Computer Vision and Society

As a member of the public

- Lots of news about computer vision
- Most coverage by journalists not necessarily well-versed
- Things to watch out for:
  - Claims of doing the impossible (“enhance!”)
  - Articles based on single papers
  - Unfounded comparisons to the human brain

The ethics of computer vision

- Does the application align with your values?
- Does the task specification / evaluation metric reflect the things you care about?
- For recognition:
  - Does the collected training / test set match your true distribution?
- Are the algorithm's errors biased?
- Are you being honest in public descriptions of your results?
- Is the accuracy/correctness sufficient for public release?

An example

- Task: Given image of workplace predict occupation

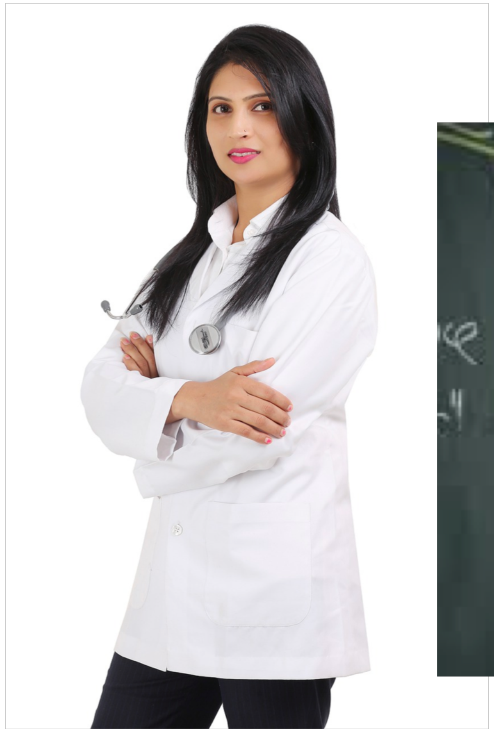

An example

Data	Gender A	Gender B
Occupation A	25%	75%
Occupation B	75%	25%

Model A	Gender A	Gender B
Occupation A	70%	70%
Occupation B	70%	70%

Model B	Gender A	Gender B
Occupation A	0%	100%
Occupation B	100%	0%

Case study: face recognition

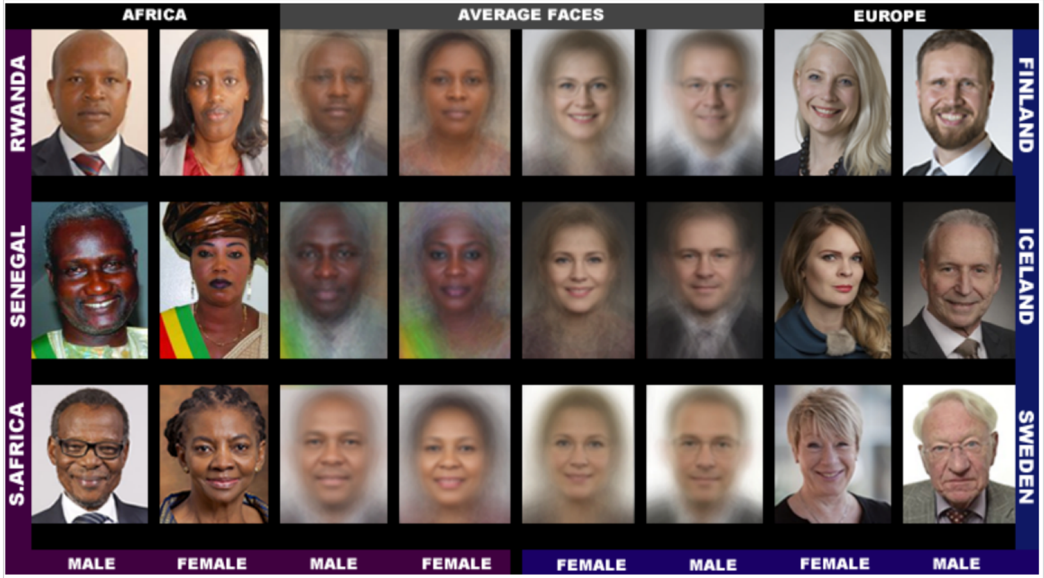

Figure 1: Example images and average faces from the new Pilot Parliaments Benchmark (PPB). As the examples show, the images are constrained with relatively little variation in pose. The subjects are composed of male and female parliamentarians from 6 countries. On average, Senegalese subjects are the darkest skinned while those from Finland and Iceland are the lightest skinned.

Classifier	Metric	All	F	M	Darker	Lighter	DF	DM	LF	LM
MSFT	PPV(%)	93.7	89.3	97.4	87.1	99.3	79.2	94.0	98.3	**100**
	Error Rate(%)	6.3	10.7	2.6	12.9	0.7	**20.8**	6.0	1.7	0.0
	TPR (%)	93.7	96.5	91.7	87.1	99.3	92.1	83.7	**100**	98.7
	FPR (%)	6.3	8.3	3.5	12.9	0.7	**16.3**	7.9	1.3	0.0
Face++	PPV(%)	90.0	78.7	99.3	83.5	95.3	65.5	**99.3**	94.0	99.2
	Error Rate(%)	10.0	21.3	0.7	16.5	4.7	**34.5**	0.7	6.0	0.8
	TPR (%)	90.0	98.9	85.1	83.5	95.3	98.8	76.6	**98.9**	92.9
	FPR (%)	10.0	14.9	1.1	16.5	4.7	**23.4**	1.2	7.1	1.1
IBM	PPV(%)	87.9	79.7	94.4	77.6	96.8	65.3	88.0	92.9	**99.7**
	Error Rate(%)	12.1	20.3	5.6	22.4	3.2	**34.7**	12.0	7.1	0.3
	TPR (%)	87.9	92.1	85.2	77.6	96.8	82.3	74.8	**99.6**	94.8
	FPR (%)	12.1	14.8	7.9	22.4	3.2	**25.2**	17.7	5.20	0.4

