


The Limits and Potentials of Deep Learning for Facial Analysis


Walter J. Scheirer

Computer Vision Research Laboratory
Department of Computer Science and Engineering



Biometrics isn't the only field
studying faces...

Imagine the following scenario:




A proprietary face recognition system purportedly solves biometric identification with human-like ability.

By all accounts, the software achieves superior performance on established computer vision benchmarks

How would you go about falsifying the claim of human-like ability?

The Chinese Room




Visual Psychophysics

Probe psychological and perceptual thresholds through controlled manipulation of stimuli.

Careful management of stimulus construction, ordering and presentation allows for precise determination of perceptual thresholds.




Garrido et al 2011



Goren et al 1975


Visual Psychophysics

Canonical Early Example: minimum threshold for stimulation of an individual photoreceptor.

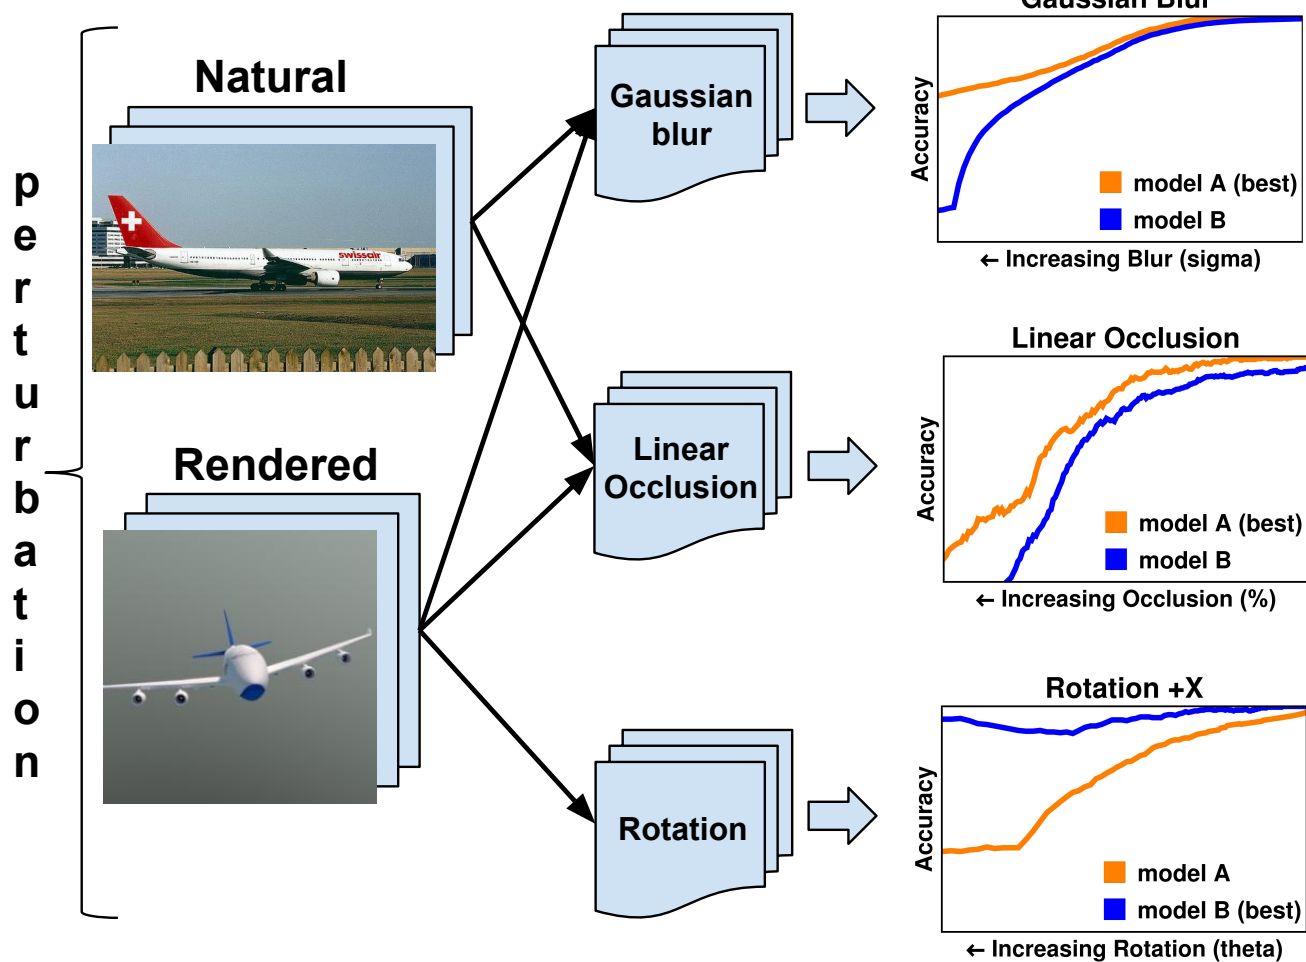


Photoreceptor cell © BY-SA 4.0 Manu5

(a) Traditional



(b) Visual Psychophysics



The limits of deep learning...



