

Introduction

- Many IoT “things” are **video cameras**
 - Have IP addresses and are configurable
 - Cheap and integrated into mobile devices.
 - High data rates and bandwidth.

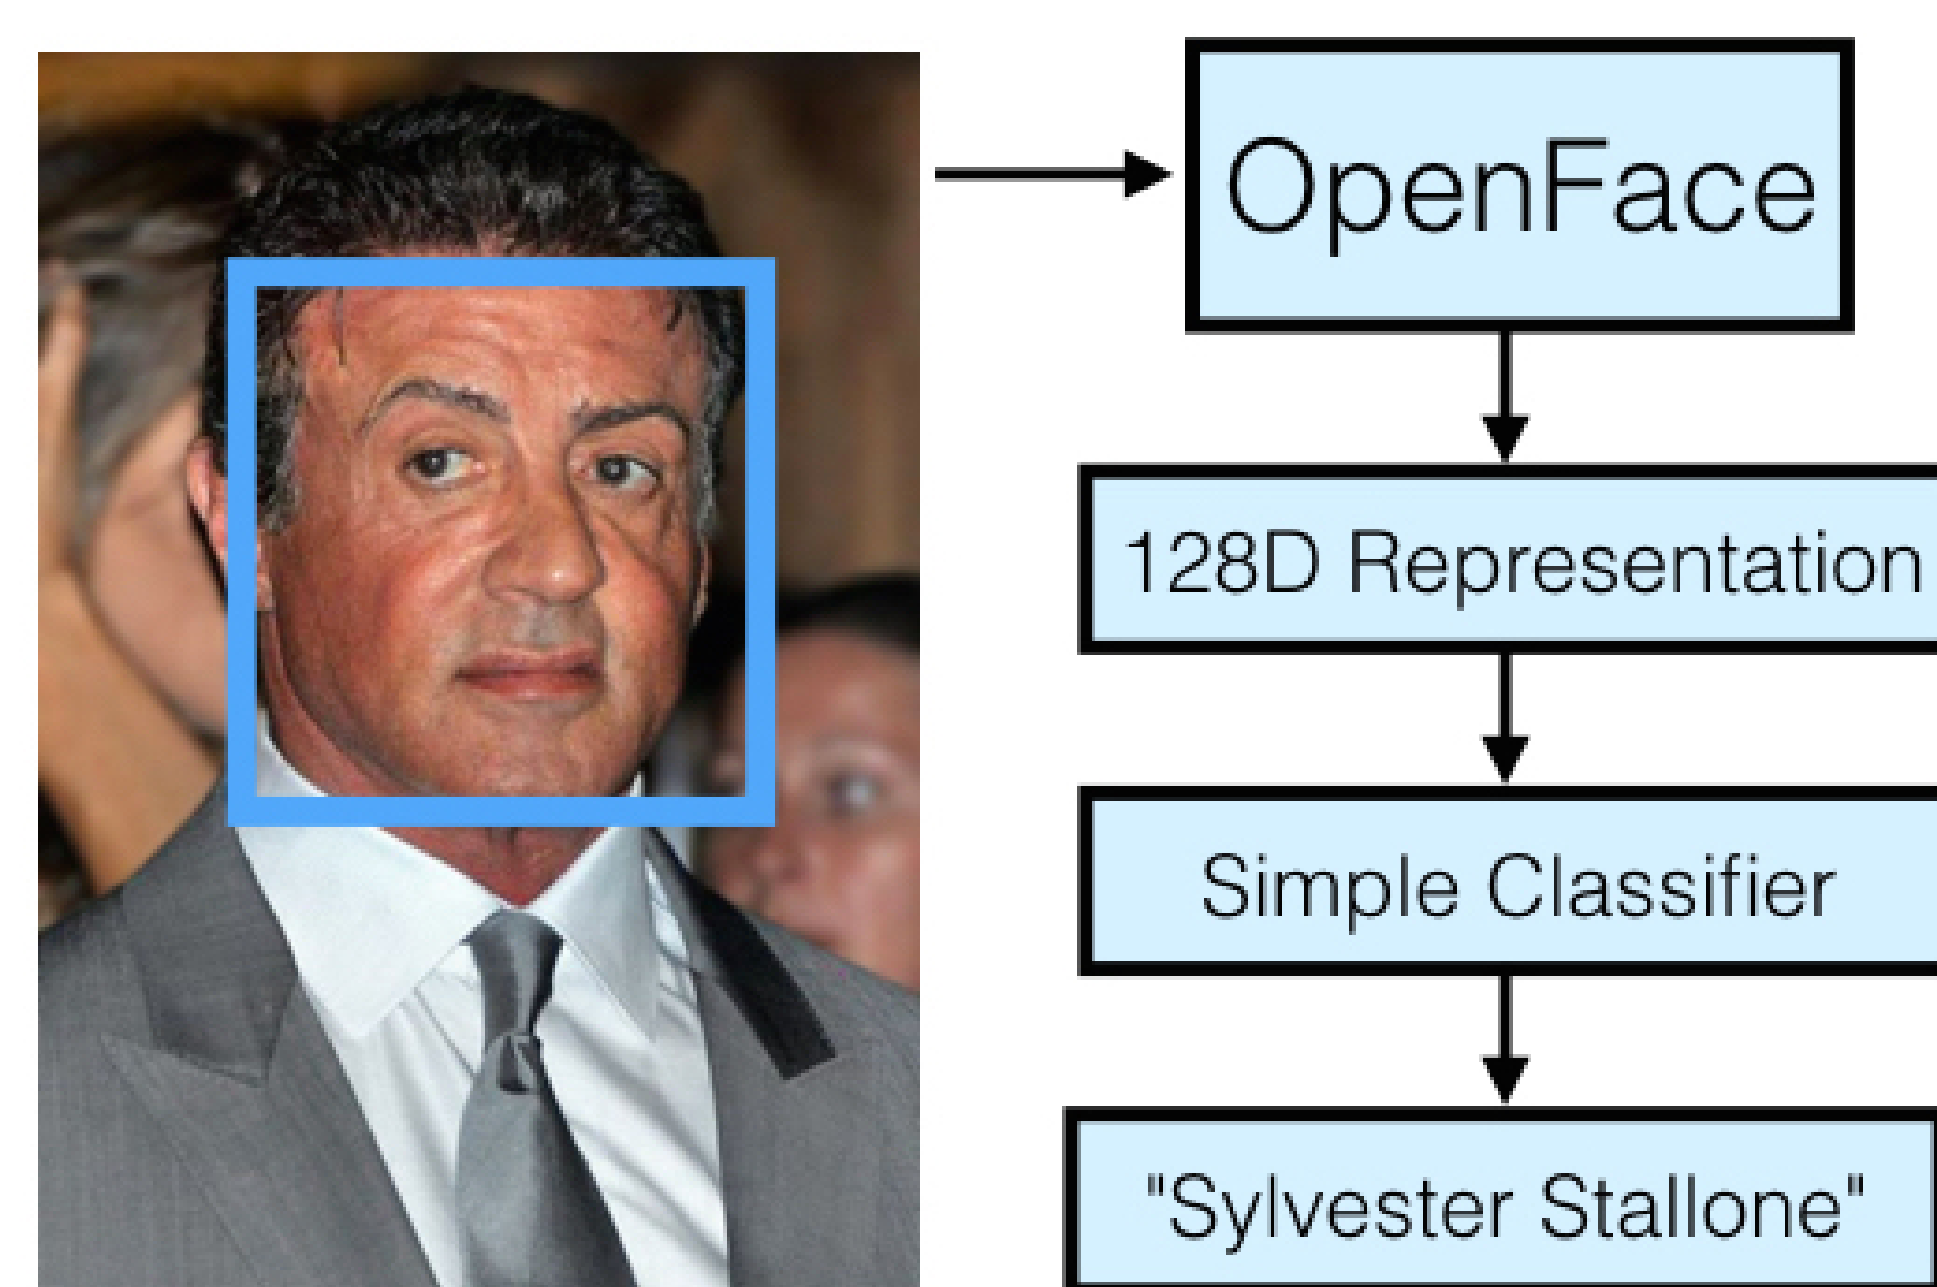
Adding context to video-based IoT systems

- Recognizing people is crucial to human context.
- It modulates what you say and how you act.
- Why not add this context to IoT systems?

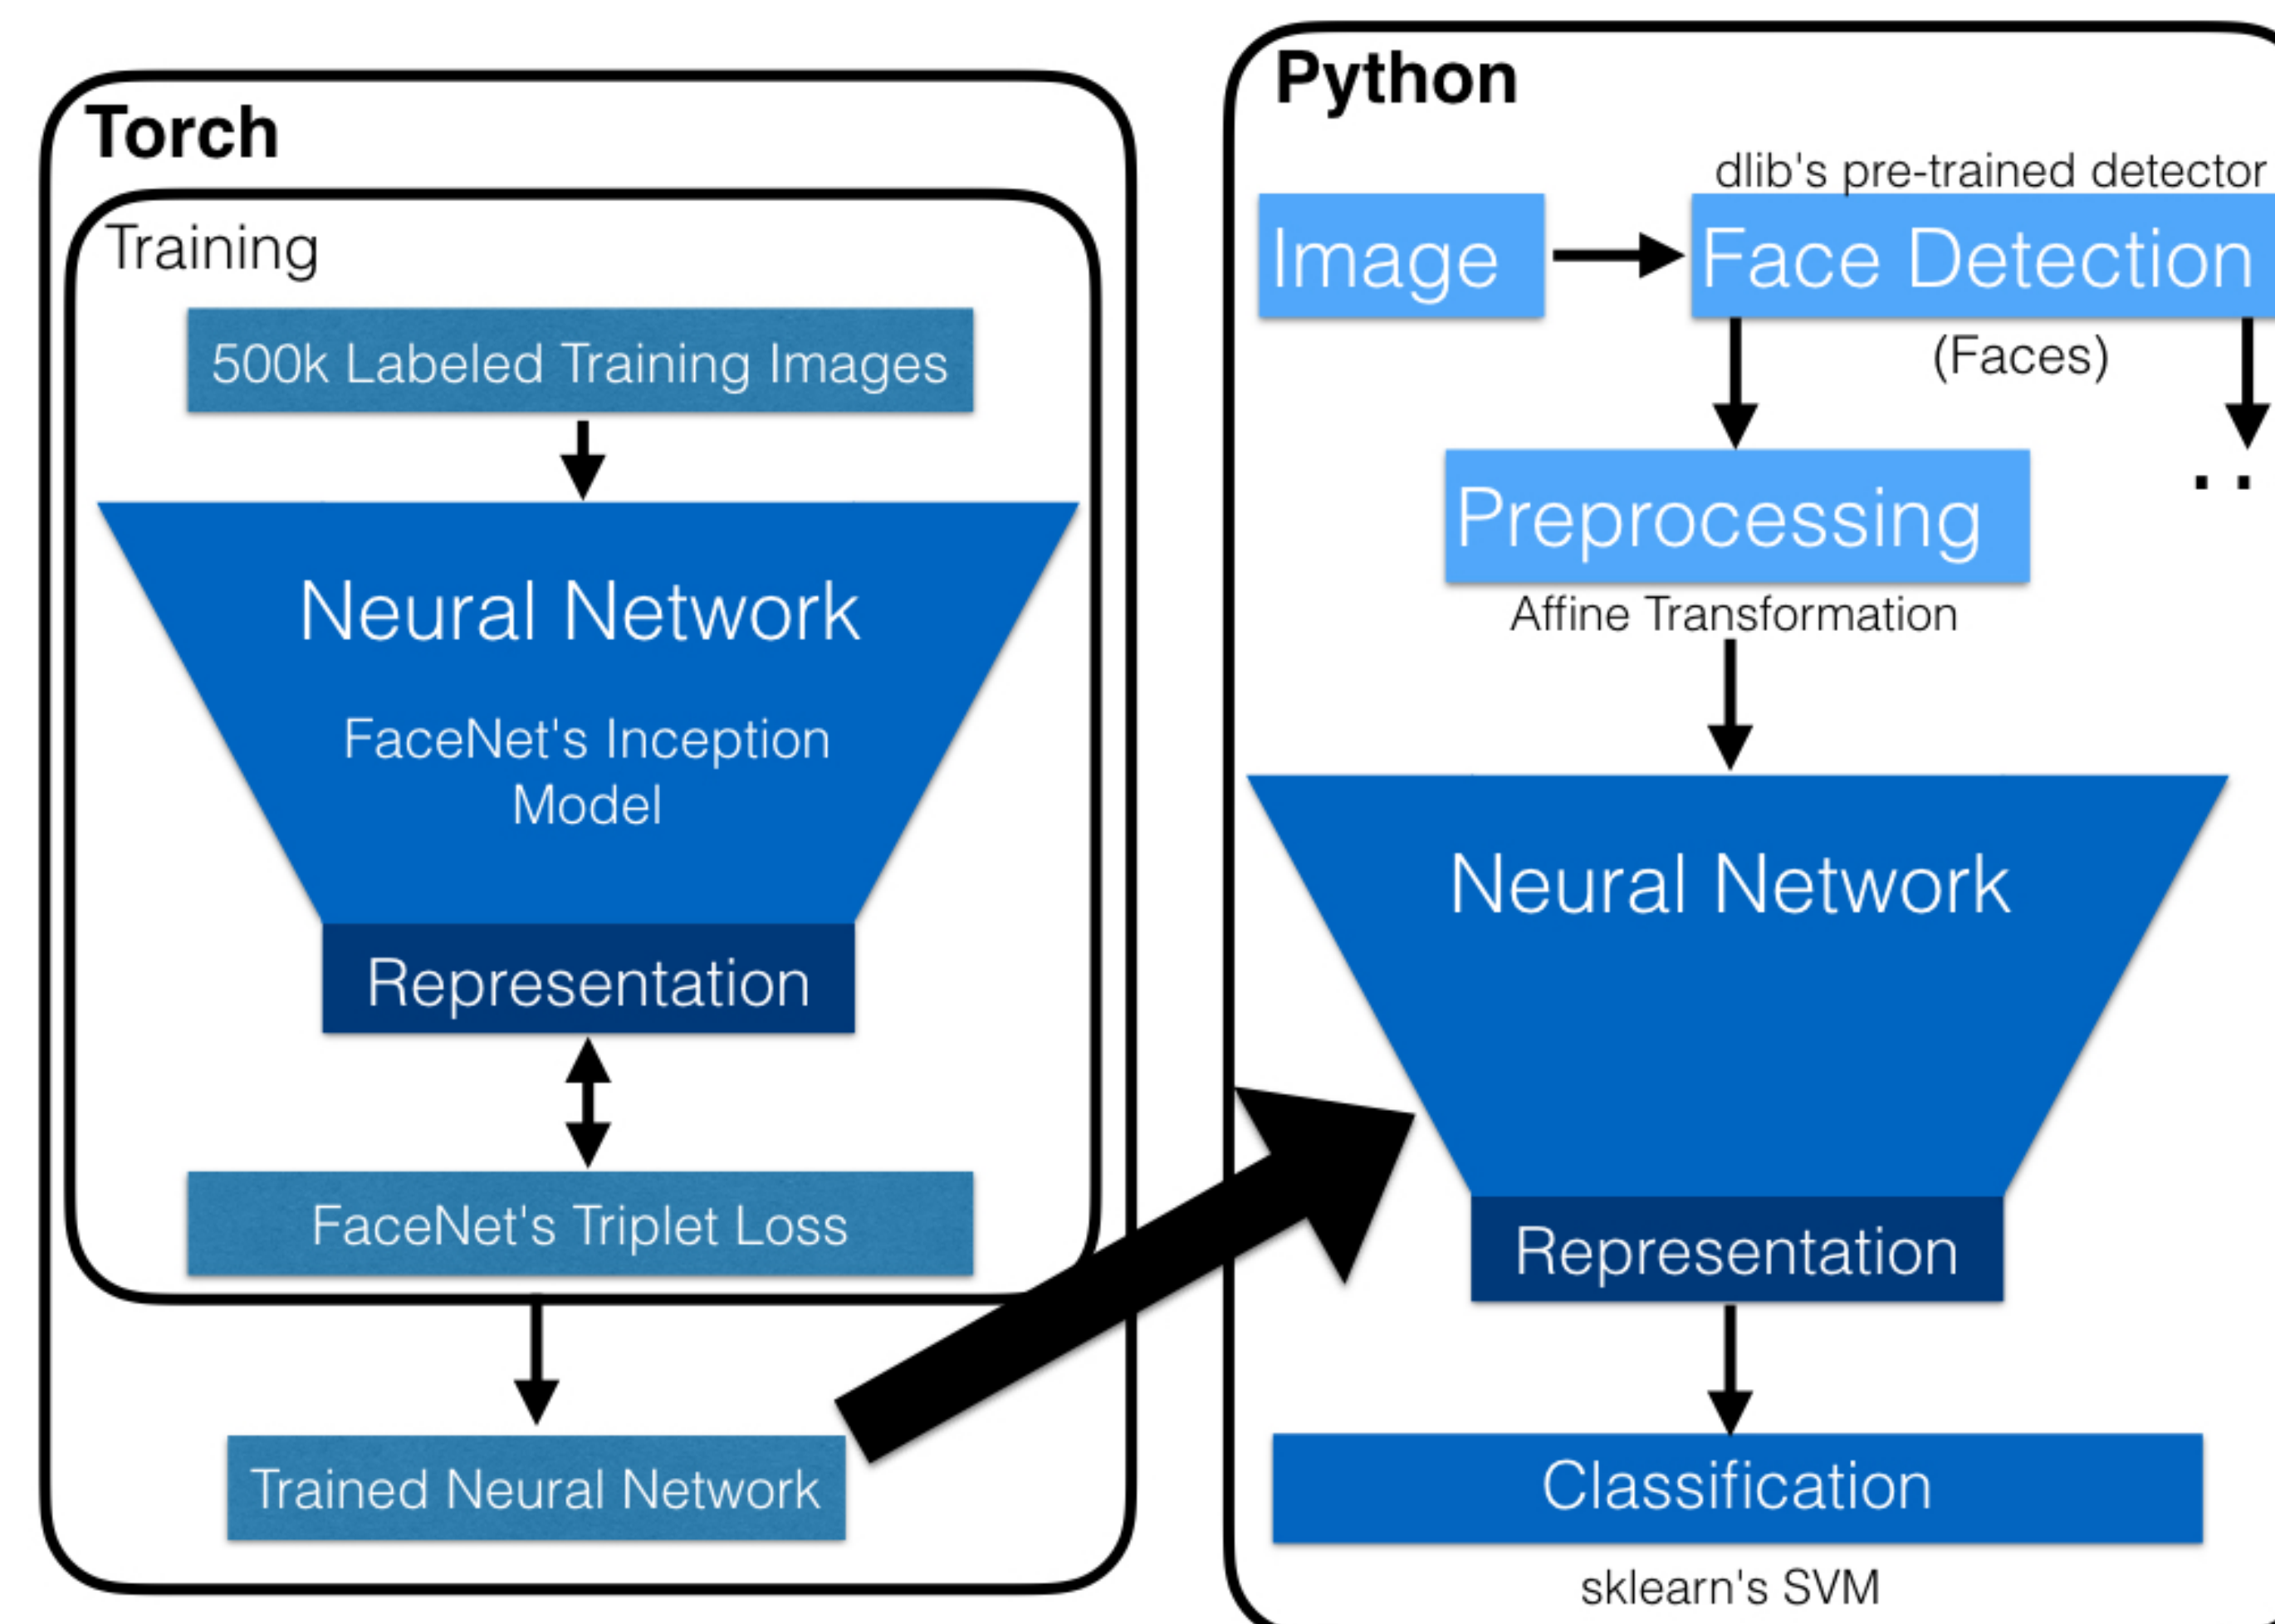
Towards exploring transient and mobile face recognition, we have created *OpenFace* as an open source face recognition library.