Buolamwini, Joy, and Timnit Gebru. "Gender shades: Intersectional accuracy disparities in commercial gender classification." *Conference on fairness, accountability and transparency*. 2018.

Case study

Public | 13 | ...

Deep neural networks are more accurate than humans at detecting sexual orientation from facial images.

Contributors: Yilun Wang, Michal Kosinski

Date created: 2017-02-15 11:37 AM | Last Updated: 2018-10-23 05:05 PM

Category: Project

Description: We show that faces contain much more information about sexual orientation than can be perceived and interpreted by the human brain. We used deep neural networks to extract features from 35,326 facial images. These features were entered into a logistic regression aimed at classifying sexual orientation. Given a single facial image, a classifier could correctly distinguish between gay and heterosexual men in 81% of cases, and in 74% of cases for women. Human judges achieved much lower accuracy: 61% for men and 54% for women. The accuracy of the algorithm increased to 91% and 83%, respectively, given five facial images per person. Facial features employed by the classifier included both fixed (e.g., nose shape) and transient facial features (e.g., grooming style). Consistent with the prenatal hormone theory of sexual orientation, gay men and women tended to have gender-atypical facial morphology, expression, and grooming styles. Prediction models aimed at gender alone allowed for detecting gay males with 57% accuracy and gay females with 58% accuracy. Those findings advance our understanding of the origins of sexual orientation and the limits of human perception. Additionally, given that companies and governments are increasingly using computer vision algorithms to detect people's intimate traits, our findings expose a threat to the privacy and safety of gay men and women.

Case study

We show that faces contain much more information about sexual orientation than can be perceived and interpreted by the human brain. … Given a single facial image, a classifier could correctly distinguish between gay and heterosexual men in 81% of cases, and in 74% of cases for women. … Consistent with the prenatal hormone theory of sexual orientation, gay men and women tended to have gender-atypical facial morphology, expression, and grooming styles … our findings expose a threat to the privacy and safety of gay men and women.