Brandon Richard Webster

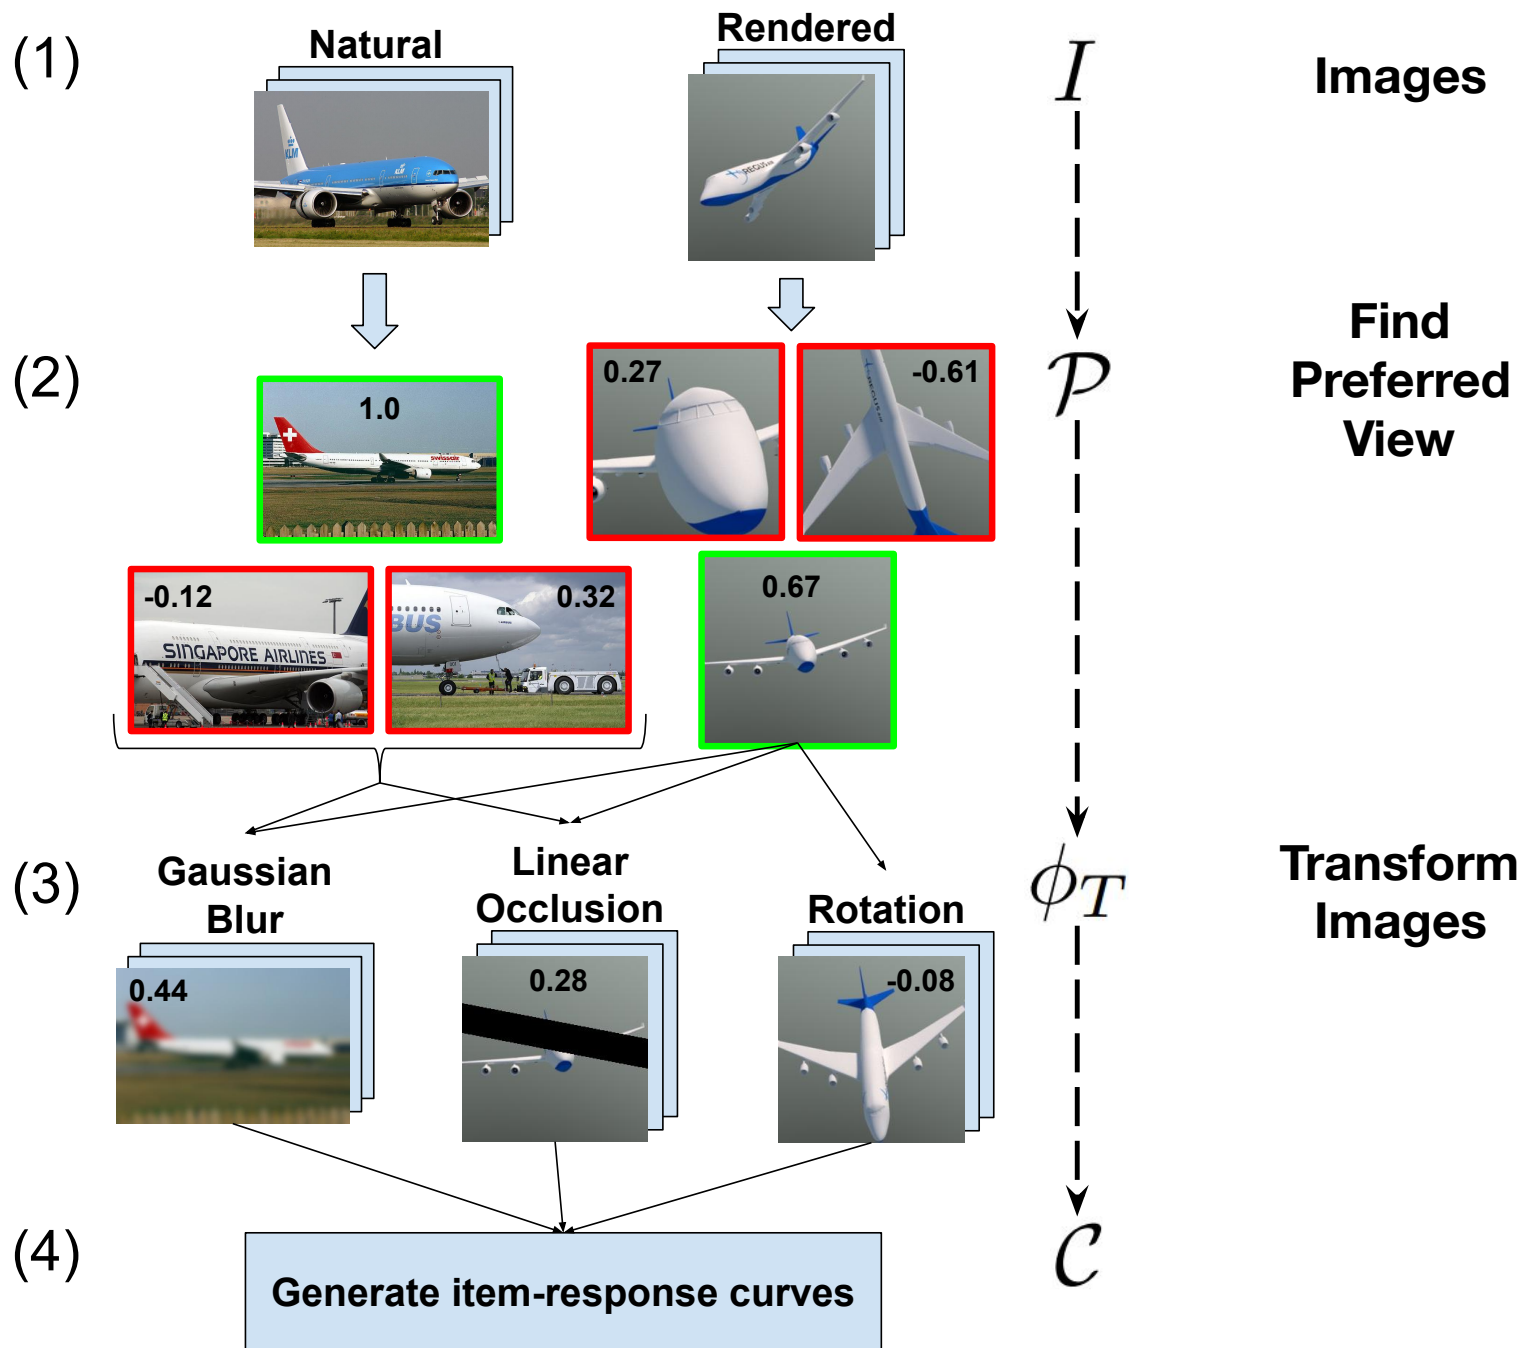


Sam Anthony

Visual Psychophysics for Object Recognition

<https://arxiv.org/abs/1611.06448>

IEEE T-PAMI, September 2019




Two-Alternative Forced Choice (2AFC) Task




