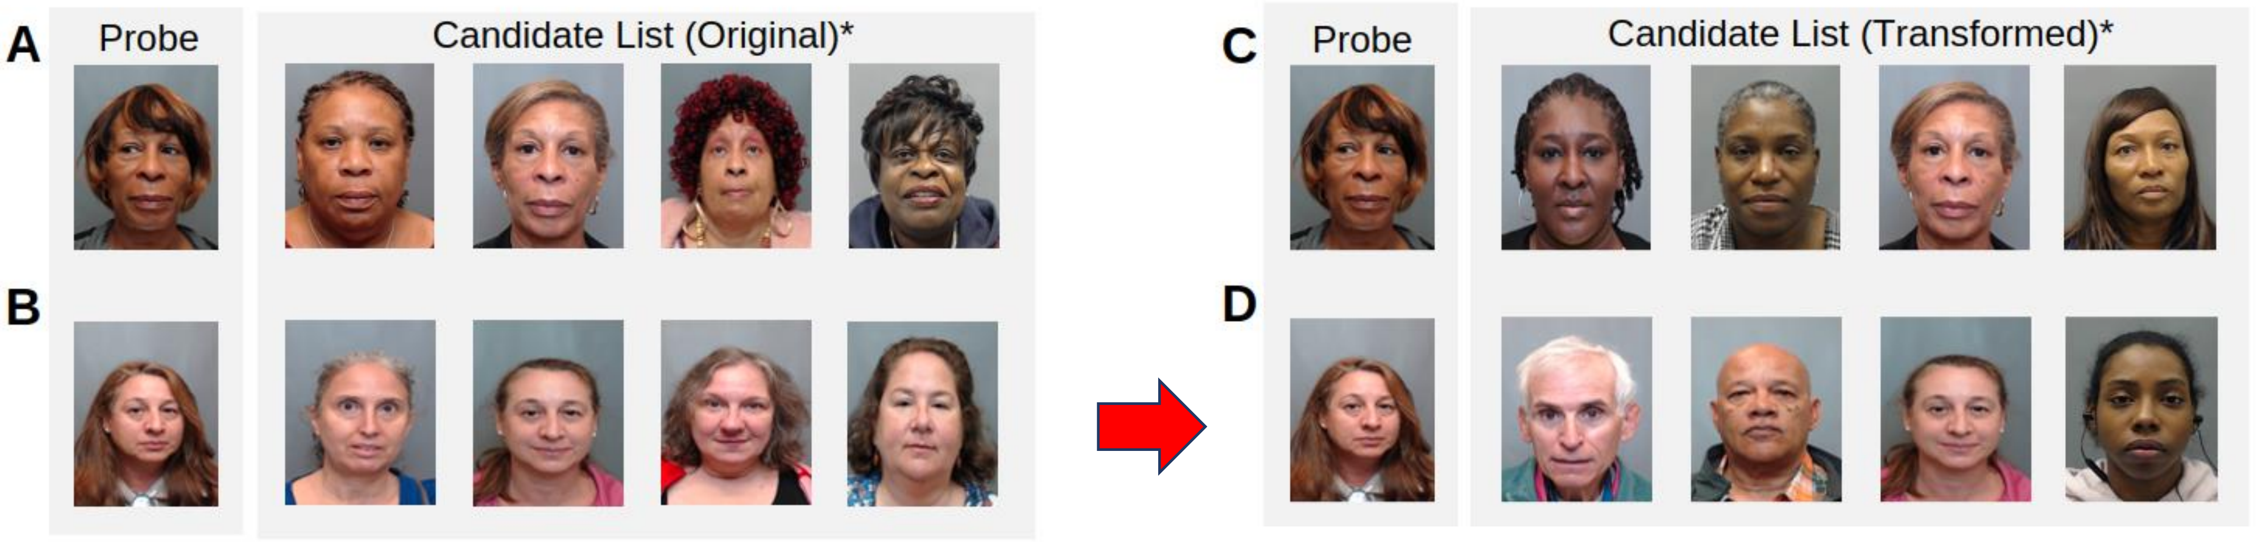

But for others, a homogenous set was replaced with a non-homogenous one

Current literature on face matching in humans work suggest these are much easier for humans to review

Future Work

- What is the best metric for results? Need something beyond false match rate.

- What is the best means to identify and remove “clustering” in feature vector space?

- How stable are these transforms across and within demographic group? Can they be made more stable?

- What is the best algorithm for a human to work with? Might not be “the best algorithm”

In Summary

- Testing face recognition algorithms for demographic effects is important
- The way we understand and measure these effects continues to evolve (because we are testing)
- “Bias” is multifaceted – comes from data, algorithmic decisions, interactions of humans with technical systems
- Better understanding will lead to better technical solutions

Questions & Answers

- Contact information
  - arun.vemury@hq.dhs.gov
  - jhoward@idslabs.org
  - peoplescreening@hq.dhs.gov

- Visit our websites for additional information
  - To see additional work DHS S&T supports, visit www.dhs.gov/science-and-technology
  - Detailed application instructions will be available in a separate document on https://mdtf.org
  - To view additional information about this year and prior Rallies, visit https://mdtf.org