OpenFace Overview


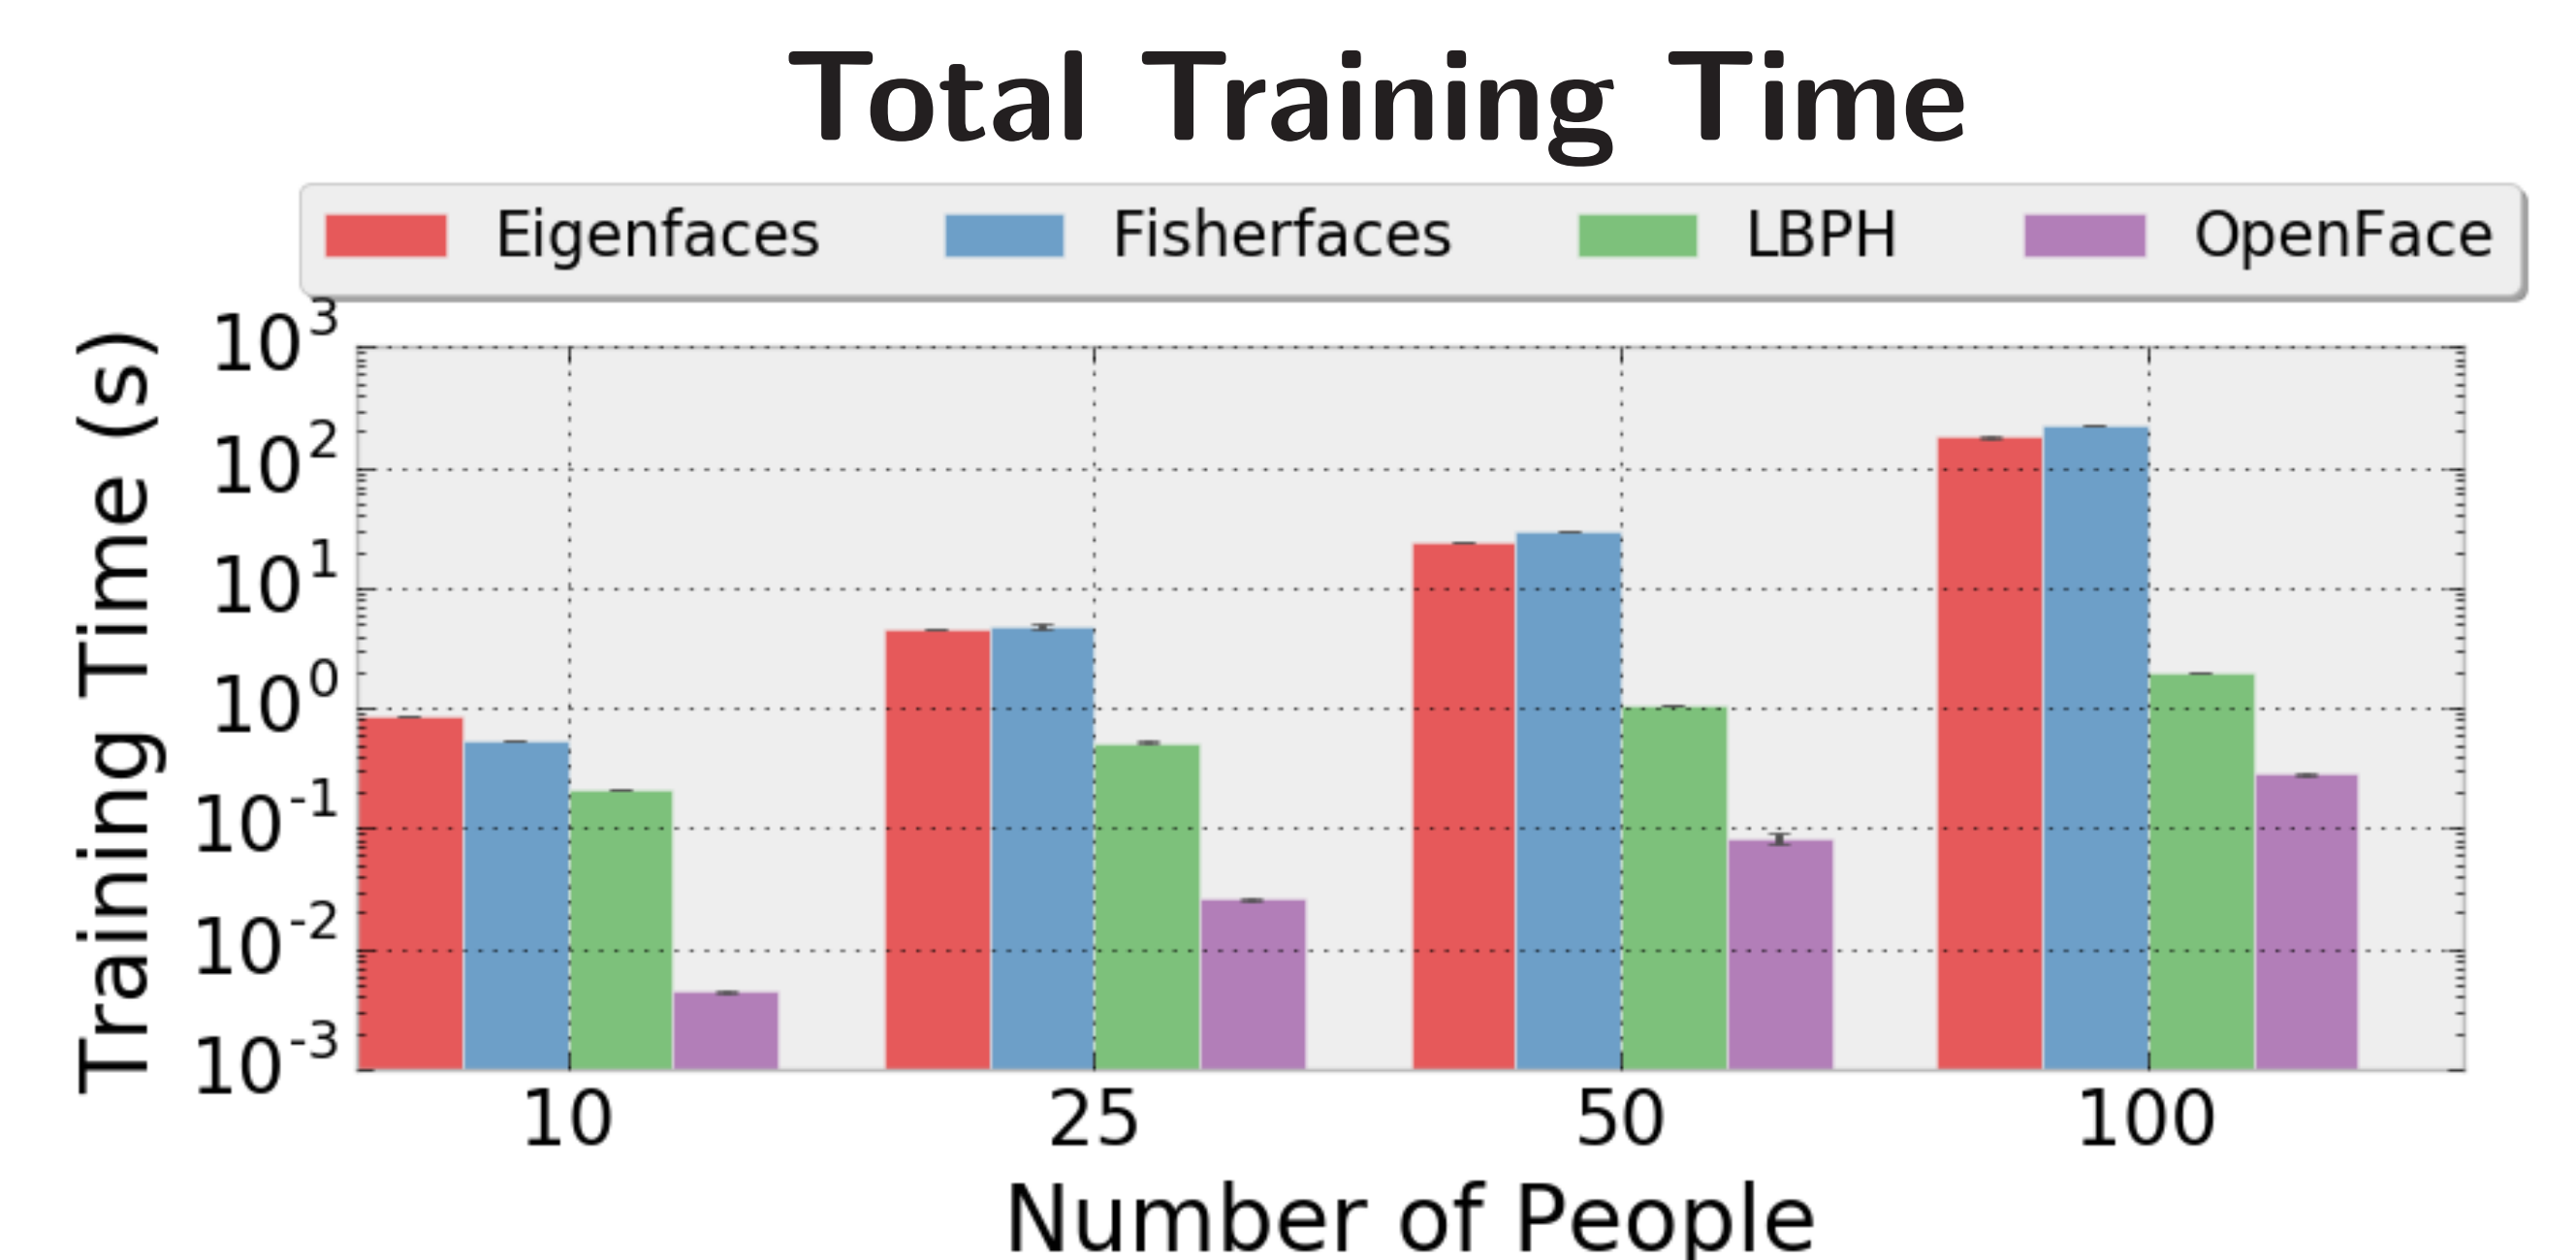

- FaceNet’s [2] neural network structure.
- Faces are embedded on the **unit hypersphere**.
- Over 4,000 GitHub stars.



OpenFace Design



Results: Face Classification



Results: Face Verification

- Predicts whether 2 faces are the same person [1].

Technique	Accuracy
Human-level (cropped)	0.9753
Eigenfaces (no outside data)	0.6002 ± 0.0079
FaceNet [2]	0.9964 ± 0.009
DeepFace-ensemble [3]	0.9735 ± 0.0025
OpenFace	0.9292 ± 0.0134

References

- [1] G. B. Huang et al. *Labeled faces in the wild: A database for studying face recognition in unconstrained environments*. Tech. rep. 07-49, University of Massachusetts, Amherst, 2007.
- [2] F. Schroff, D. Kalenichenko, and J. Philbin. “FaceNet: A Unified Embedding for Face Recognition and Clustering”. In: *CVPR*. 2015, pp. 815–823.
- [3] Y. Taigman et al. “Deepface: Closing the gap to human-level performance in face verification”. In: *CVPR*. 2014, pp. 1701–1708.

Acknowledgments: Hervè Bredin, Zhuo Chen, Khalid Elgazzar, Kiryong Ha, Wenlu Hu, Bartosz Ludwiczuk, Padmanabhan (Babu) Pillai, Rahul Sukthankar, and Junjue Wang.