Matching Alternate Stimulus



Non-Matching Alternate Stimulus





MAFC Task





Probe Image

Classification Model



A Curious Contrast Deficit





Brandon Richard Webster



Alli Kwon



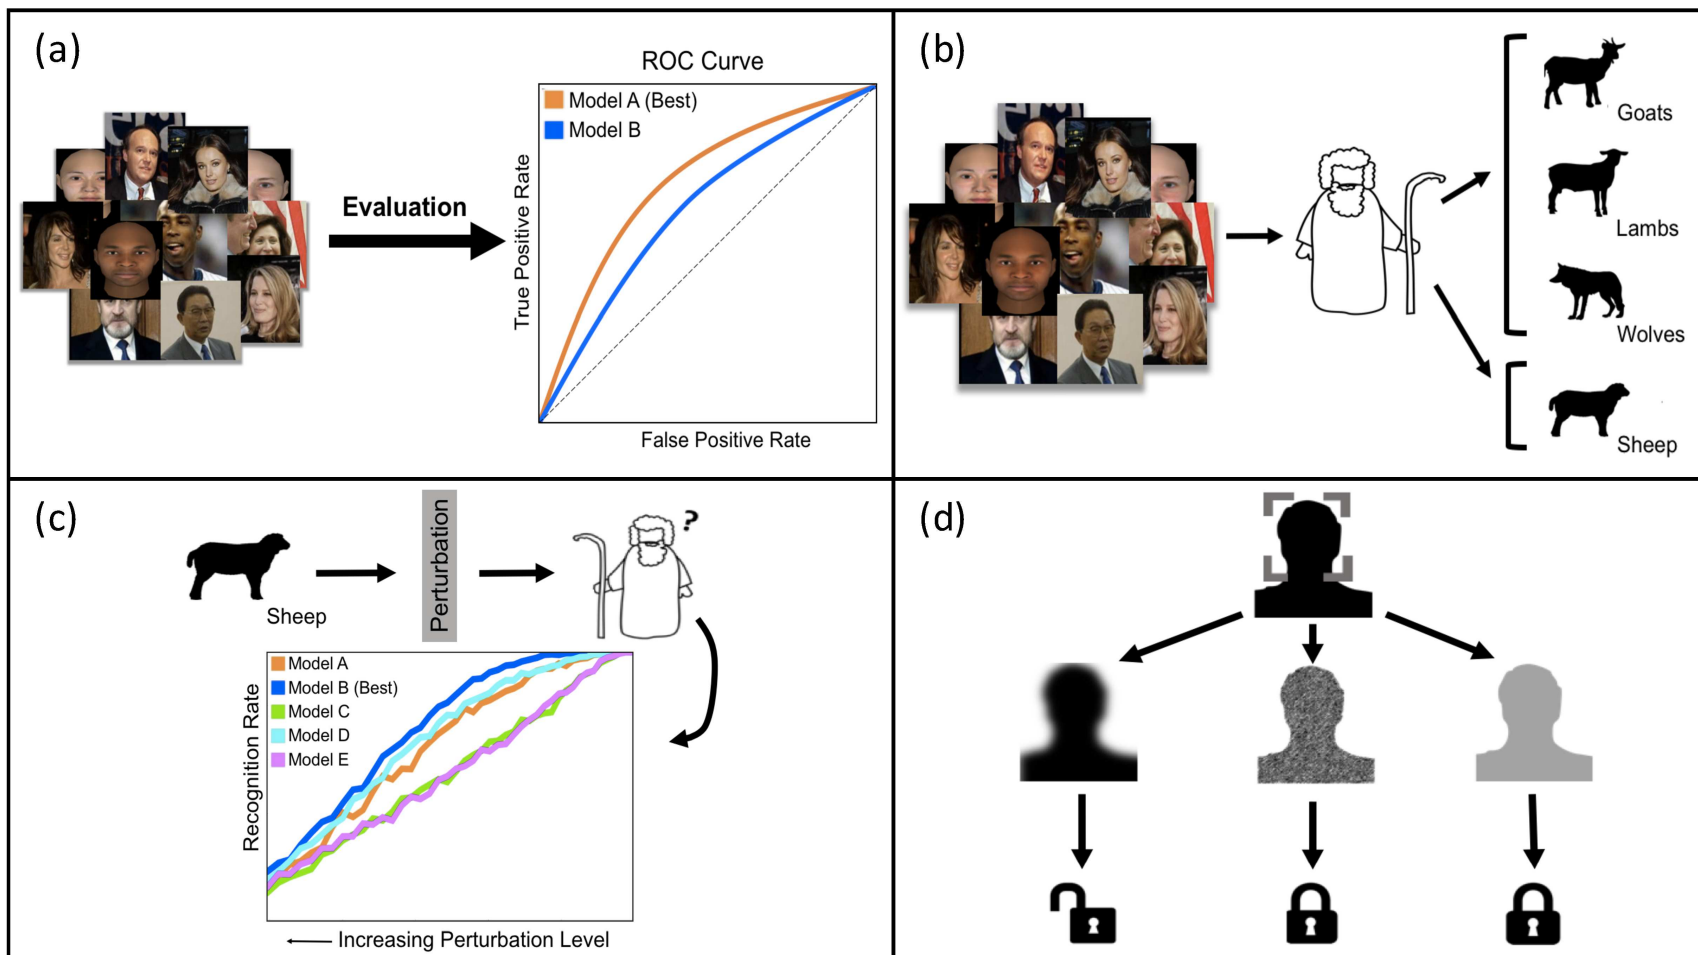
Sam Anthony


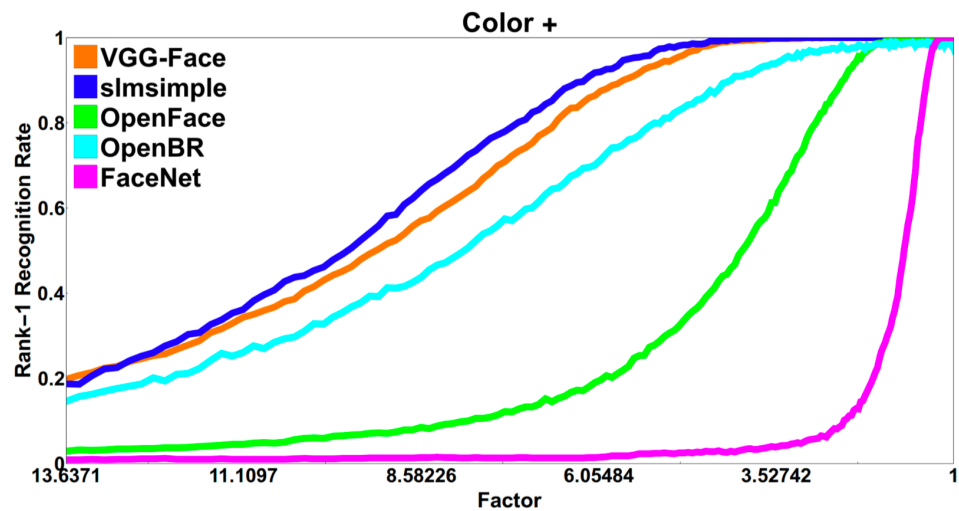
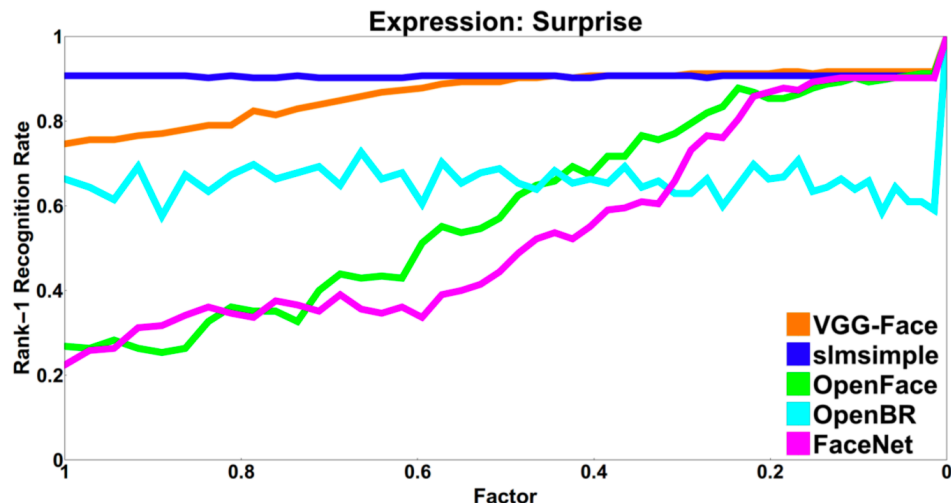

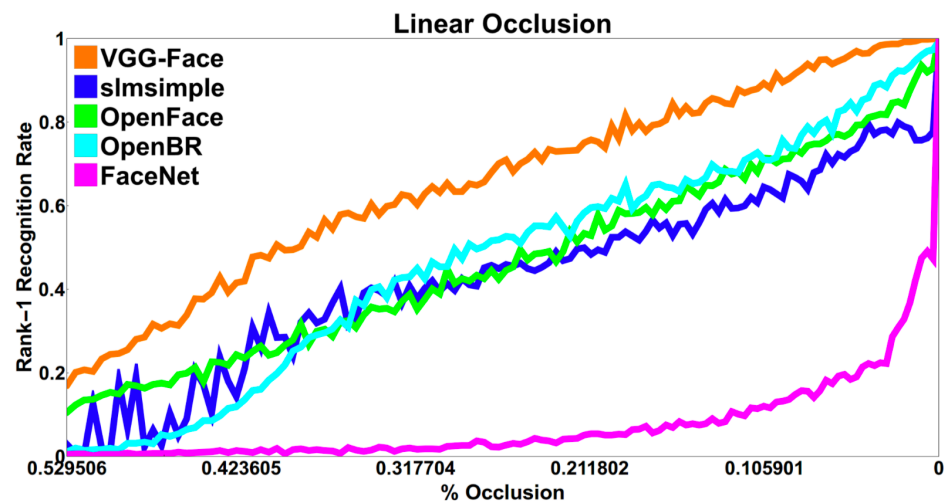
Visual Psychophysics for Face Recognition