Some questions

- Does the application align with your values?
- Does the task specification / evaluation metric reflect the things you care about?
- For recognition:
  - Does the collected training / test set match your true distribution?
- Are the algorithm's errors biased?
- Are you being honest in public descriptions of your results?
- Is the accuracy/correctness sufficient for public release?

Some answers

- Training and testing set?
  - 35,326 images from public profiles on a US dating website
- ”average” images of straight/gay people:

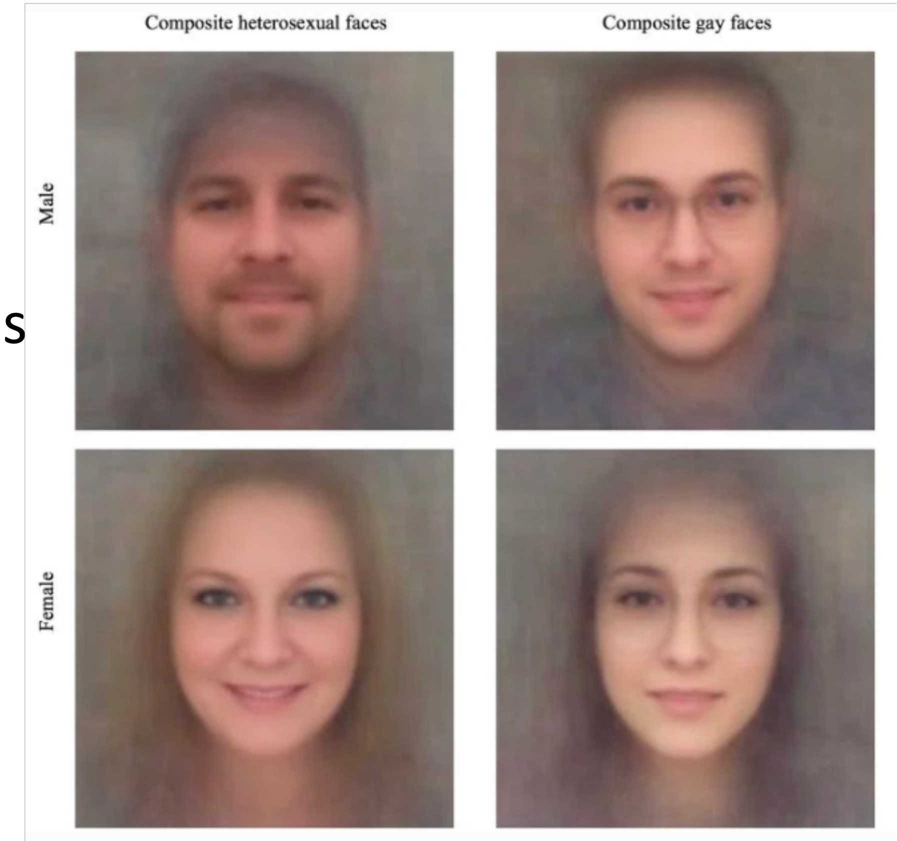

- Question:
  - Are differences caused by actual differences in faces
  - Or how people choose to present themselves in dating websites?

https://medium.com/@blaisea/do-algorithms-reveal-sexual-orientation-or-just-expose-our-stereotypes-d998fafdf477

Some answers

- Men on average taller than women
- Straight men take selfies from slightly below, straight women from slightly above
- Profiles in dating websites reflect cultural norms
  - Straight women more likely to wear eyeshadow/make-up
  - Straight men more likely to keep facial hair
- Physiology of sexual orientation more nuanced
- Hard to say what exactly neural network was detecting
- Goal: raise privacy concerns. Side-effects?
  - Reinforces potentially harmful stereotypes
  - Provides ostensibly “objective” criteria for discrimination