Social Networks:
Analyzing Social Information in Deep Convolutional Neural Networks Trained for Face Identification

Connor J. Parde[1], Ying Hu[1], Carlos Castillo[2], Swami Sankaranarayanan[2], and Alice J. O'Toole[1]

[1]The University of Texas at Dallas, [2]University of Maryland

Introduction & Goals

Social Traits

- Humans make social trait inferences from faces readily [1] and rapidly [2]
- Trait inferences predict important decisions (e.g., voting preferences) [3, 4]
- Social traits can be generated from models of face structure and reflectance [5, 6]

DCNNs for Face Identification

- State-of-the-art for face identification [7] and generalize over viewpoint, illumination, etc.
- "**Top-level**" DCNNs features retain non-identity information (e.g., pose, image quality) [8]
- Do face-identification features ***also*** retain social trait information?

Goal 1: Measure similarity between human and computer trait predictions made from identity-trained DCNNs

Goal 2: Measure accuracy of trait predictions using DCNN features from non-frontal images

Goal 3: Predict *individual* social trait inferences from top-level DCNN features

Social Trait Ratings

Human ratings of social traits for faces

- 280 face images
- 18 traits from Big Five Factors of Personality [9]
- 20 sets of ratings per face
- responses averaged across participants

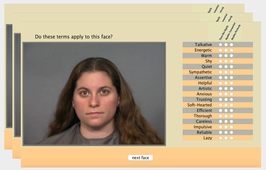

Participants:

- *n* = 80 (60 female)
- *Mean age* = 21

Stimuli:

- 280 images, 194 identities
- 204 female, 76 male
- Caucasian
- neutral expression
- Ratings collected for **front-facing** images

Collected ratings for 18 Social Traits

Talkative	Anxious
Energetic	Trusting
Warm	Soft-hearted
Shy	Efficient
Quiet	Thorough
Sympathetic	Careless
Assertive	Impulsive
Helpful	Reliable
Artistic	Lazy

11 Unique Dimensions

Averaged highly correlated traits:

- talkative, energetic
- warm, sympathetic, soft-hearted, trusting, helpful, reliable
- efficient, thorough

Identity Descriptors

DCNNs modeled after primate visual cortex

- Early layers model V1-V4, final layers model IT cortex
- For face identification, final DCNN layer stores abstract identity code <- ***face representation***

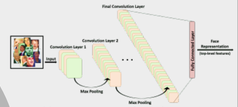

- Network used in this study contains 6 convolutional layers, 3 fully connected layers [9]
- State-of-the-art performance on challenging, unconstrained IJB-A dataset [10]

Deep Network Architecture		
Layer	***Kernel Size/Stride***	***#Parameters***
Conv1	11x11 / 4	35k
Conv2	5x5 / 2	614k
Conv3	3x3 / 2	885k
Conv4	3x3 / 2	1.3M
Conv5	3x3 / 1	885k
Conv6	3x3 / 1	590k
Fc6	1024	9.4M
Fc7	512	524k
Fc8	10575	5.5M
Softmax Loss		19.8M

Verify Structure of Face Trait Space (e.g. [5])

- principal component analysis of human trait ratings
  - created "trait space"
- 2 significant principal components:
  - 1st component interpreted as ***approachability***
  - 2nd component interpreted as ***dominance***

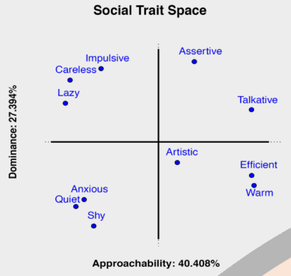

Predict Social Trait Inferences

- *N* x *K* "feature matrix" obtained from DCNN
- *N* x *M* "trait matrix" obtained from averaged participant responses
- Predict trait matrix from feature matrix using linear regression

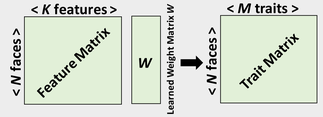

- Remove *L* features (*n* = 140) with low learned weights, re-train model
  - Keep only features important for trait prediction
- Columns in the final trait matrix are computer predictions of columns from original data matrix

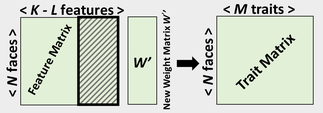

- Similarity between human-generated and computer-predicted trait vectors measured using cosine distance
- Accuracy of individual trait predictions measured using R^2 between human-generated and computer-predicted values

Results

Trait-Profile Predictions

- Cosine similarity between human-generated trait profiles and computer predictions:

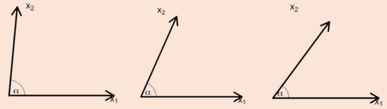

Null distribution: **α** =arccos(**0.078**)

Using *K features*: **α** = arccos(**0.353**)

Using *K – L* features: **α** = arccos(**0.533**)

- Predictions made from non-frontal DCNN features

Average Cosine Similarity Between Human-Generated Traits and Computer-Predicted Traits		
Direction	***Left-facing Pose***	***Right-facing Pose***
22.5°	**0.4912**	**0.5556**
45°	**0.5578**	**0.5181**
67.5°	**0.5153**	**0.5390**
90°	**0.5173**	**0.5411**

Individual Trait Predictions

- Error between human ratings and predicted traits, plotted against a null distribution
  - All traits predicted significantly above chance
  - Blue line: α = 0.002
  - Red line: predicted value

Trait-Prediction Error

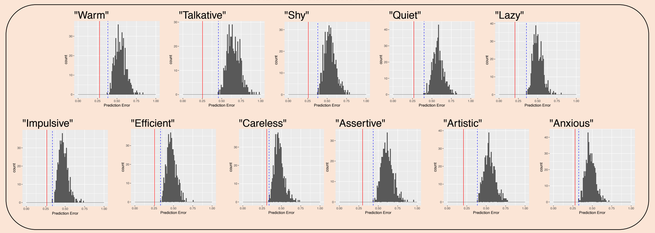

R^2 Between Human Inferences and Computer Predictions

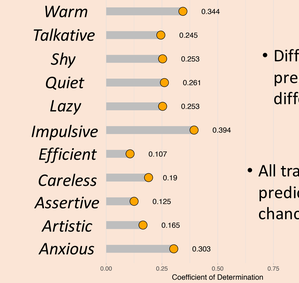

- Different traits predicted to different extents
- All trait inferences predicted above chance

Conclusions

Conclusion 1

Human trait inferences can be predicted from the top-level features of a DCNN trained for face identification

Conclusion 2

Trait inferences assigned to frontal faces can be predicted from DCNN features generated for both frontal and non-frontal faces

Conclusion 3

Top-level DCNN features for face identification retain robust trait representation – each individual trait predicted above chance

- DCNN representation allows for state-of-the-art identification
- Not independent of image information, social traits

References

[1] Bruce, V., & Young, A. (1986). Understanding face recognition. *British journal of psychology, 77*(3), 305-327.

[2] Bar, M., Neta, M., & Linz, H. (2006). Very first impressions. *Emotion, 6*(2), 269.

[3] Todorov, A., Mandisodza, A. N., Goren, A., & Hall, C. C. (2005). Inferences of competence from faces predict election outcomes. *Science, 308*(5728), 1623-1626.

[4] Rule, N. O. Ambady, N. (2008). The face of success: Inferences from chief executive officers' appearance predict company profits. *Psychological Science: A Journal of the American Psychological Society/APS*, 19, 109–111.

[5] Oosterhof, N. N., & Todorov, A. (2008). The functional basis of face evaluation. *Proceedings of the National Academy of Sciences, 105*(32), 11087-11092.

[6] Walker, M., & Vetter, T. (2009). Portraits made to measure: Manipulating social judgments about individuals with a statistical face model. *Journal of Vision, 9*(11), 12-12.

[7] Taigman, Y., Yang, M., Ranzato, M. A., & Wolf, L. (2014). Deepface: Closing the gap to human-level performance in face verification. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 1701-1708).

[8] Parde, C. J., Castillo, C., Hill, M. Q., Colon, Y. I., Sankaranarayanan, S., Chen, J. C., & O'Toole, A. J. (2017, May). Face and Image Representation in Deep CNN Features. In *Automatic Face & Gesture Recognition (FG 2017), 2017 12th IEEE International Conference on* (pp. 673-680). IEEE.

[9] Gosling, S. D., Rentfrow, P. J., & Swann Jr, W. B. (2003). A very brief measure of the Big-Five personality domains. *Journal of Research in personality, 37*(6), 504-528.

[10] Sankaranarayanan, S., Alavi, A., Castillo, C. D., & Chellappa, R. (2016, September). Triplet probabilistic embedding for face verification and clustering. In *Biometrics Theory, Applications and Systems (BTAS), 2016 IEEE 8th International Conference on* (pp. 1-8). IEEE.

Acknowledgements

This research is based upon work supported by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA), via IARPA R&D Contract No. 2014-14071600012. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the ODNI, IARPA, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon.