Human Behavior Understanding with Machine Learning

Albert Ali Salah (Utrecht University)

What is this lecture about?

- **Human Behavior Modeling via Machine Learning**
  - Individual behavior – e.g. facial expression analysis
  - Dyadic behavior – e.g. social interactions
  - Aggregate behavior -- computational social science – e.g. mobility modeling

Two Set of Challenges

Technical Challenges

Are we able to automatically interpret and predict complex human behavior using machine learning techniques?

Human(ity) Challenges

What are the social implications and ethical considerations in the deployment and wide-spread use of these tools?

Computational Analysis of Human Behavior

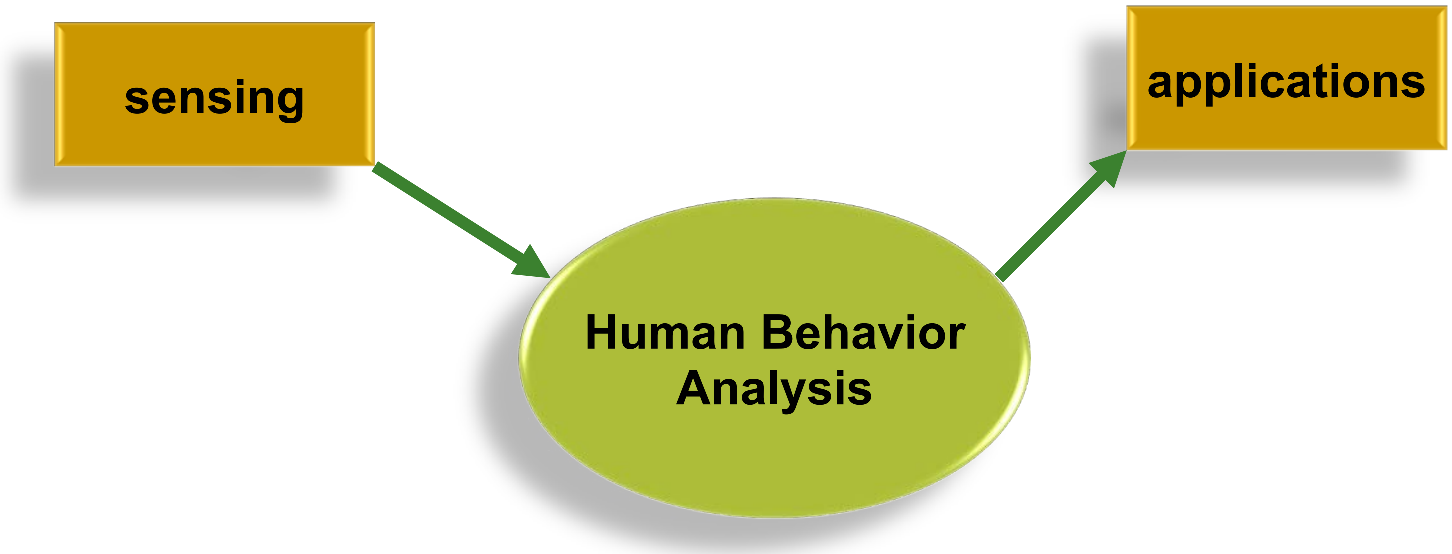

A.A. Salah, T. Gevers, (eds.) Computer Analysis of Human Behavior, Springer Verlag, 2011

Computational Analysis of Human Behavior

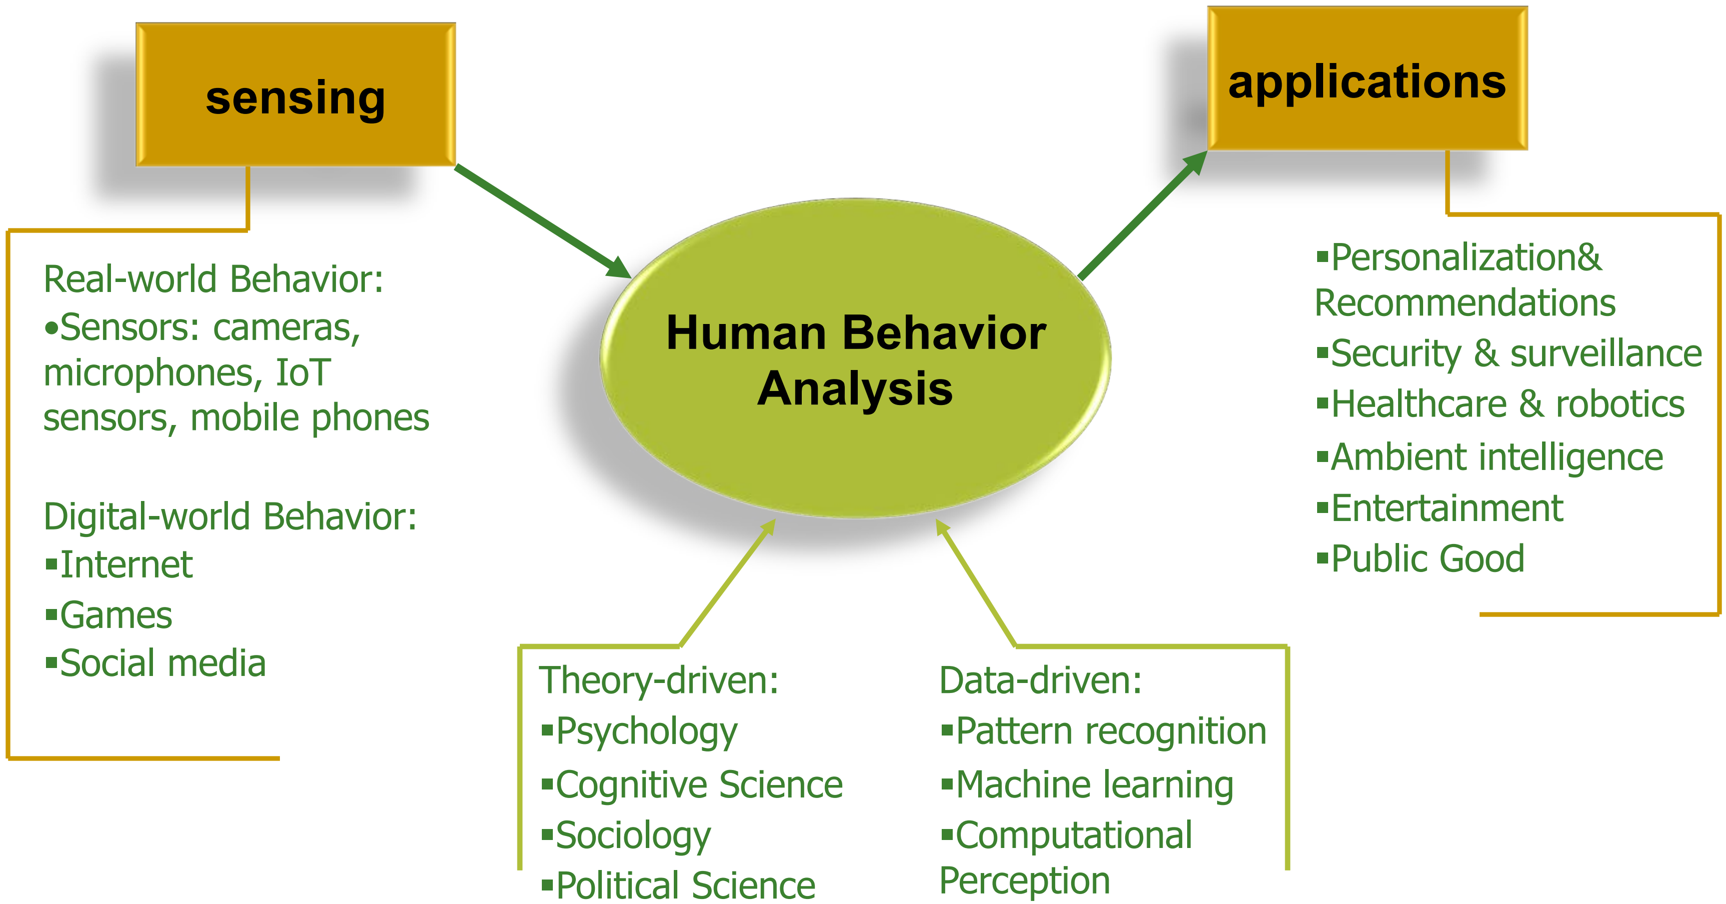

A.A. Salah, T. Gevers, (eds.) Computer Analysis of Human Behavior, Springer Verlag, 2011

Example: Facial expressions

- Both **spontaneous** and **planned** behavior
- Unique to individuals, but recognizable in cultural contexts
- Rich signals, linked with emotions, personality, deception, mental health…

Fundamental questions

- How do we **represent** the behavior?

Representation: Features

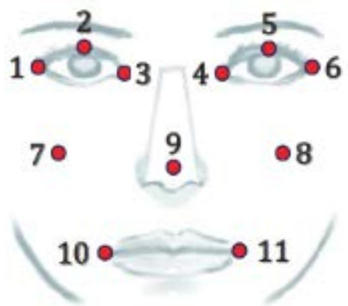

Facial Landmarks

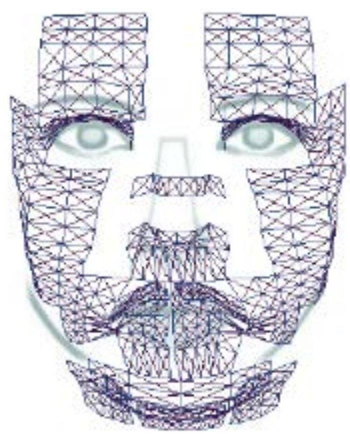

Deformable Patches

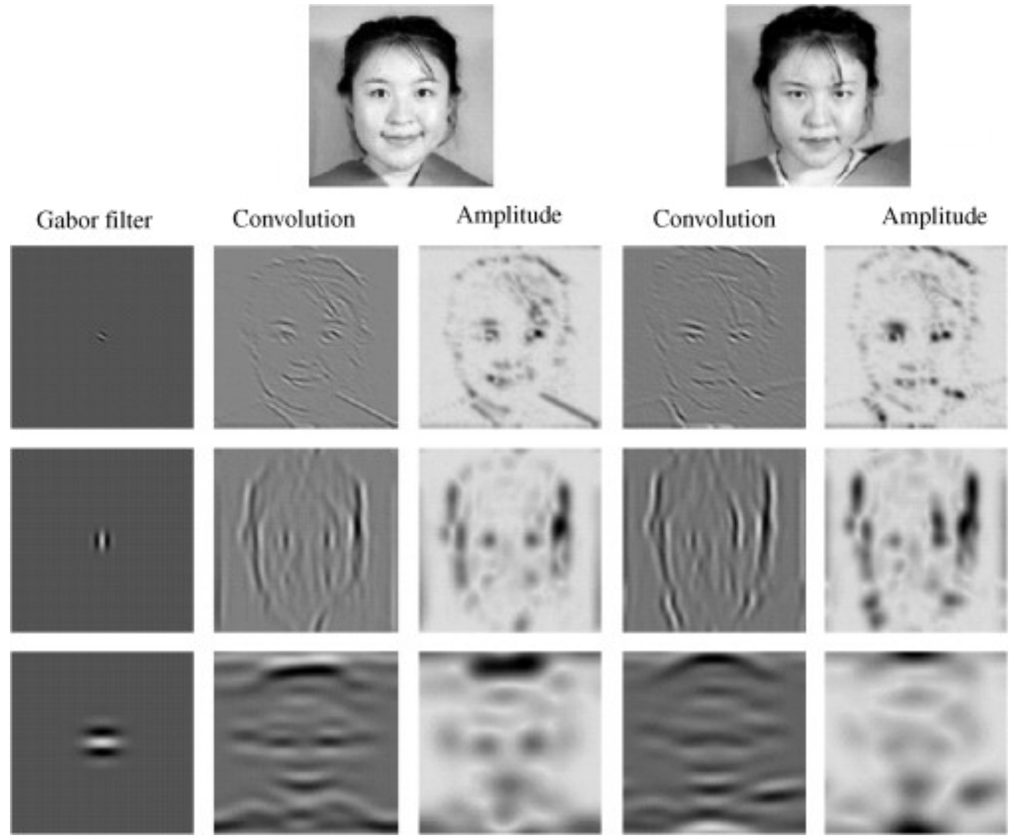

Gabor Wavelets

H. Dibeklioglu, A.A. Salah, and T. Gevers. "A statistical method for 2-d facial landmarking". *IEEE Trans. Image Proc.*, 21(2):844–858, 2012.

Bashyal, S., & Venayagamoorthy, G. K. (2008). "Recognition of facial expressions using Gabor wavelets and learning vector quantization". Eng. Apps of AI, 21(7), 1056-1064.

Representation: Features

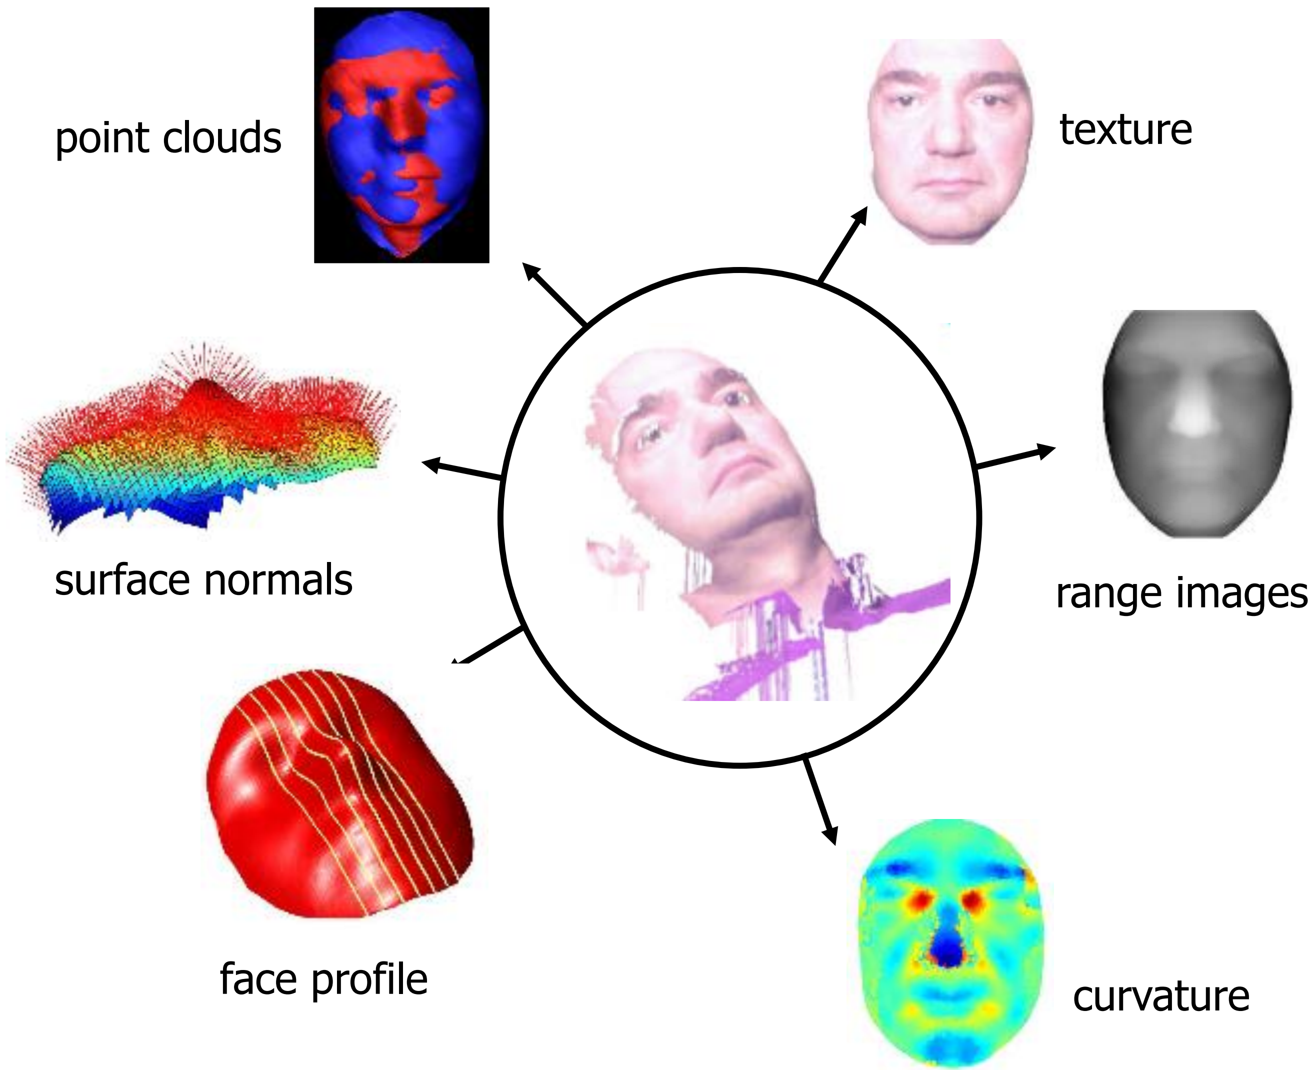

Slide credits: Berk Gökberk

Representation: Features

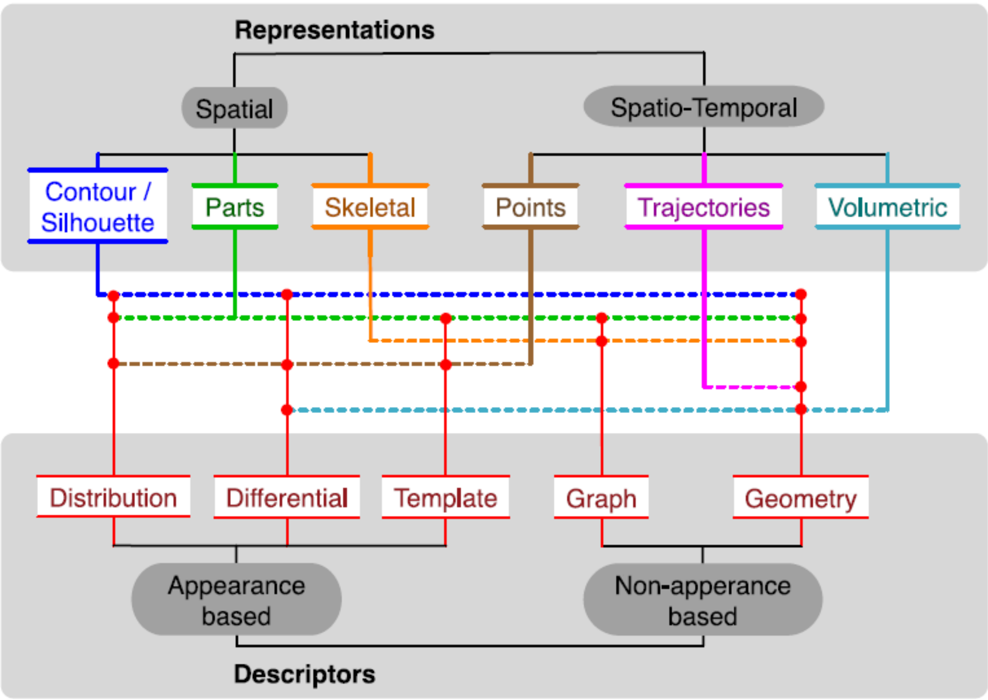

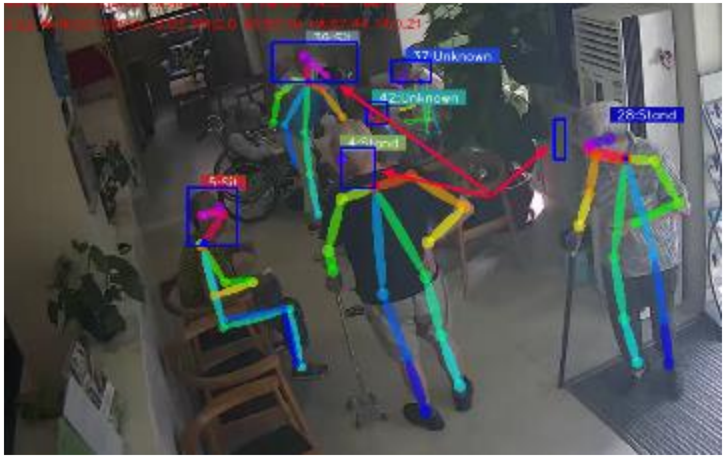

https://github.com/CMU-Perceptual-Computing-Lab/openpose/

Yilmaz, A. (2011). "Detecting and Tracking Action Content". In Salah, Gevers (eds.) Computer Analysis of Human Behavior (pp. 41-68). Springer, London.

Fundamental questions

- How do we represent the behavior?
- How do we establish **ground truth**?

Annotation: Categorical (Discrete)

- Anger
- Fear
- Disgust
- Surprise
- Happiness
- Sadness

Ekman, P. “Basic Emotions”, Handbook of cognition and emotion, 1999

Annotation: Continuous

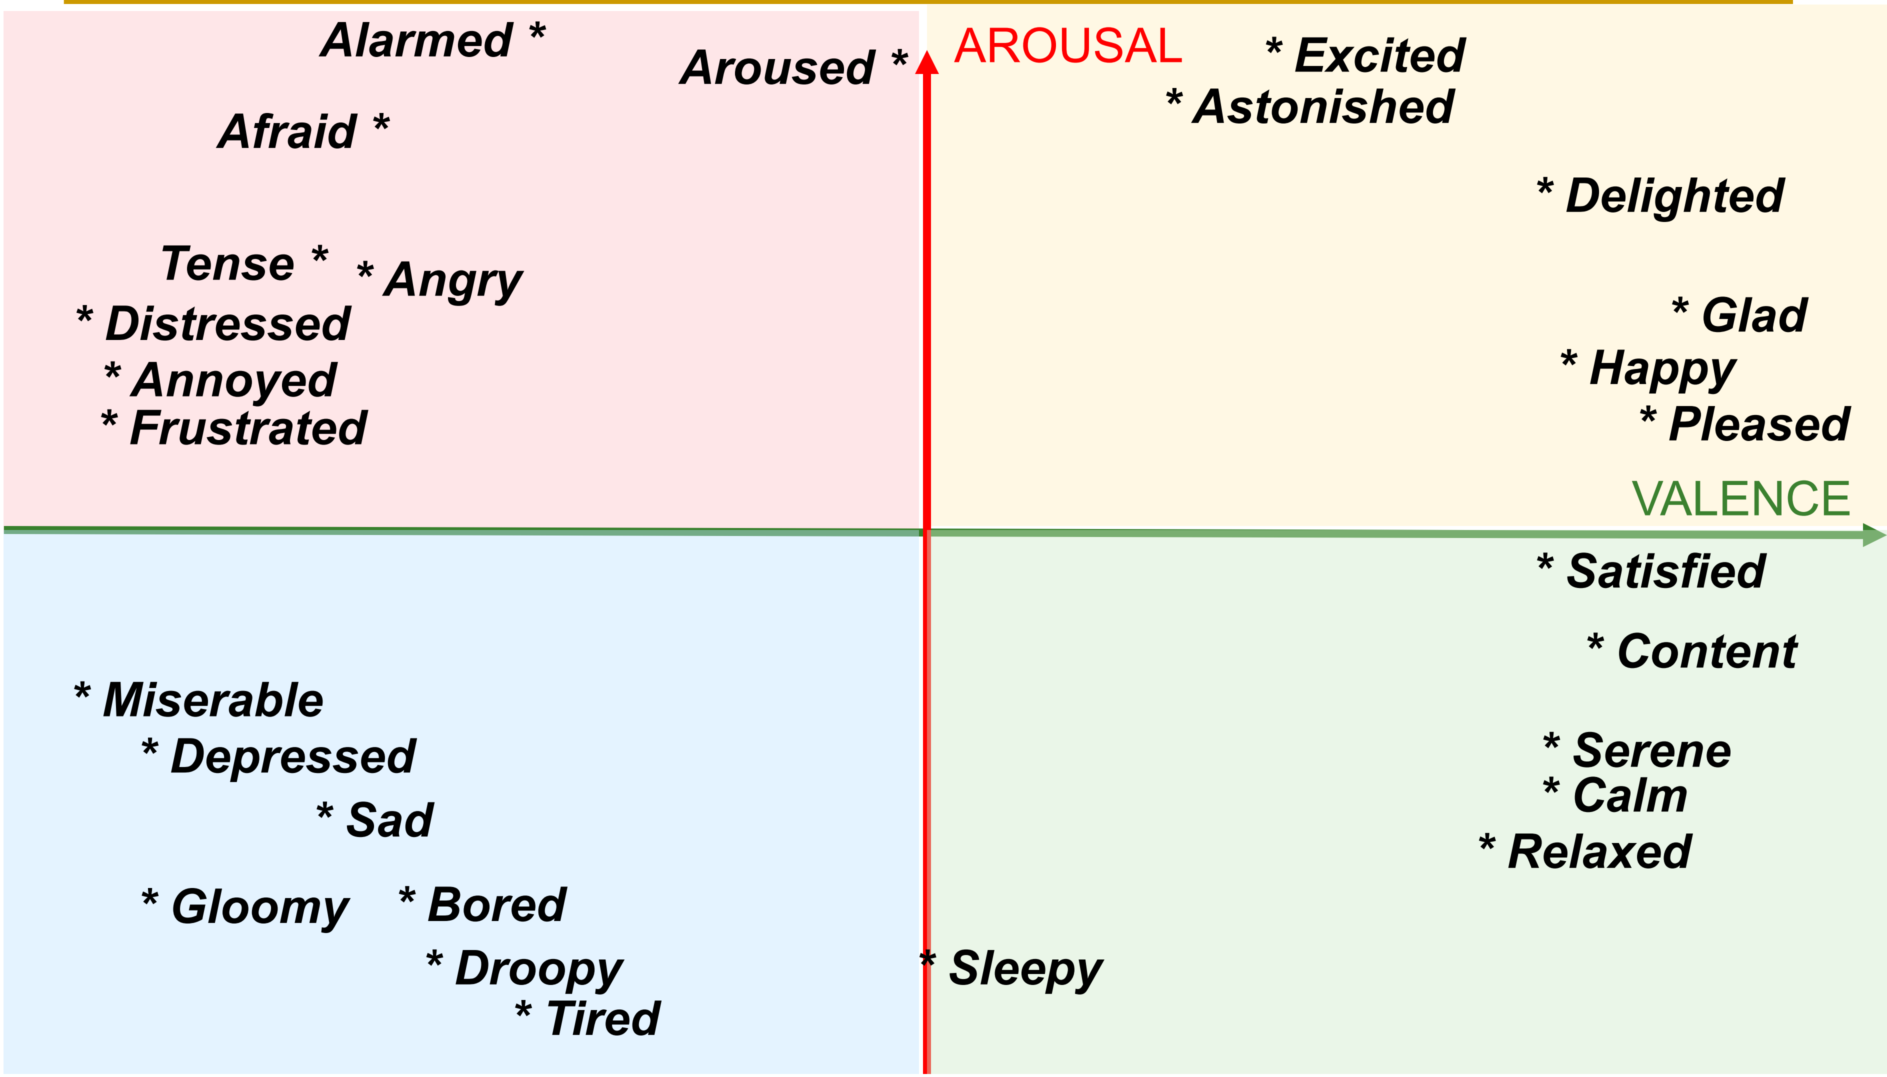

Russell, J. A. (1980). "A circumplex model of affect". *Journal of personality and social psychology*, *39*(6), 1161.

Annotation: Objective vs. Subjective

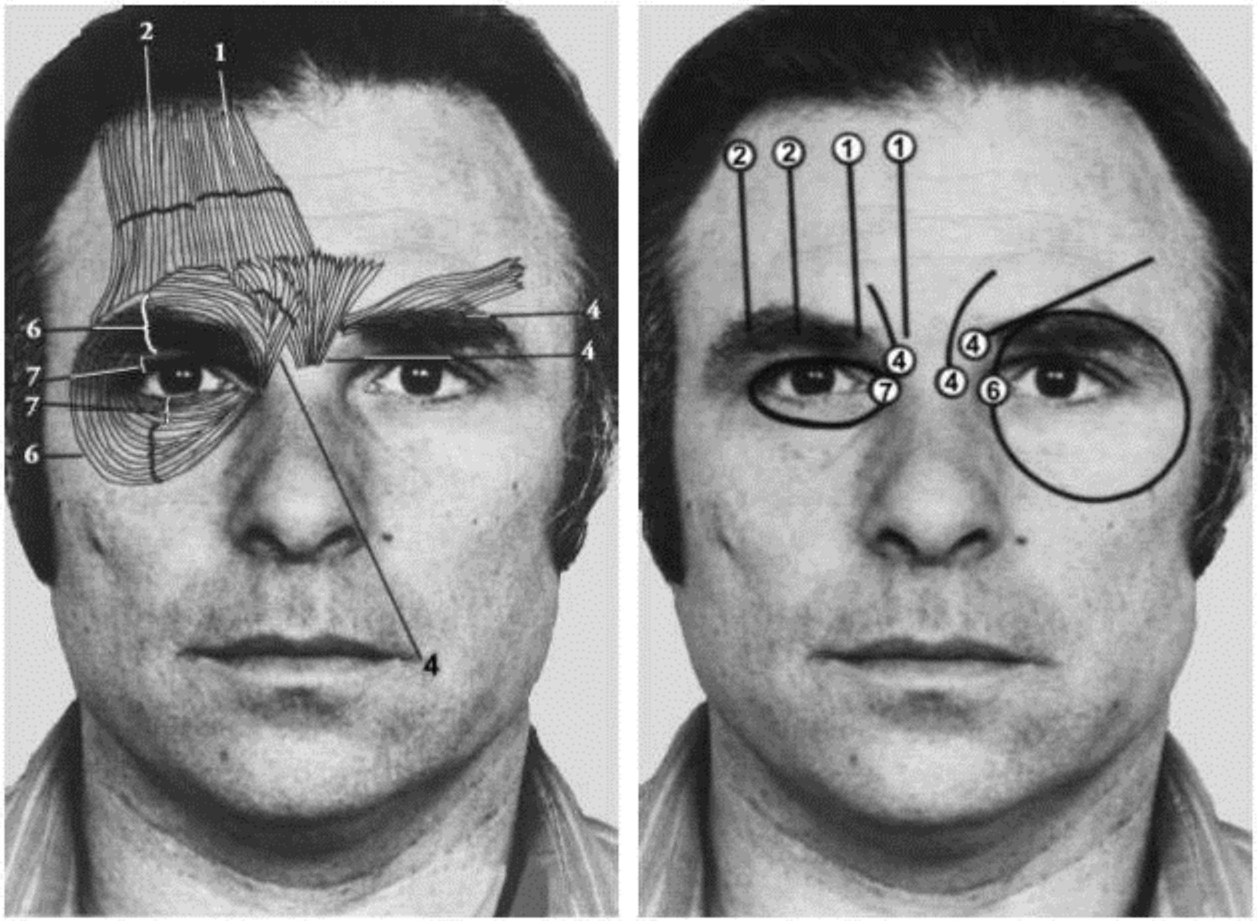

Annotation: Objective vs. Subjective

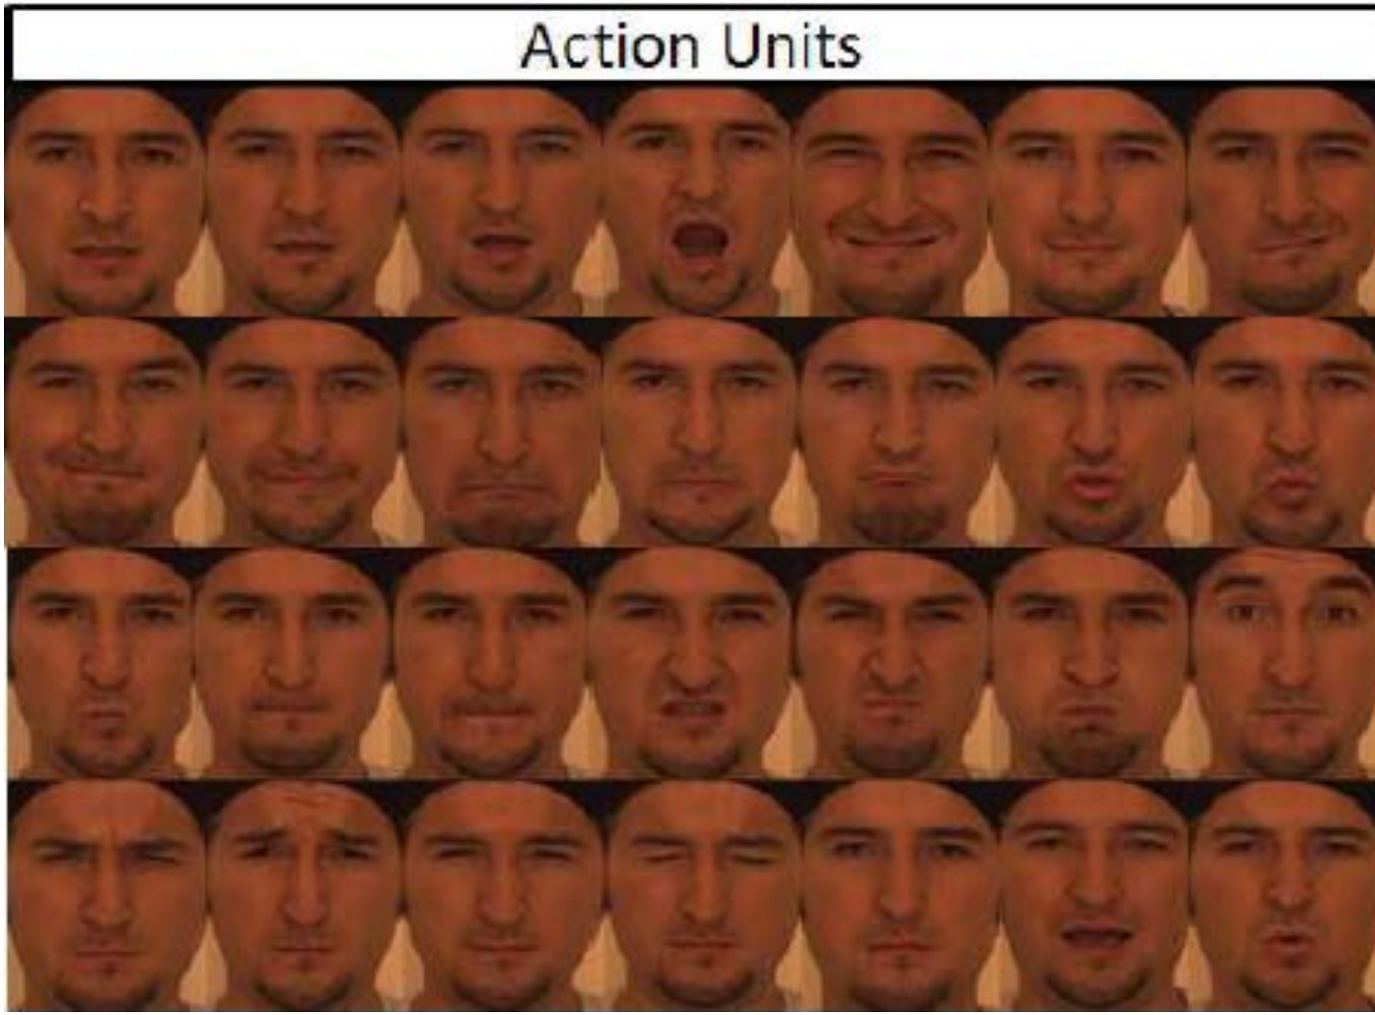

Alyüz, N., Gökberk, B., Dibeklioğlu, H., Savran, A., Salah, A. A., Akarun, L., & Sankur, B. (2008, May). "3D face recognition benchmarks on the Bosphorus database with focus on facial expressions". In European workshop on biometrics and identity management (pp. 57-66)

Application example: real vs. fake pain

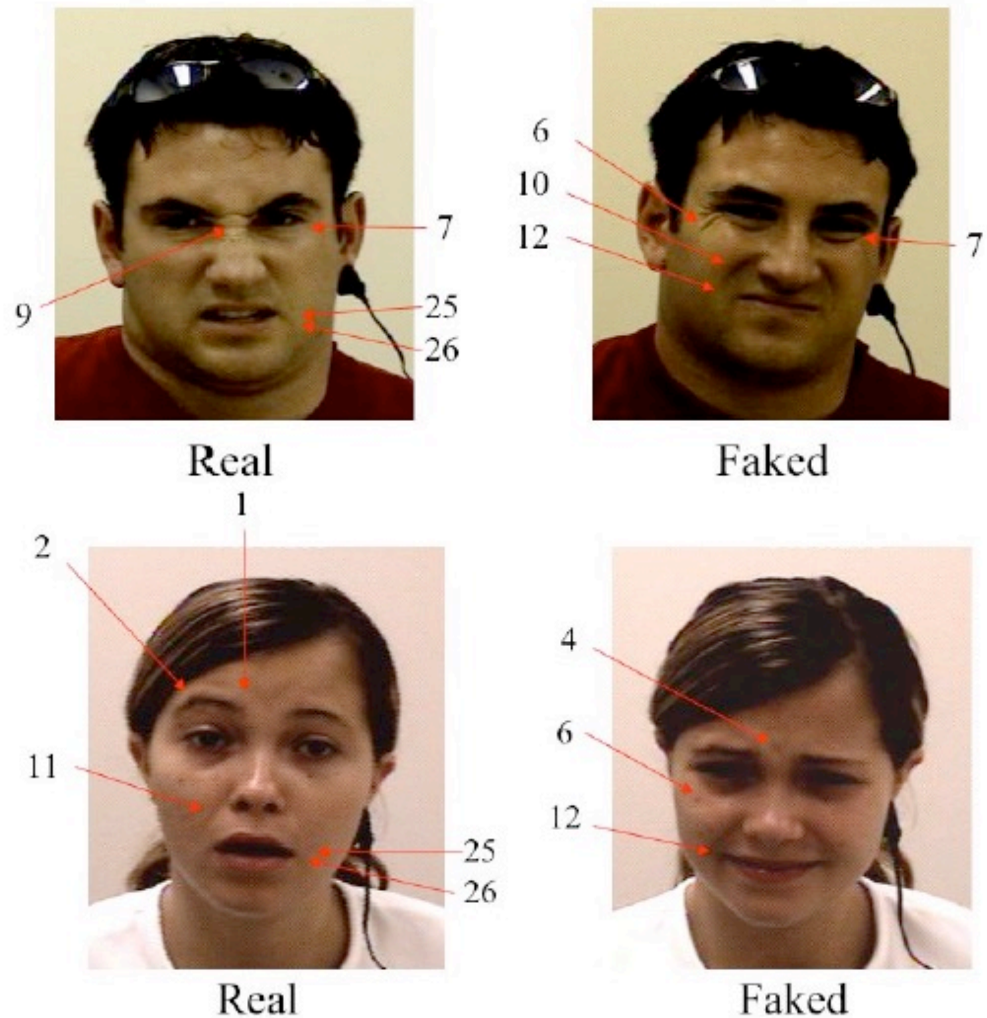

Bartlett, M., Littlewort, G., Vural, E., Lee, K., Cetin, M., Ercil, A., & Movellan, J. (2008). "Data mining spontaneous facial behavior with automatic expression coding". In *Verbal and Nonverbal Features of Human-Human and Human-Machine Interaction* (pp. 1-20). Springer, Berlin, Heidelberg.

Annotation: Custom labels

frustration disappointment anger triumph

relief elation confusion

Schimmel, A., M. Doyran, P. Baki, K. Ergin, B. Türkmen, A. Akdag Salah, S. Bakkes, H. Kaya, R. Poppe, AA Salah, MP-BGAAD: "Multi-Person Board Game Affect Analysis Dataset", Proc. eNTERFACE, 2019.

Annotation: Context

Aviezer, H., Trope, Y., & Todorov, A. (2012). “Body cues, not facial expressions, discriminate between intense positive and negative emotions”. *Science*, *338*(6111), 1225-1229.

Annotation: Context

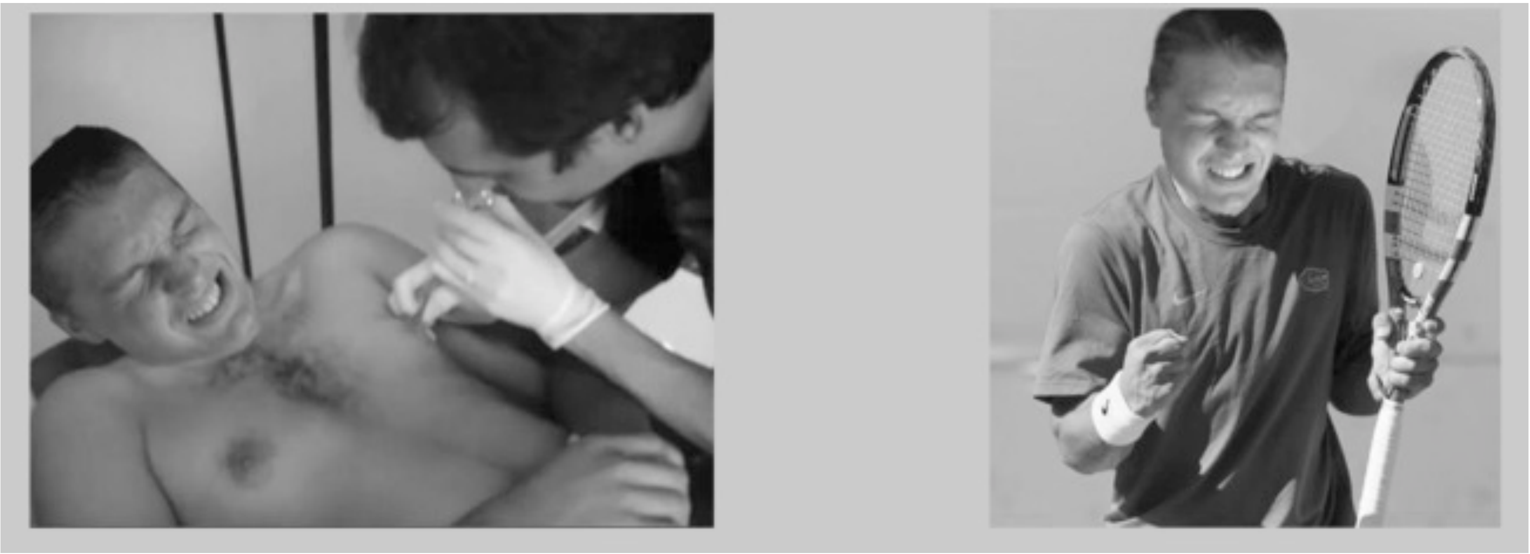

Aviezer, H., Trope, Y., & Todorov, A. (2012). "Body cues, not facial expressions, discriminate between intense positive and negative emotions". *Science*, *338*(6111), 1225-1229.

Annotation: Ground truth

- Create **controlled** situations

Example: Controlled ground truth

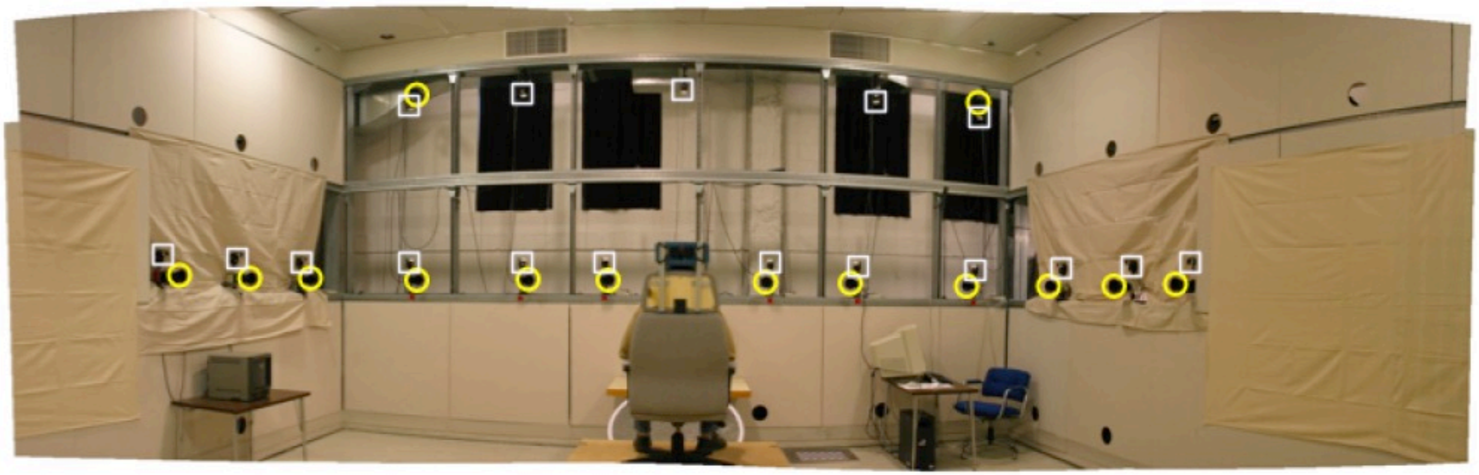

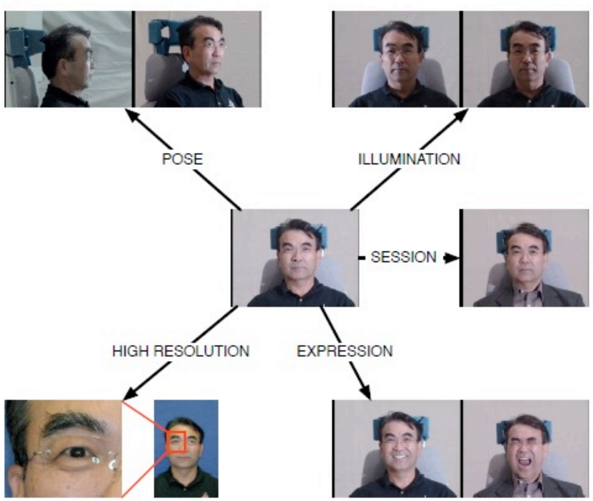

Gross, R., Matthews, I., Cohn, J., Kanade, T., & Baker, S. (2010). “Multi-pie”. *Image and Vision Computing*, *28*(5), 807-813.

Example: Controlled ground truth

- 2D/3D face database for
  - Facial expression classification
  - Spontaneous/Posed smile classification
  - Age estimation
  - Head pose estimation (3D data)
  - Color constancy
- 1240 videos (597 spontaneous, 643 posed) from 400 subjects (185 female, 215 male)
- 1920×1080 pixels resolution @50 fps
- Age varies from 8 to 76
- www.face2age.com

H. Dibeklioglu, A.A. Salah, and T. Gevers. Are you really smiling at me? Spontaneous versus posed smiles. *Proc. ECCV* (3):525-538, 2012.

Example: Controlled ground truth

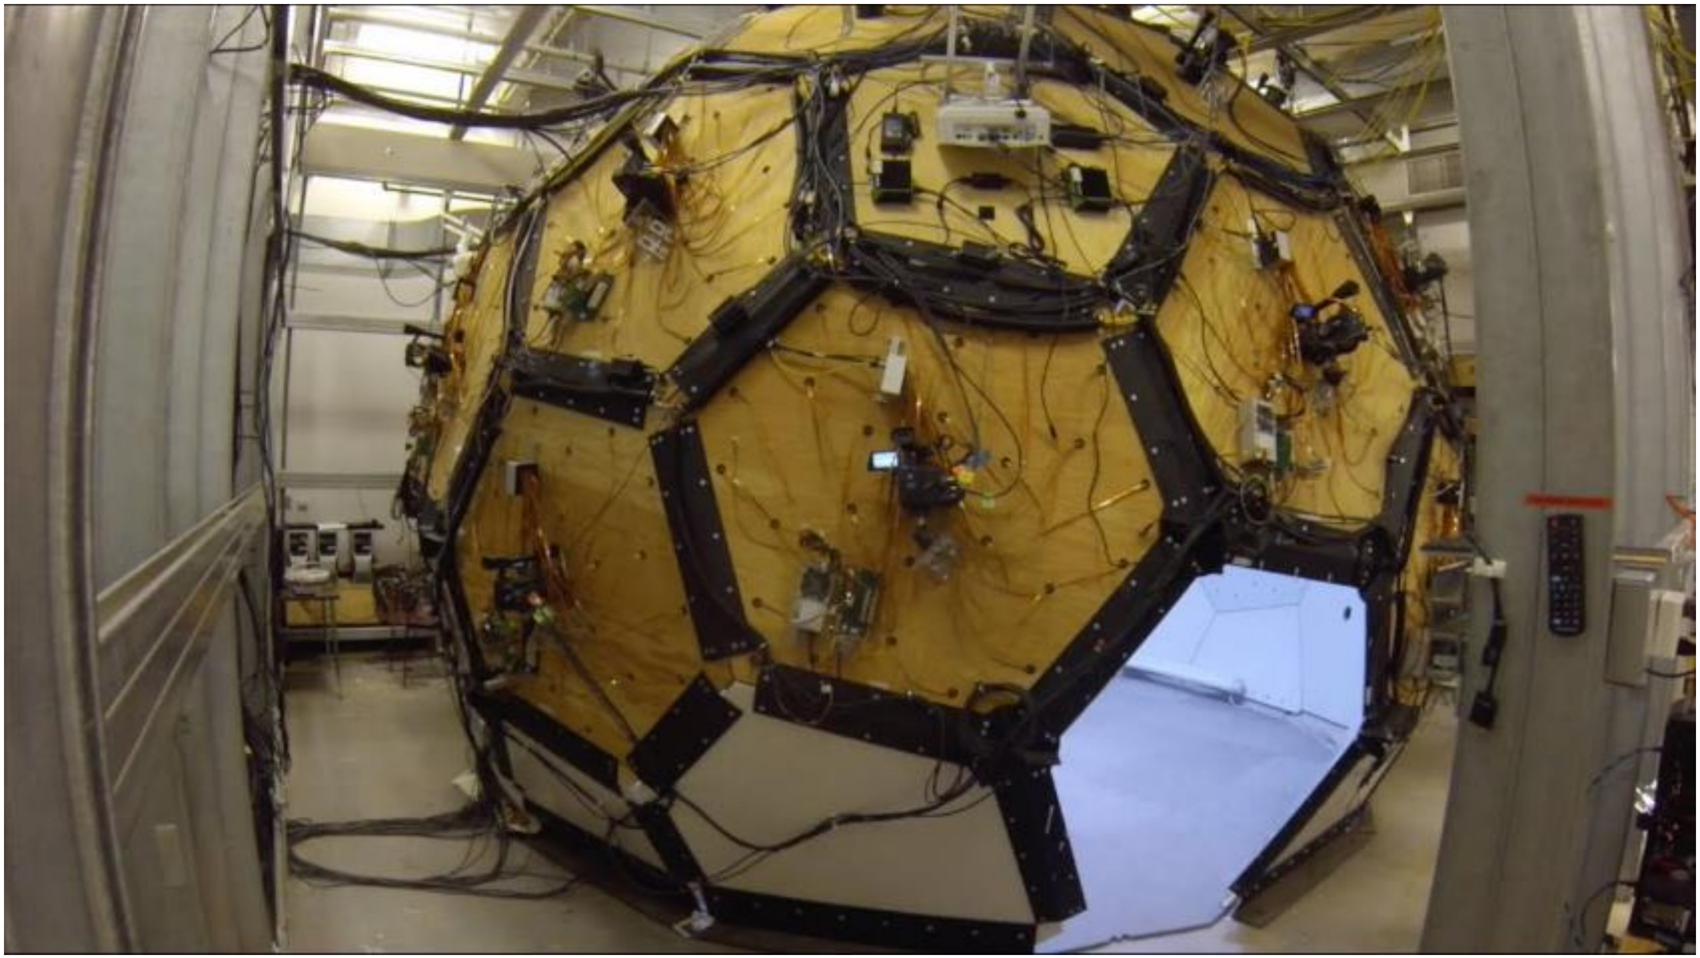

Joo, H., Simon, T., Li, X., Liu, H., Tan, L., Gui, L., ... & Kanade, T. (2017). Panoptic studio: A massively multiview system for social interaction capture. IEEE TPAMI, 41(1), 190-204.

Example: Controlled ground truth

Joo, H., Simon, T., Li, X., Liu, H., Tan, L., Gui, L., ... & Kanade, T. (2017). Panoptic studio: A massively multiview system for social interaction capture. IEEE TPAMI, 41(1), 190-204.

Annotation: Ground truth

- Create controlled situations
- Find situations with **genuine behavioral displays**

Annotation: Ecological validity

Matsumoto, D., & Hwang, H. S. (2012). "Evidence for a nonverbal expression of triumph". *Evolution and Human Behavior*, *33*(5), 520-529.

Annotation: Ground truth

- Create controlled situations
- Find situations with genuine behavioral displays
- Get **expert annotations** from trained people and ensure high interrater agreement

Annotation: Expert annotation

- Children engaging in play therapy with a psychotherapist
- A validated behavior assessment tool is used for annotations.

Doyran, M., Türkmen, B., Oktay, E. A., Halfon, S., & Salah, A. A. (2019). "Video and Text-Based Affect Analysis of Children in Play Therapy". In *ACM ICMI.*

Annotation: Ground truth

- Create controlled situations
- Find situations with genuine behavioral displays
- Get expert annotations from trained people and ensure high interrater agreement
- Use **semi-automatic** approaches to **pre-select** what to **annotate**

Annotation: Semi-automatic

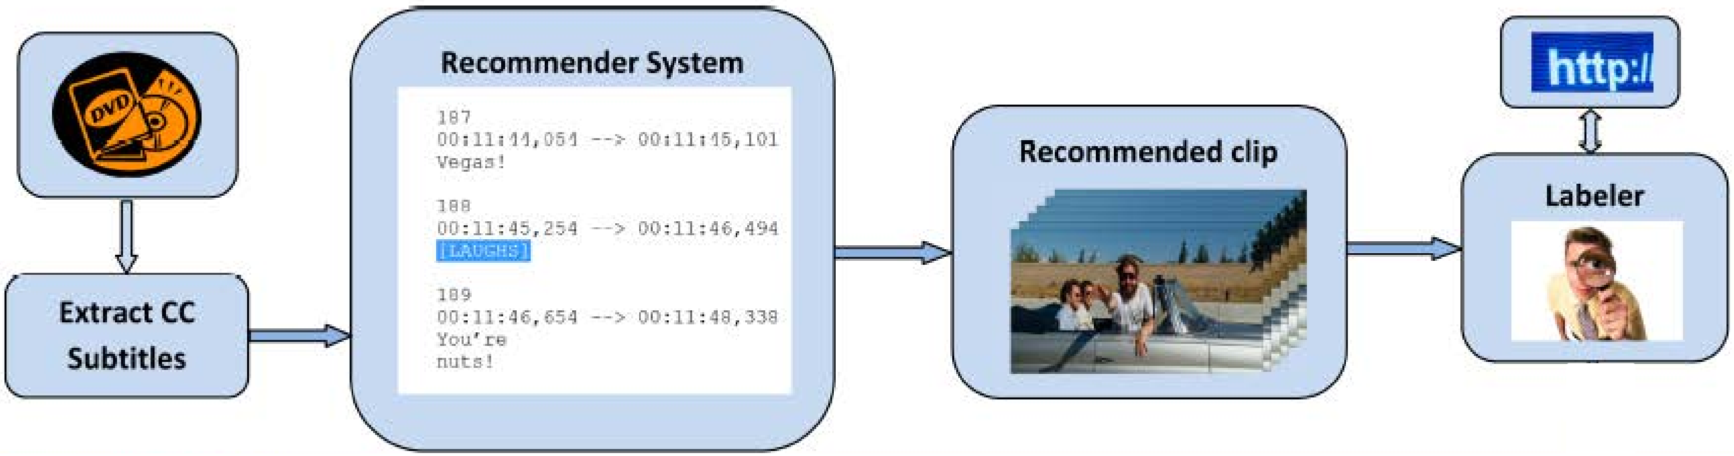

A. Dhall, R. Goecke, S. Lucey and T. Gedeon, "Collecting large, richly annotated facial-expression databases from movies", IEEE Multimedia 2012

Annotation: Ground truth

- Create controlled situations
- Find situations with genuine behavioral displays
- Get expert annotations from trained people and ensure high interrater agreement
- Use semi-automatic approaches to pre-select what to annotate
- Use **additional sensors** that can provide ground truth