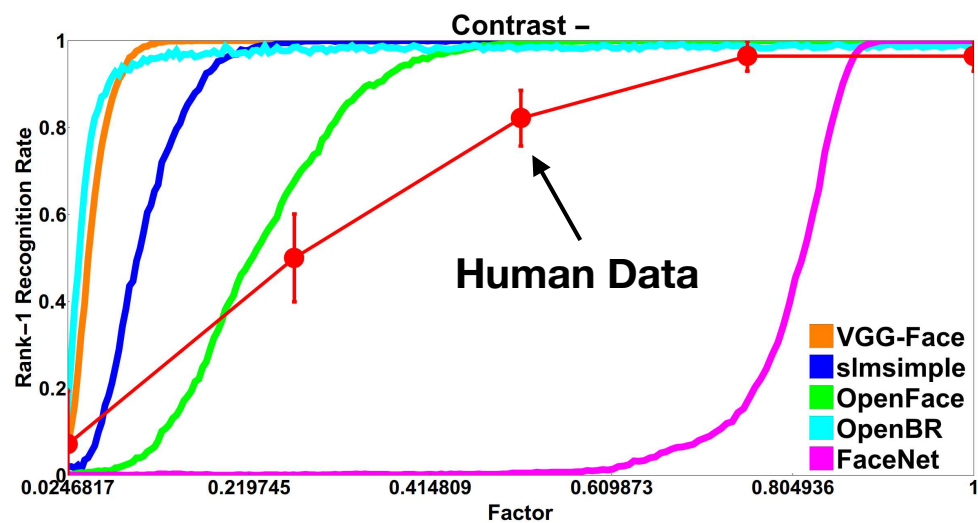

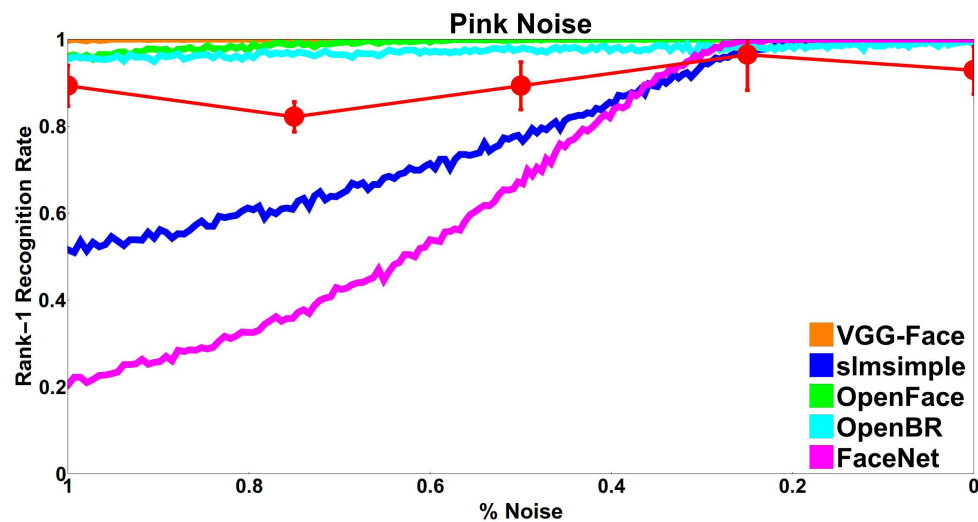
<https://arxiv.org/abs/1803.07140>

European Conference on Computer Vision, 2018

Instead of finding a preferred view, find “sheep” in a biometrics context







The potentials of deep learning...



Mel McCurrie





Sam Anthony

CNNs for Subjective Face Attributes

<https://arxiv.org/abs/1610.08119>

Image and Vision Computing, October 2018

Subjective Attribute Predictions




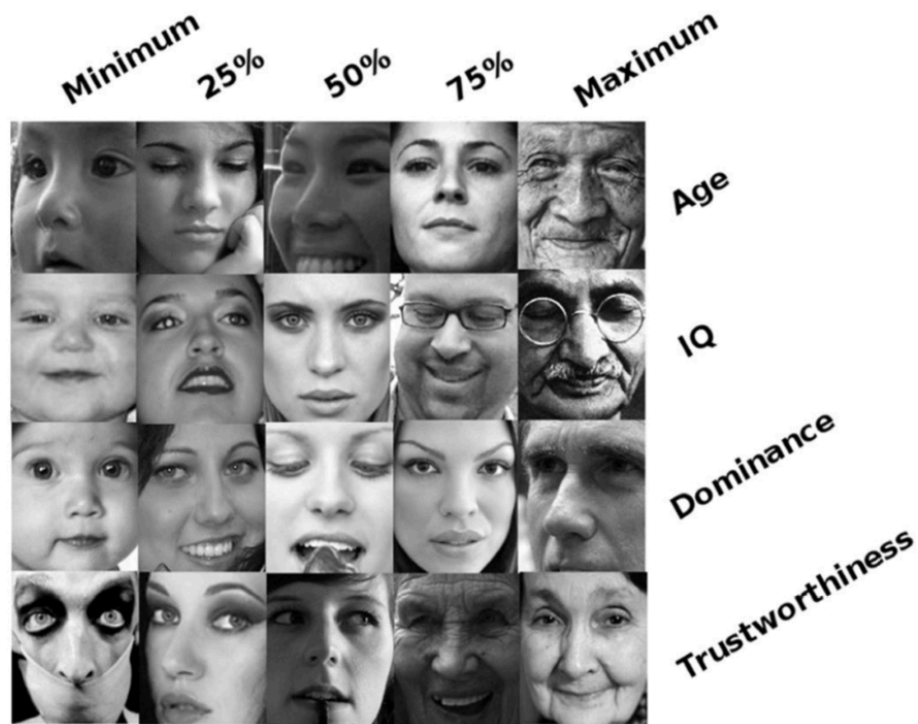
Data collection via TestMyBrain.org

Click one of the buttons below to rate this face from 1 to 7,
where 1 is the least **DOMINANT** and 7 is the most.



least ☐ 1 ☐ 2 ☐ 3 ☐ 4 ☐ 5 ☐ 6 ☐ 7 most

Data Distributions

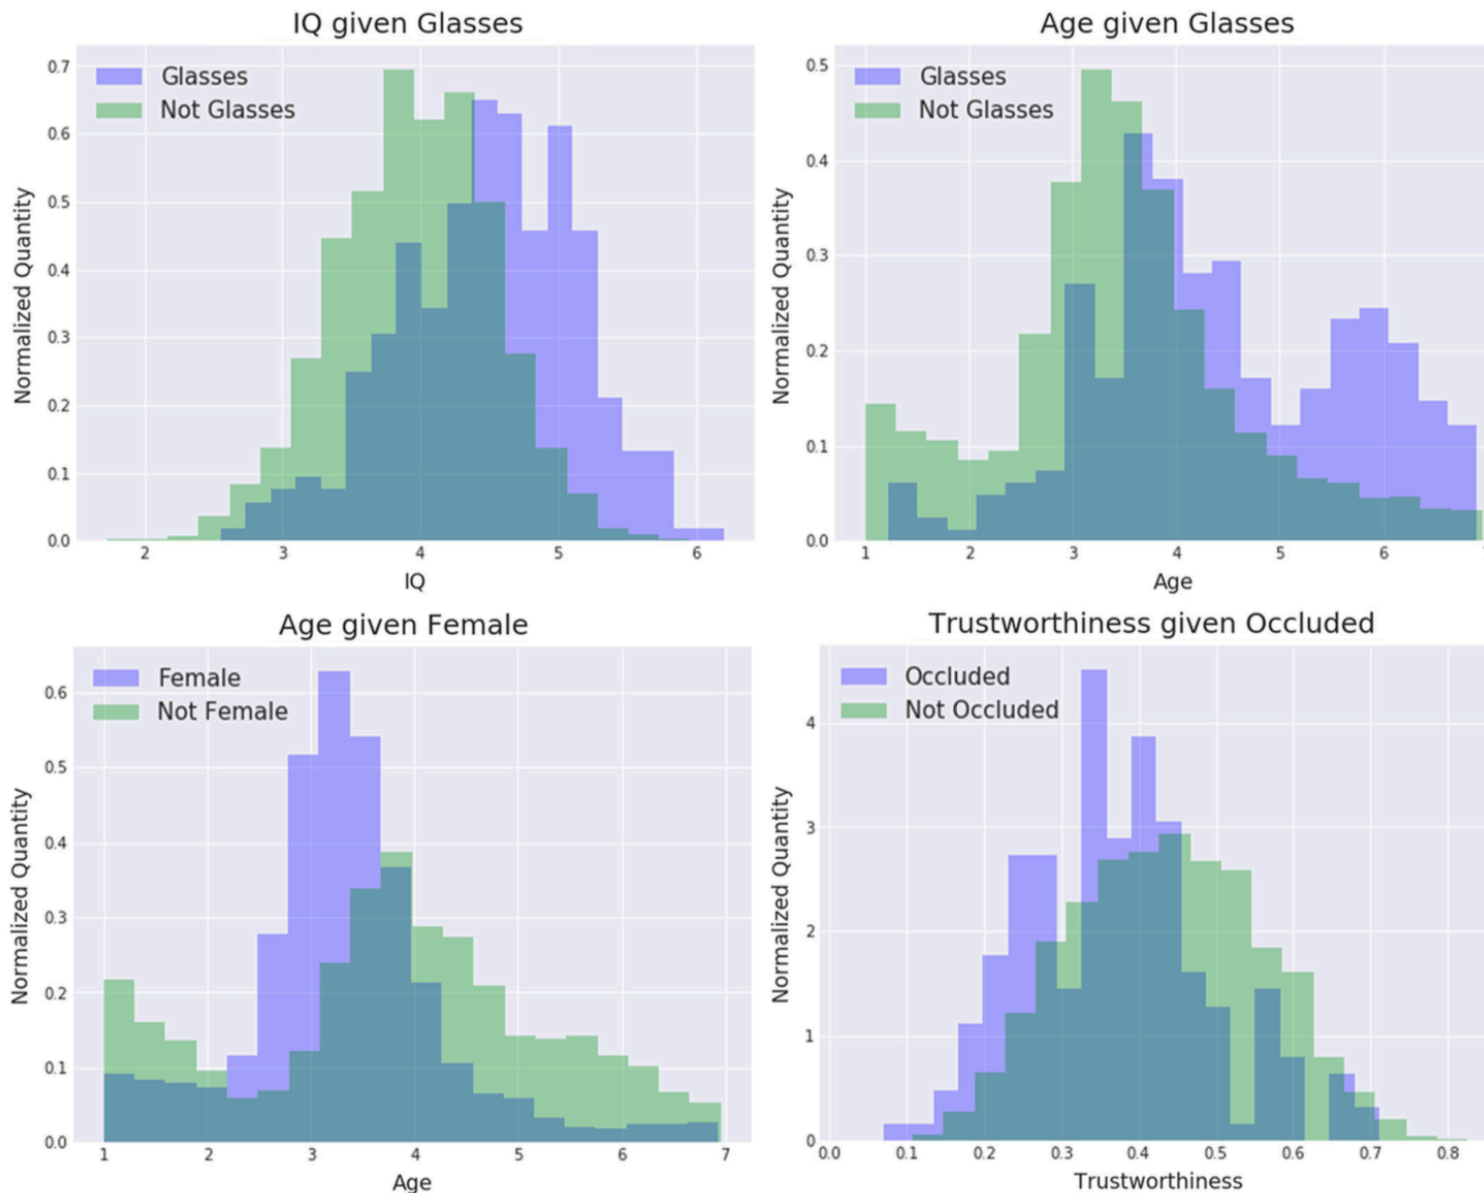


Dataset Statistics

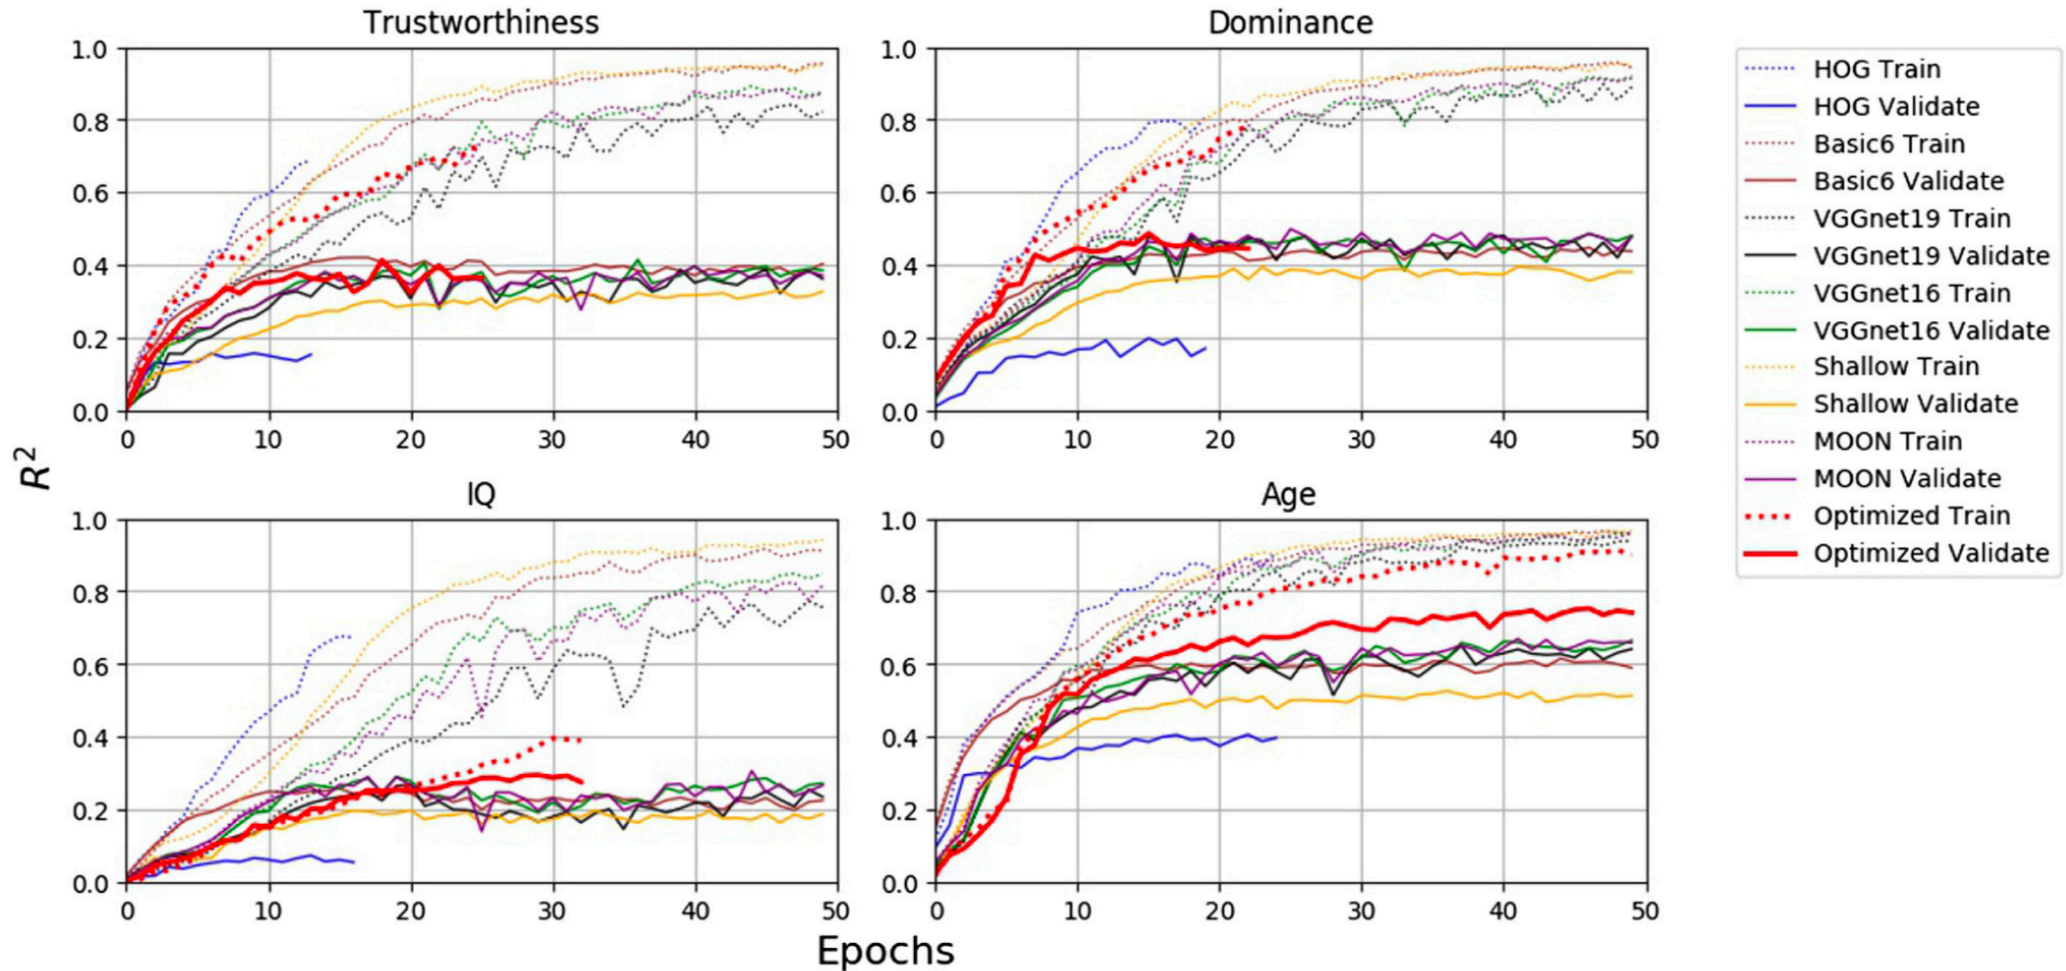
5040 Total Images

	Trust.	Dom.	Age	IQ
Mean of ratings	0.48	0.47	0.42	0.48
Std. of ratings	0.16	0.16	.20	0.14
Mean Std. of ratings	0.34	0.32	0.13	0.27
Mean Num. of ratings	32.47	32.19	15.80	15.79