Annotation: Additional sensors

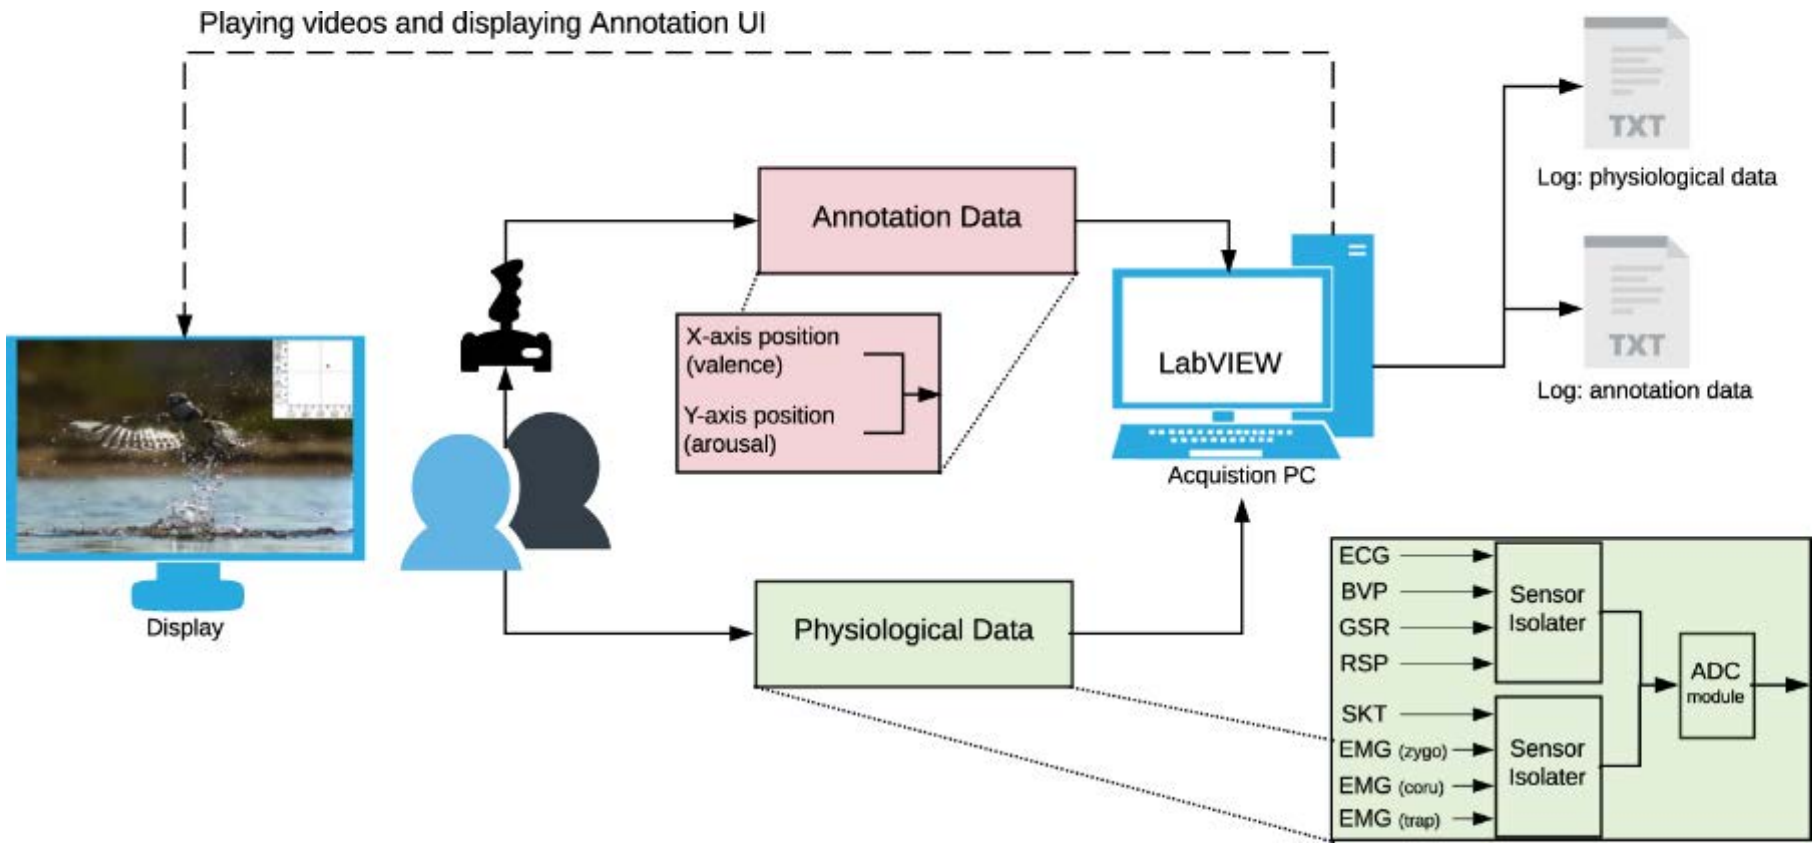

K. Sharma, C. Castellini, E.L. van den Broek, A. Albu-Schaeffer & F. Schwenker, “A dataset of continuous affect annotations and physiological signals for emotion analysis”, *Scientific Data,* 6:196 (2019)

Annotation: The smile of Mona Lisa

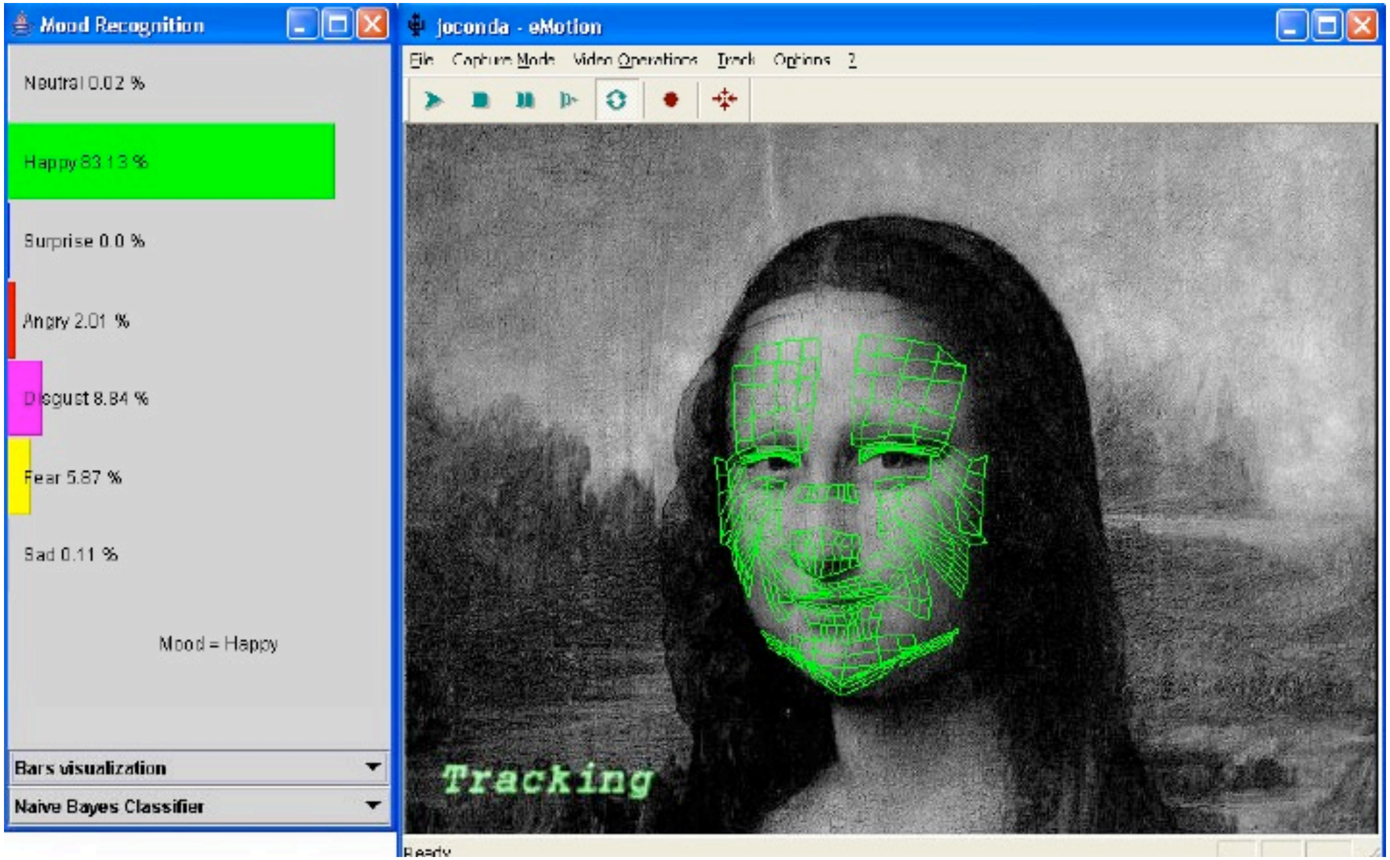

https://edition.cnn.com/2005/TECH/12/16/mona.lisa.smile/

Cohen, I., Sebe, N., Garg, A., Chen, L. S., & Huang, T. S. (2003). "Facial expression recognition from video sequences: temporal and static modeling". CVIU, 91(1-2), 160-187.

Fundamental questions

- How do we represent the behavior?
- How do we establish ground truth?
- How do we approach the problem with **machine learning**?

Learning: ML pipeline

Images / videos

Feature Extraction

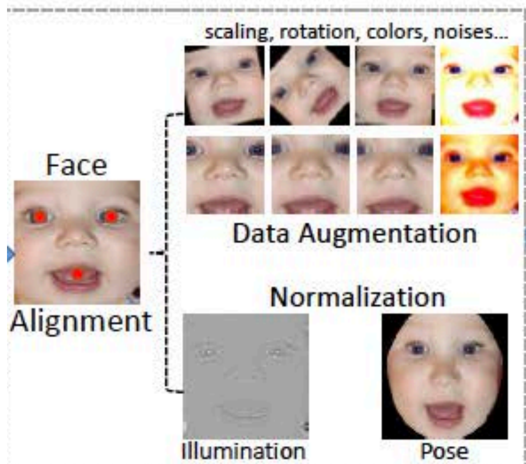

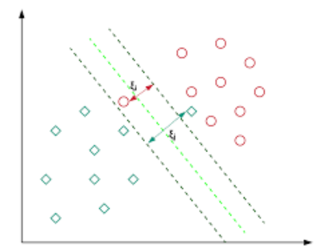

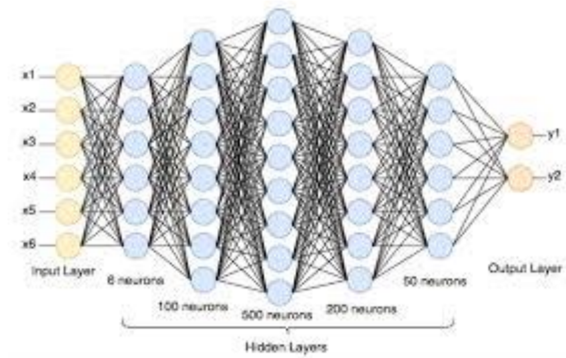

- Happy
- Sad
- Afraid
- Angry
- Surprised
- Disgusted

Input | Pre-processing | ML-model | Output

Learning: Sources of variance

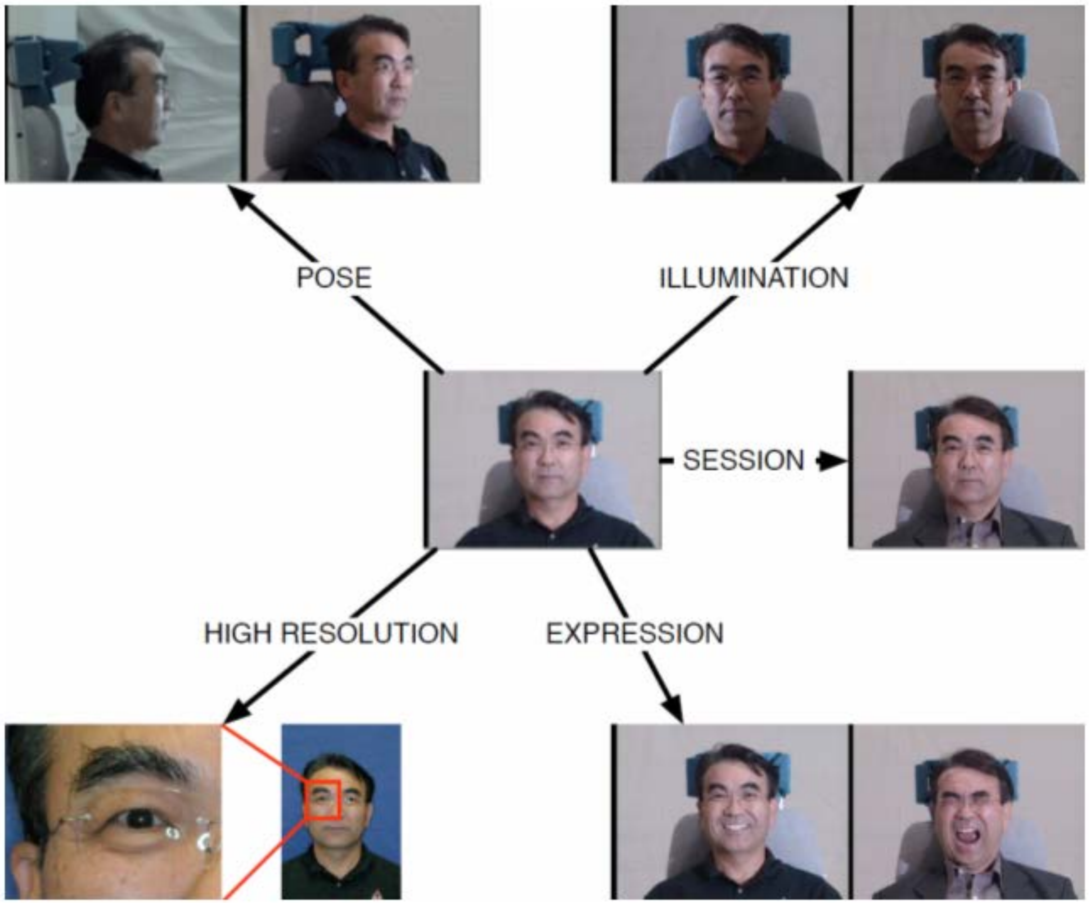

Gross, R., Matthews, I., Cohn, J., Kanade, T., & Baker, S. (2010). “Multi-pie”. *Image and Vision Computing*, *28*(5), 807-813.

EMPATH: First NeurIPS paper on facial expressions

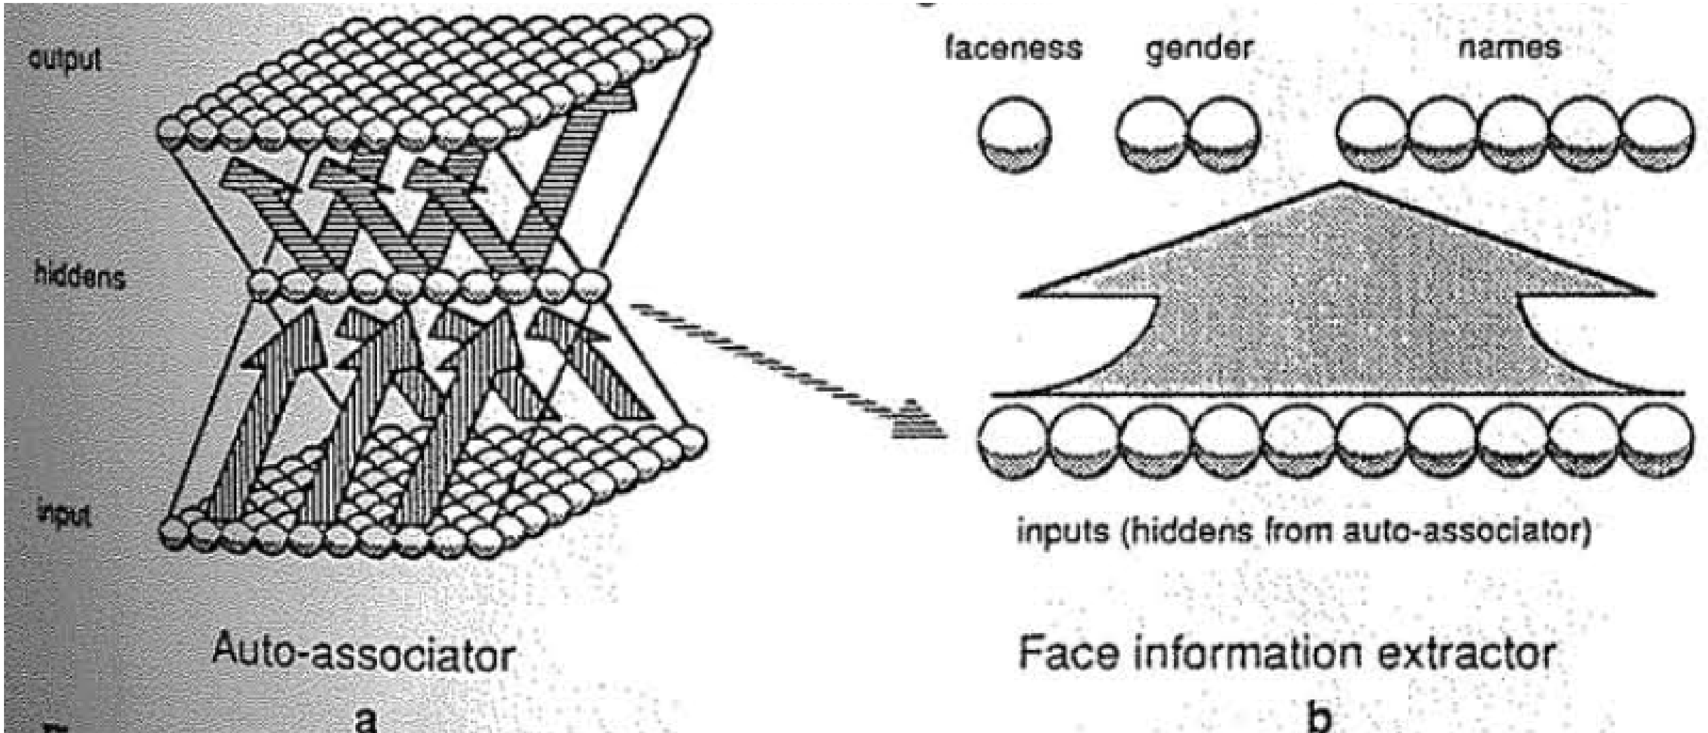

Figure 1: The face recognition model. (a) An image compression network is trained first to compress the 4096 inputs into 40 hidden units. (b) The hidden units from the first network are used as inputs to various recognition networks.

Cottrell, Garrison W., and Janet Metcalfe. "EMPATH: Face, emotion, and gender recognition using holons." *Advances in neural information processing systems*. 1991.

EMPATH: First NeurIPS paper on facial expressions

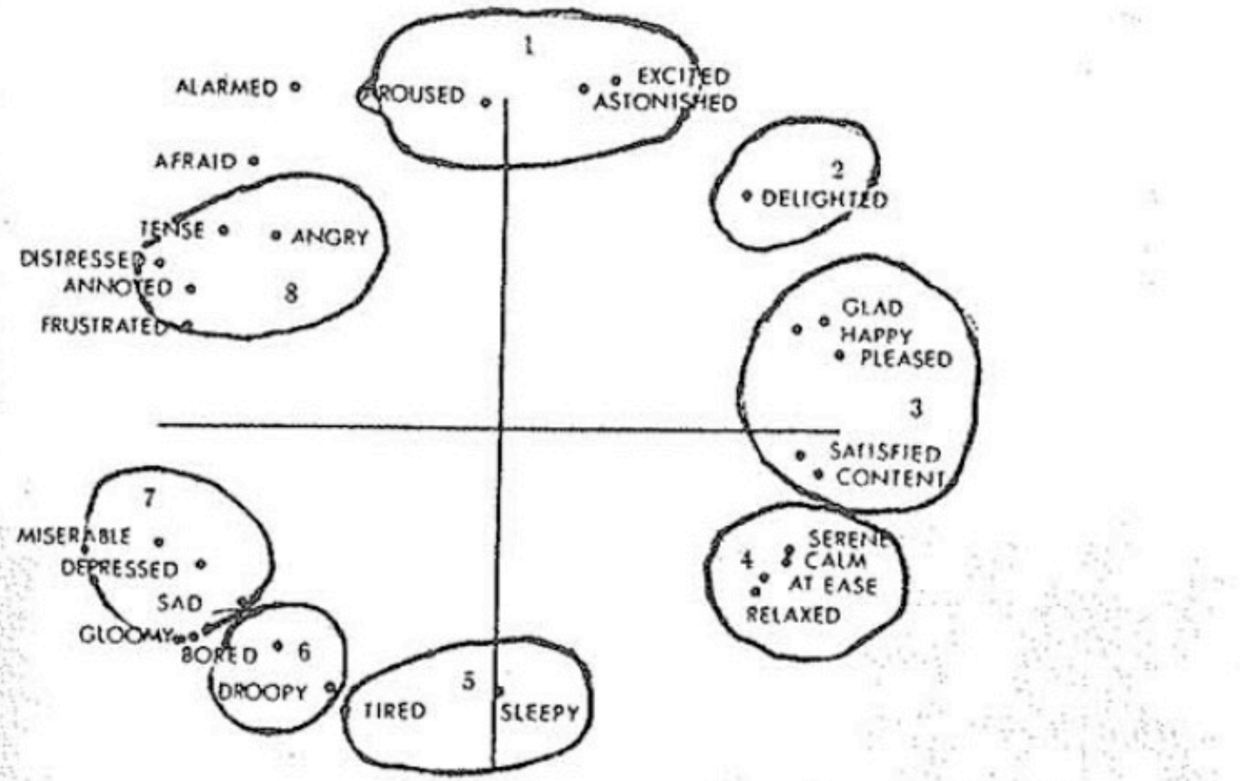

Multidimensional scaling solution for 28 affect words.

The two-dimensional emotion space extracted via multi-dimensional scaling from similarity ratings. Data from Russell (1980).

Cottrell, Garrison W., and Janet Metcalfe. "EMPATH: Face, emotion, and gender recognition using holons." *Advances in neural information processing systems*. 1991.

EMPATH: First NeurIPS paper on facial expressions

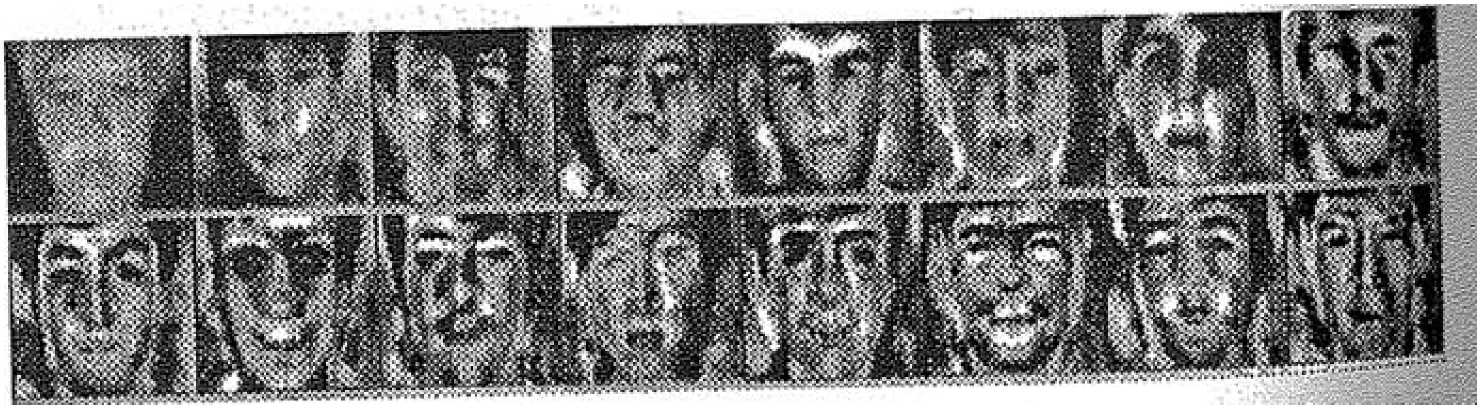

Holons derived by PCA from hidden unit responses.

“These are similar to the «eigenfaces» found by Turk & Pentland (submitted) in their principal components analysis of faces.”

Cottrell, Garrison W., and Janet Metcalfe. "EMPATH: Face, emotion, and gender recognition using holons." *Advances in neural information processing systems*. 1991.

Example: Image based shaving

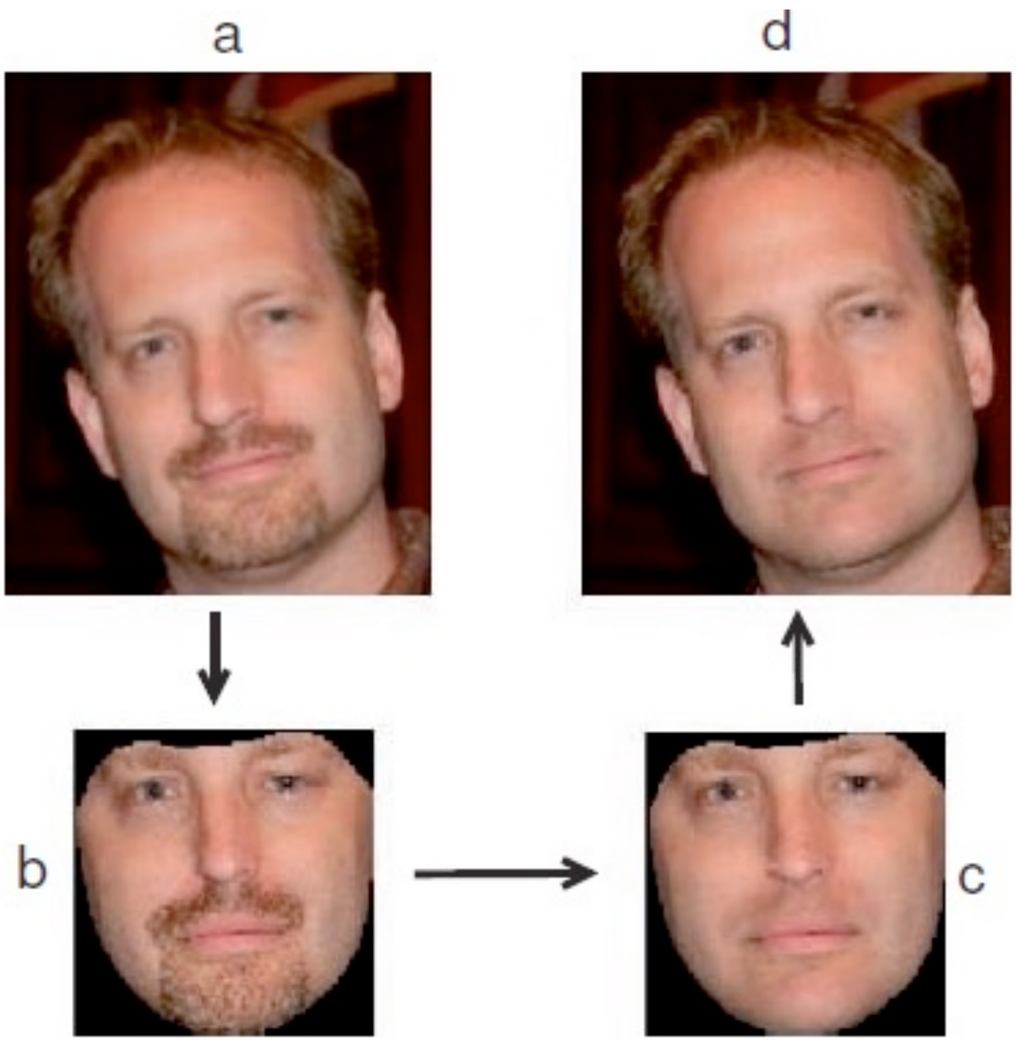

Nguyen, MH, J-F Lalonde, AA Efros, F de la Torre. "Image-based shaving." In *Computer graphics forum*, vol. 27, no. 2, pp. 627-635, 2008.

Learning: Multimodality

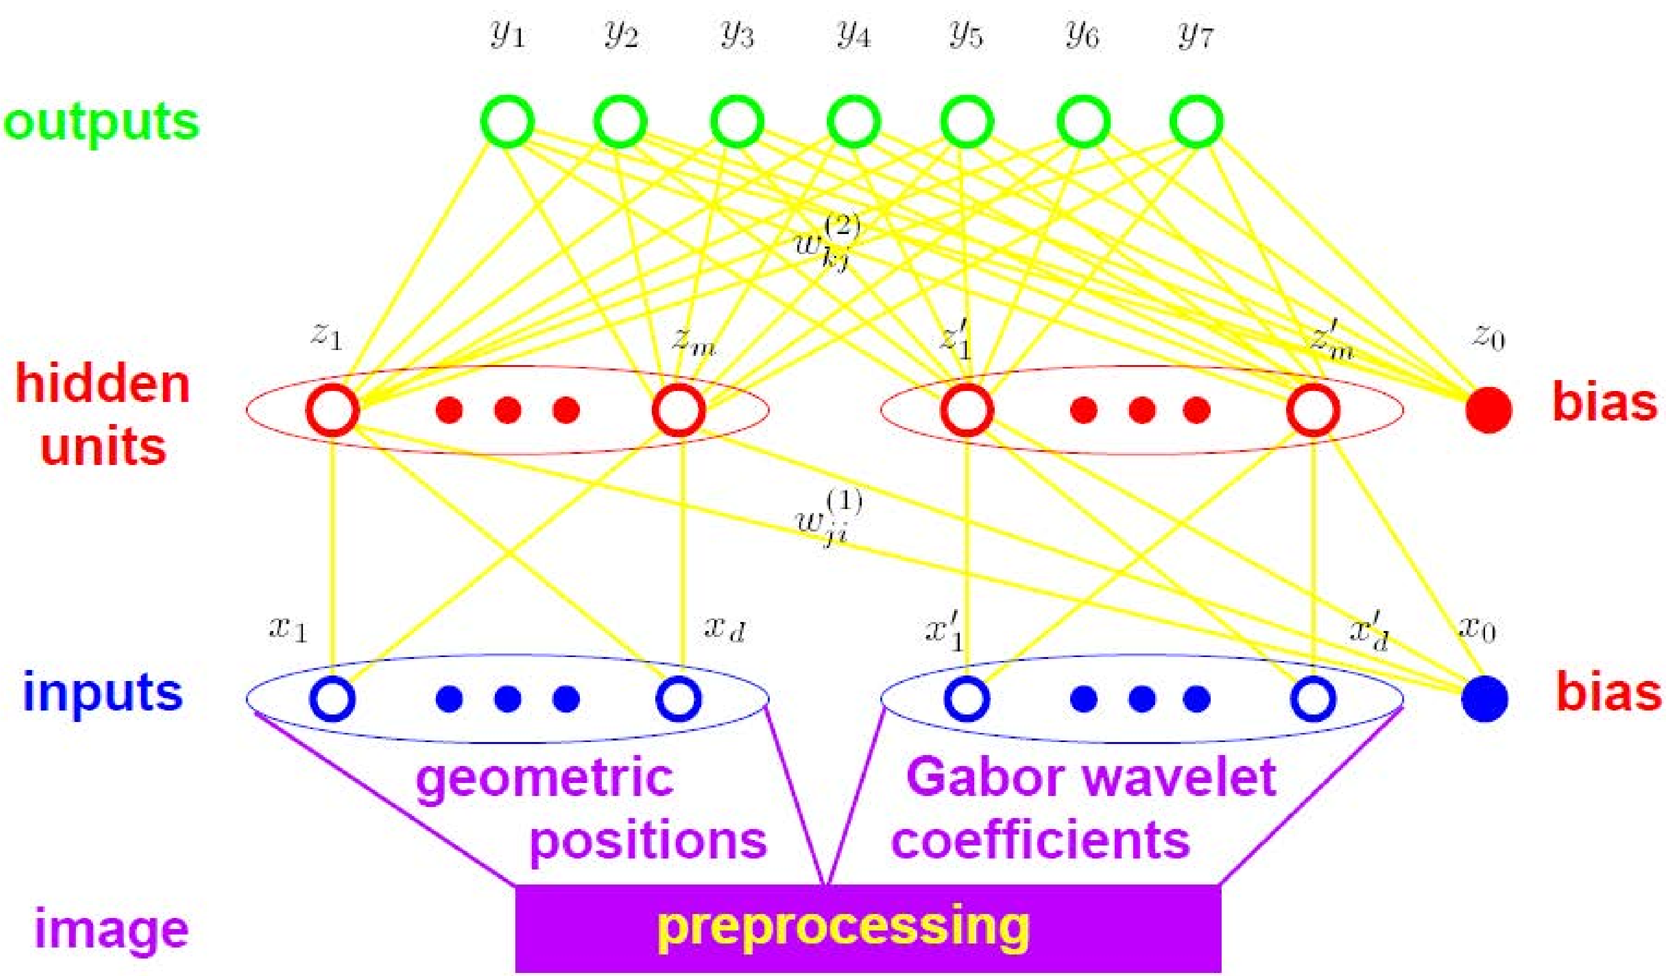

Zhang, Z., Lyons, M., Schuster, M., & Akamatsu, S. (1998). "Comparison between geometry-based and gabor-wavelets-based facial expression recognition using multi-layer perceptron". In *3rd IEEE AFGR*

Learning: Multimodality

*«Experiments show that facial expression recognition is mainly a **low frequency process**, and a spatial resolution of 64 pixels x 64 pixels or lower is probably enough.»*

*«It turns out that **five to seven hidden units** are probably enough to represent the space of feature expressions.»*

Zhang, Z., Lyons, M., Schuster, M., & Akamatsu, S. (1998). "Comparison between geometry-based and gabor-wavelets-based facial expression recognition using multi-layer perceptron". In *3rd IEEE AFGR*

Learning: Multimodality

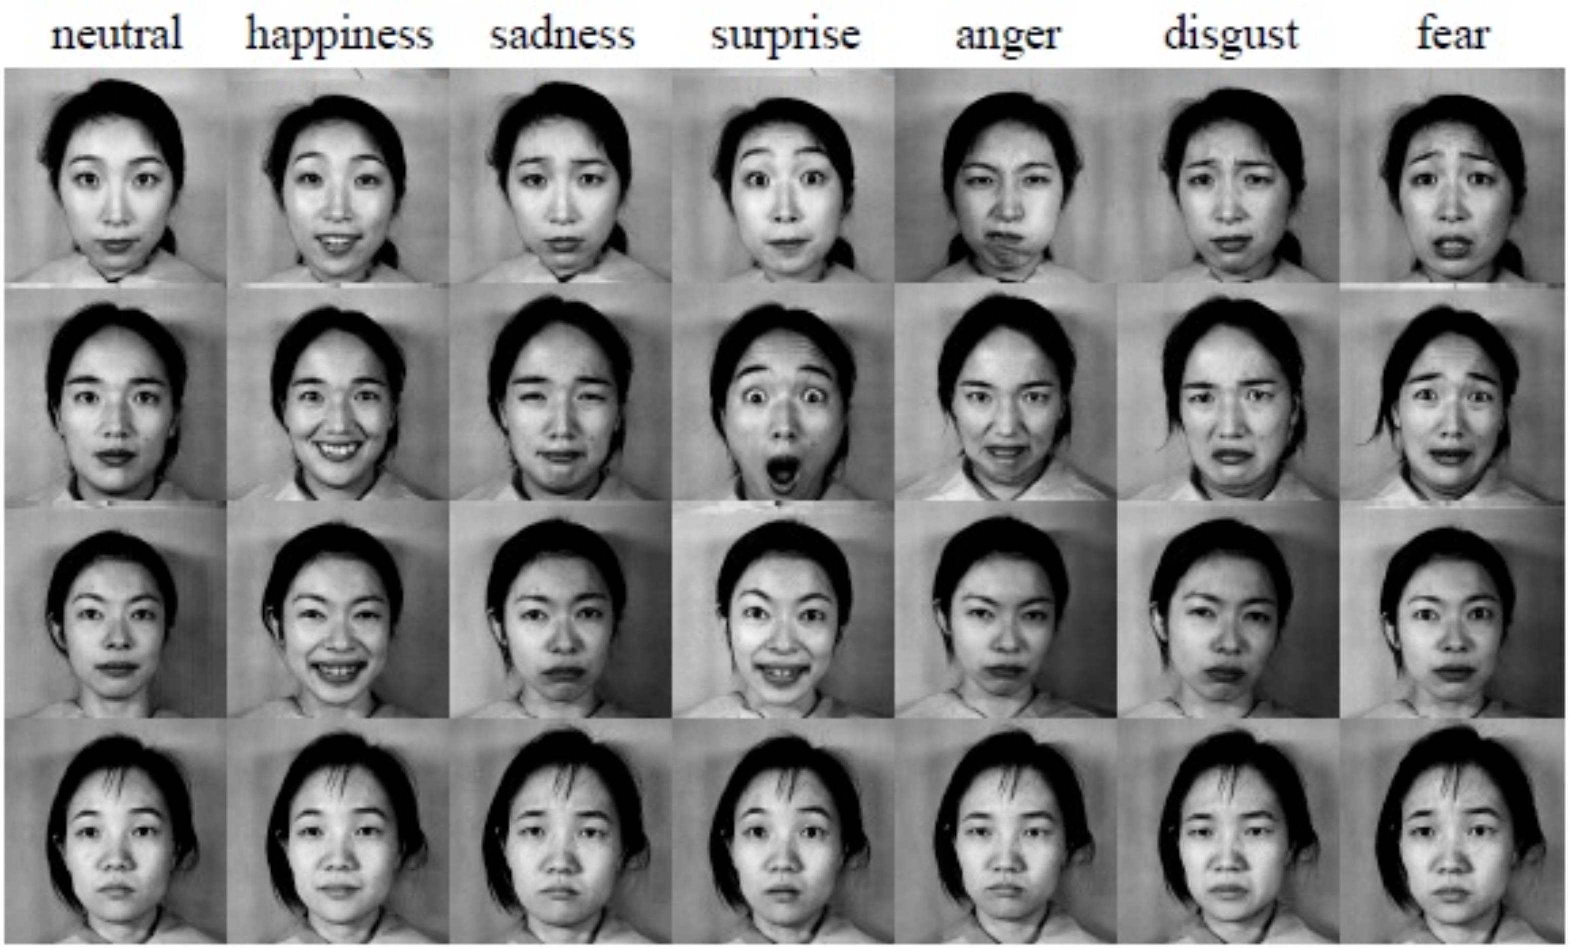

Zhang, Z., Lyons, M., Schuster, M., & Akamatsu, S. (1998). "Comparison between geometry-based and gabor-wavelets-based facial expression recognition using multi-layer perceptron". In *3rd IEEE AFGR*

Learning: Multimodality

Multiple representations of the same modality are also useful...

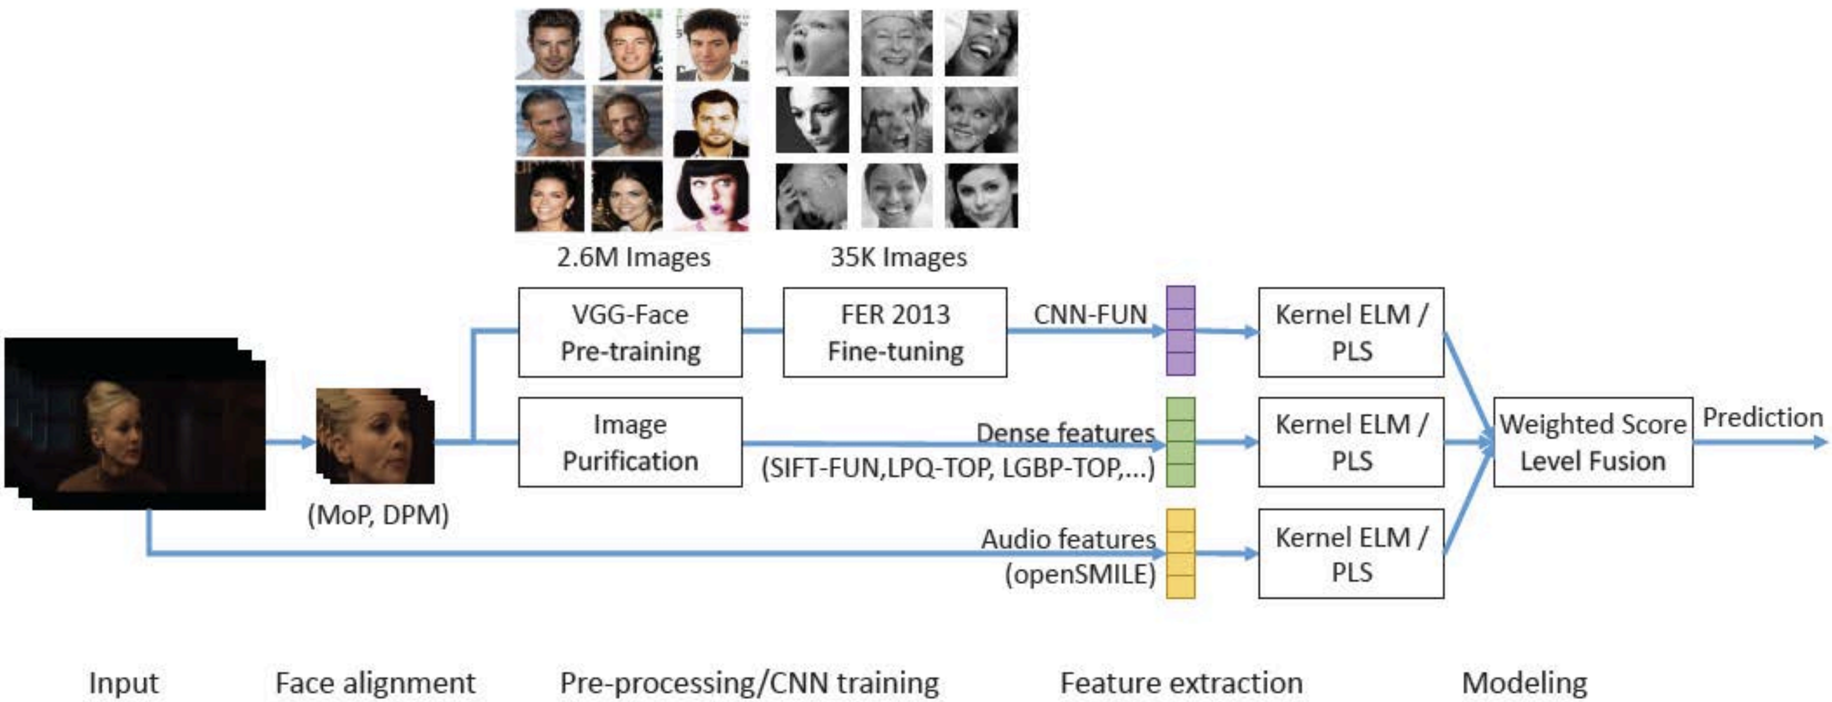

Kaya, H., F. Gürpınar, A.A., Salah "Video Based Emotion Recognition in the Wild using Deep Transfer Learning and Score Fusion," *Image and Vision Computing*, vol.65, pp. 66-75, 2017.