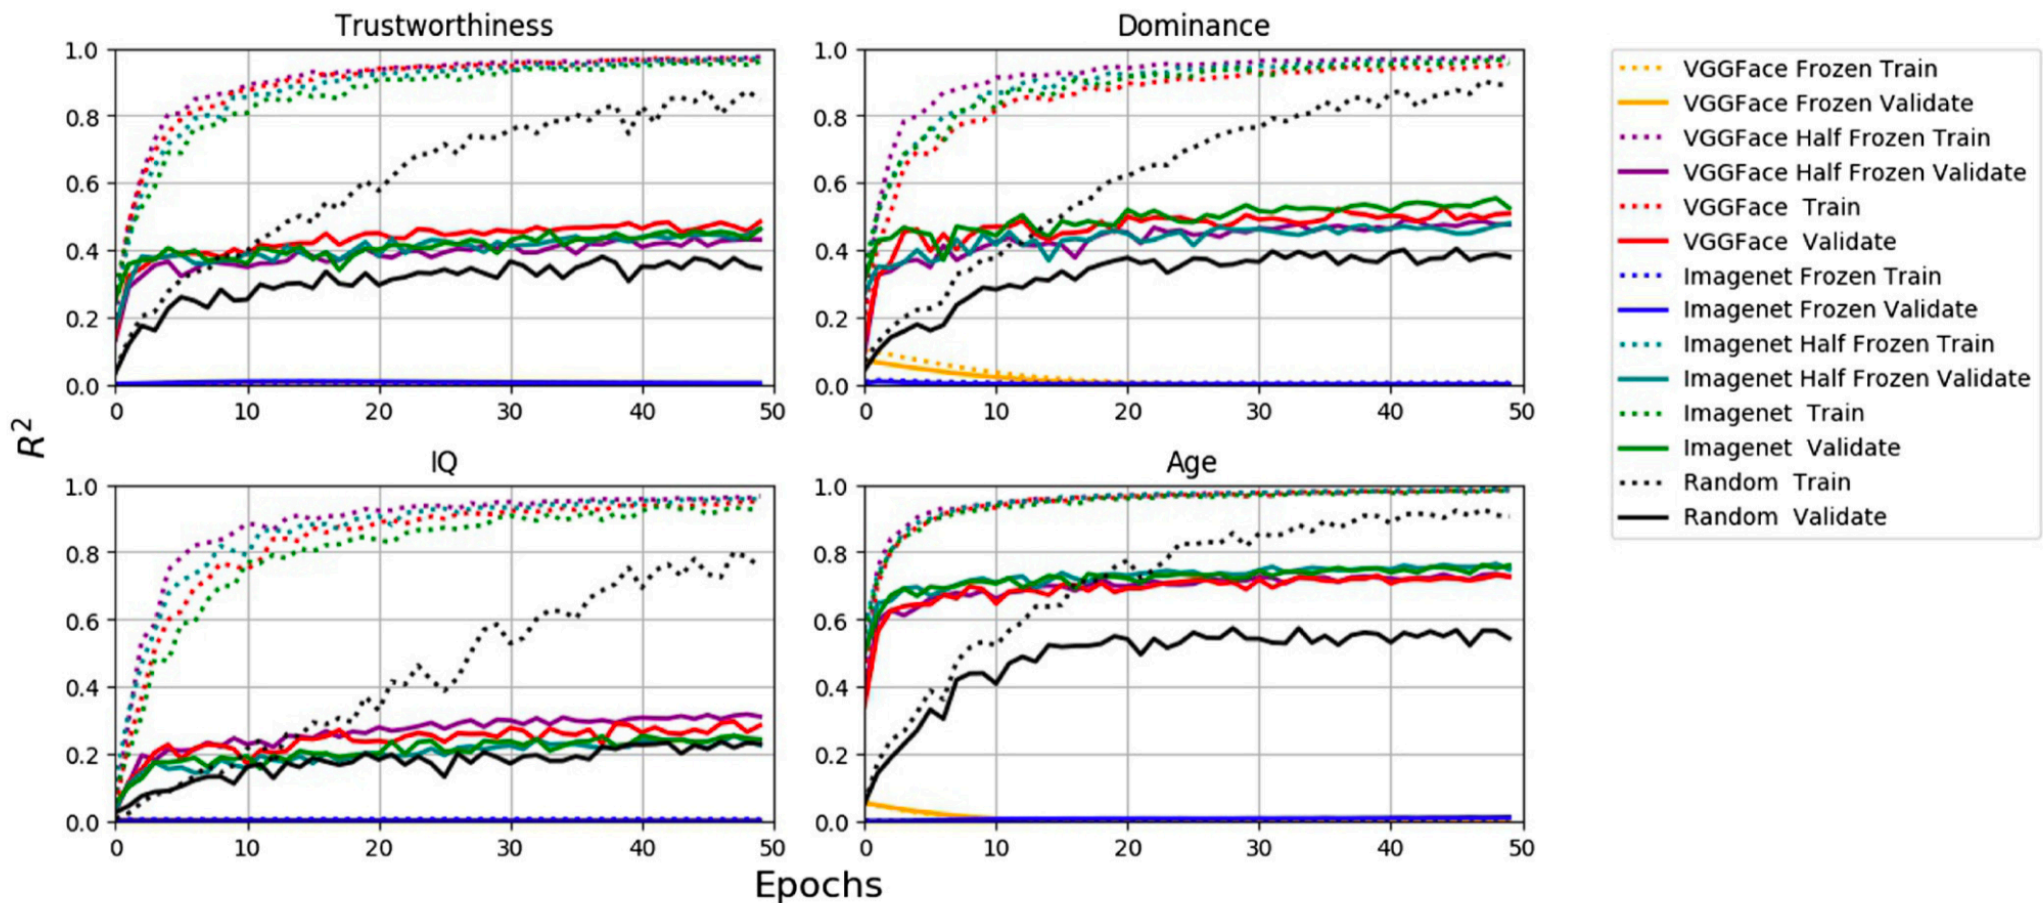
Dataset Patterns




Which architecture is the best?




The dataset is small. Does transfer learning help?




Regions of the face most important to the models (transfer learning)



Regions of the face most important to the models (saliency from occlusion)



Trustworthiness



What other attributes are interesting?



Thank you!