Learning: Analysis by synthesis

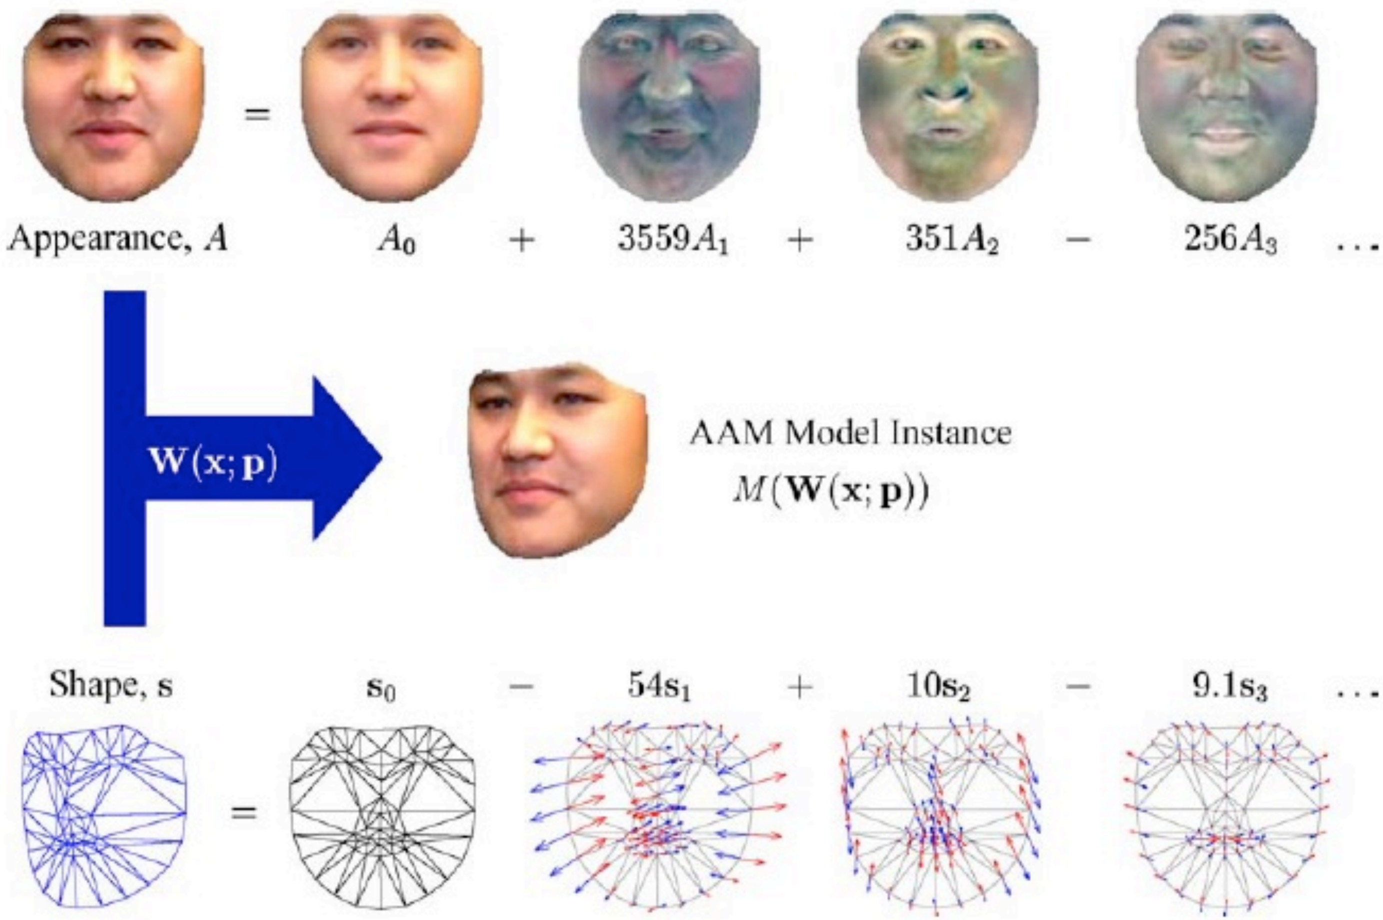

Cootes, T., Edwards, G., Taylor, C. (2001) "Active appearance models". TPAMI 23, 681-685.
Matthews, I., & Baker, S. (2004) "Active appearance models revisited". IJCV 60(2), 135-164.

Learning: Analysis by synthesis

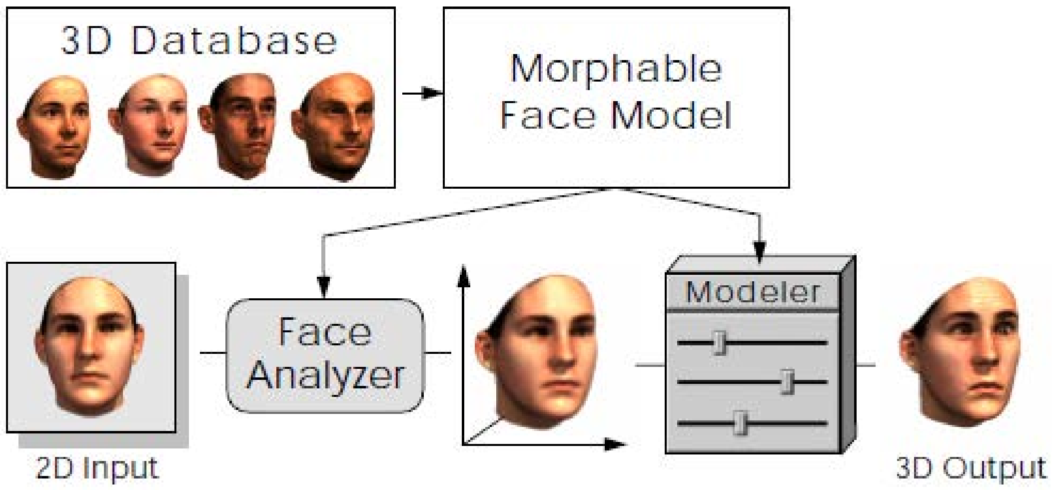

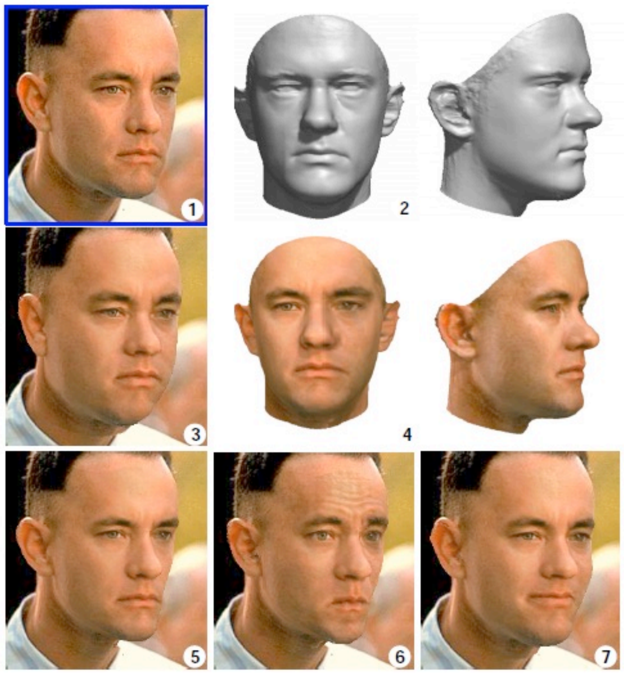

Figure 6: Matching a morphable model to a single image (1) of a face results in a 3D shape (2) and a texture map estimate. The texture estimate can be improved by additional texture extraction (4). The 3D model is rendered back into the image after changing facial attributes, such as gaining (3) and loosing weight (5), frowning (6), or being forced to smile (7).

$$\mathbf{S}_{mod} = \sum_{i=1}^{m} a_i \mathbf{S}_i \,, \quad \mathbf{T}_{mod} = \sum_{i=1}^{m} b_i \mathbf{T}_i \,, \quad \sum_{i=1}^{m} a_i = \sum_{i=1}^{m} b_i = 1.$$

Blanz, V., & Vetter, T. (1999, July). “A morphable model for the synthesis of 3D faces”. In *Siggraph* (Vol. 99, No. 1999, pp. 187-194).

Learning: Analysis by synthesis

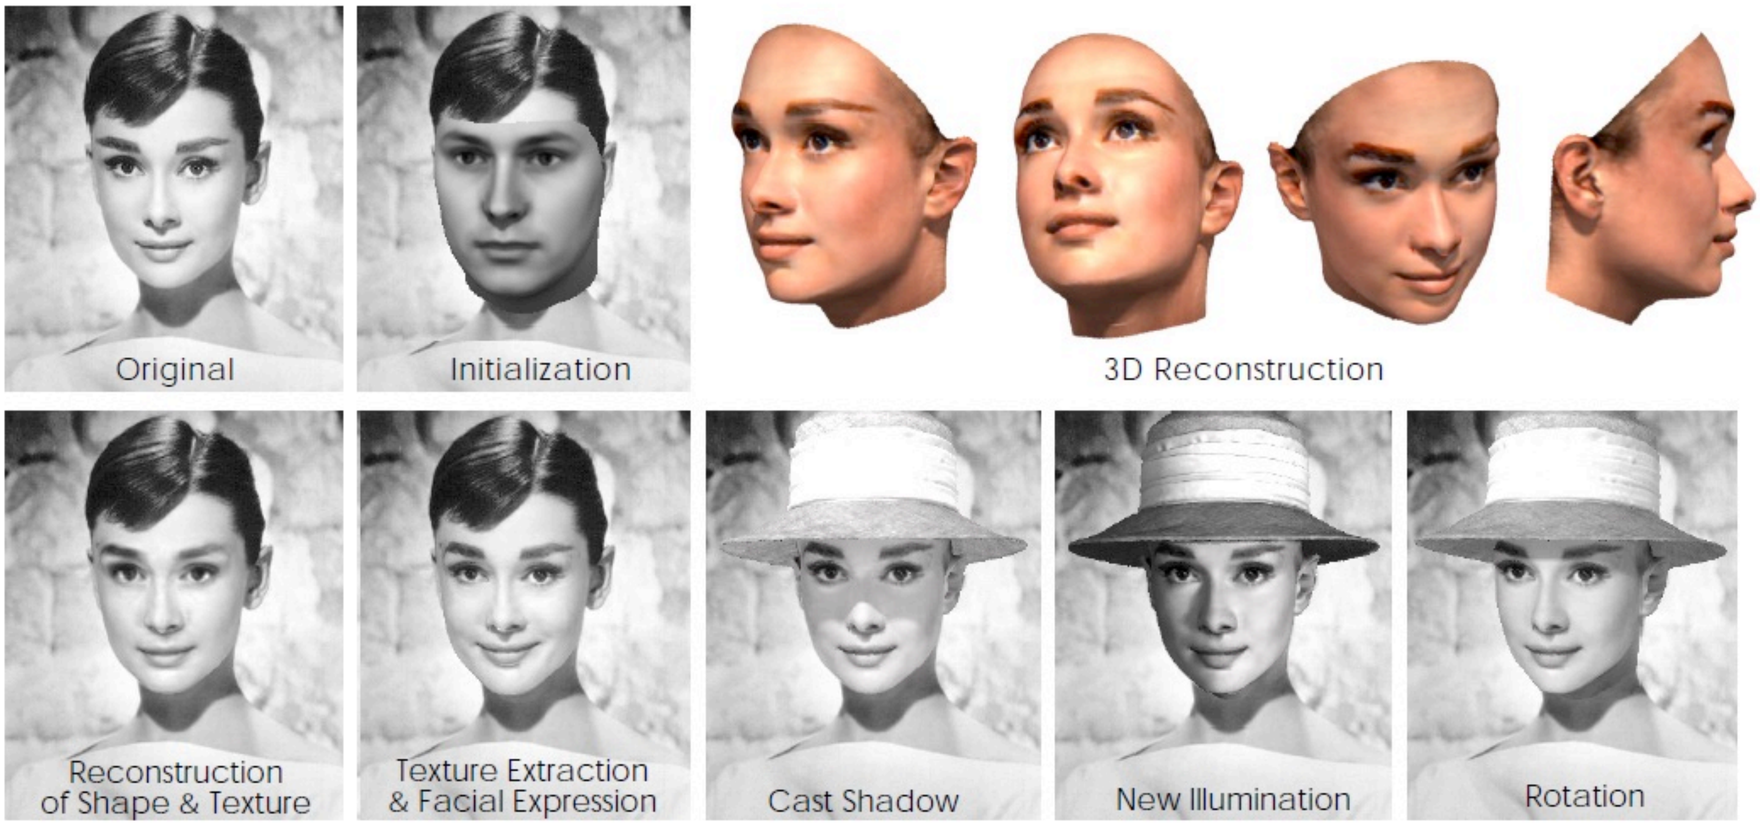

Blanz, V., & Vetter, T. (1999, July). "A morphable model for the synthesis of 3D faces". In *Siggraph* (Vol. 99, No. 1999, pp. 187-194).

Learning: Analysis by synthesis

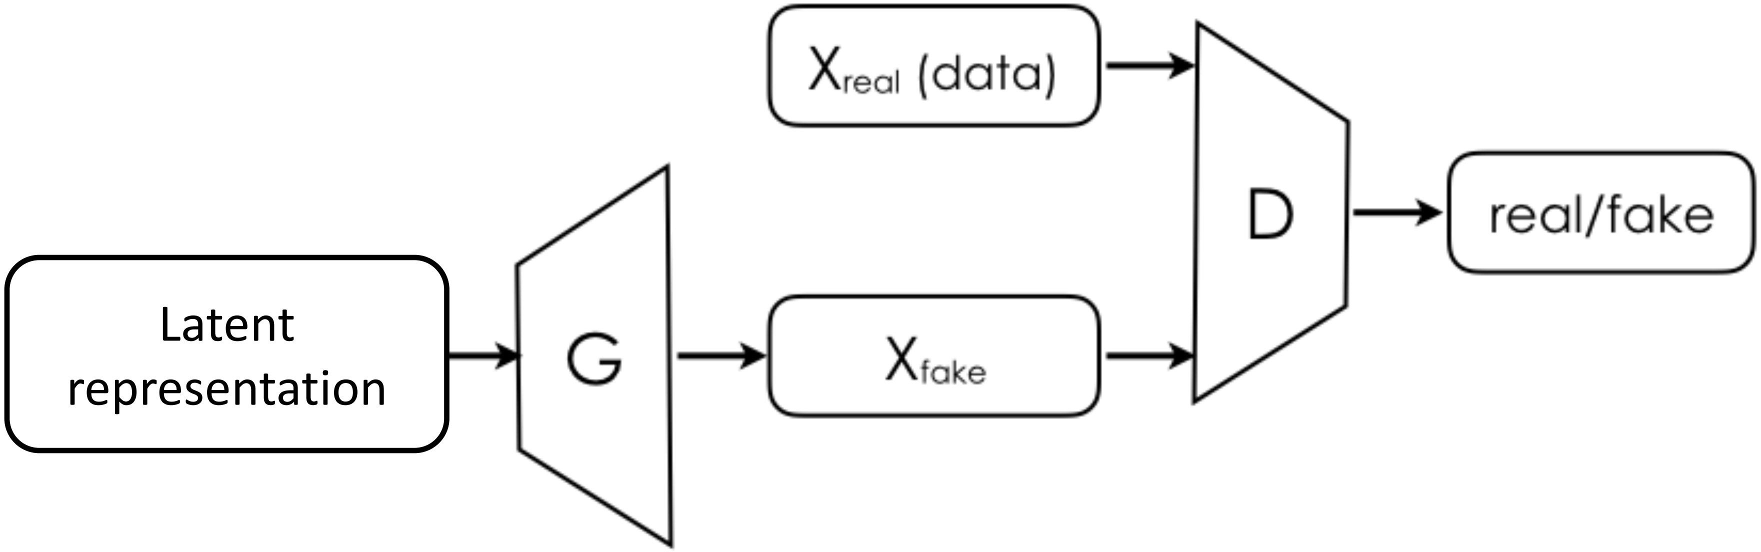

Generative adversarial networks

Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A. and Bengio, Y., 2014. “Generative adversarial nets”. In NeurIPS (pp. 2672-2680)”

Learning: Analysis by synthesis

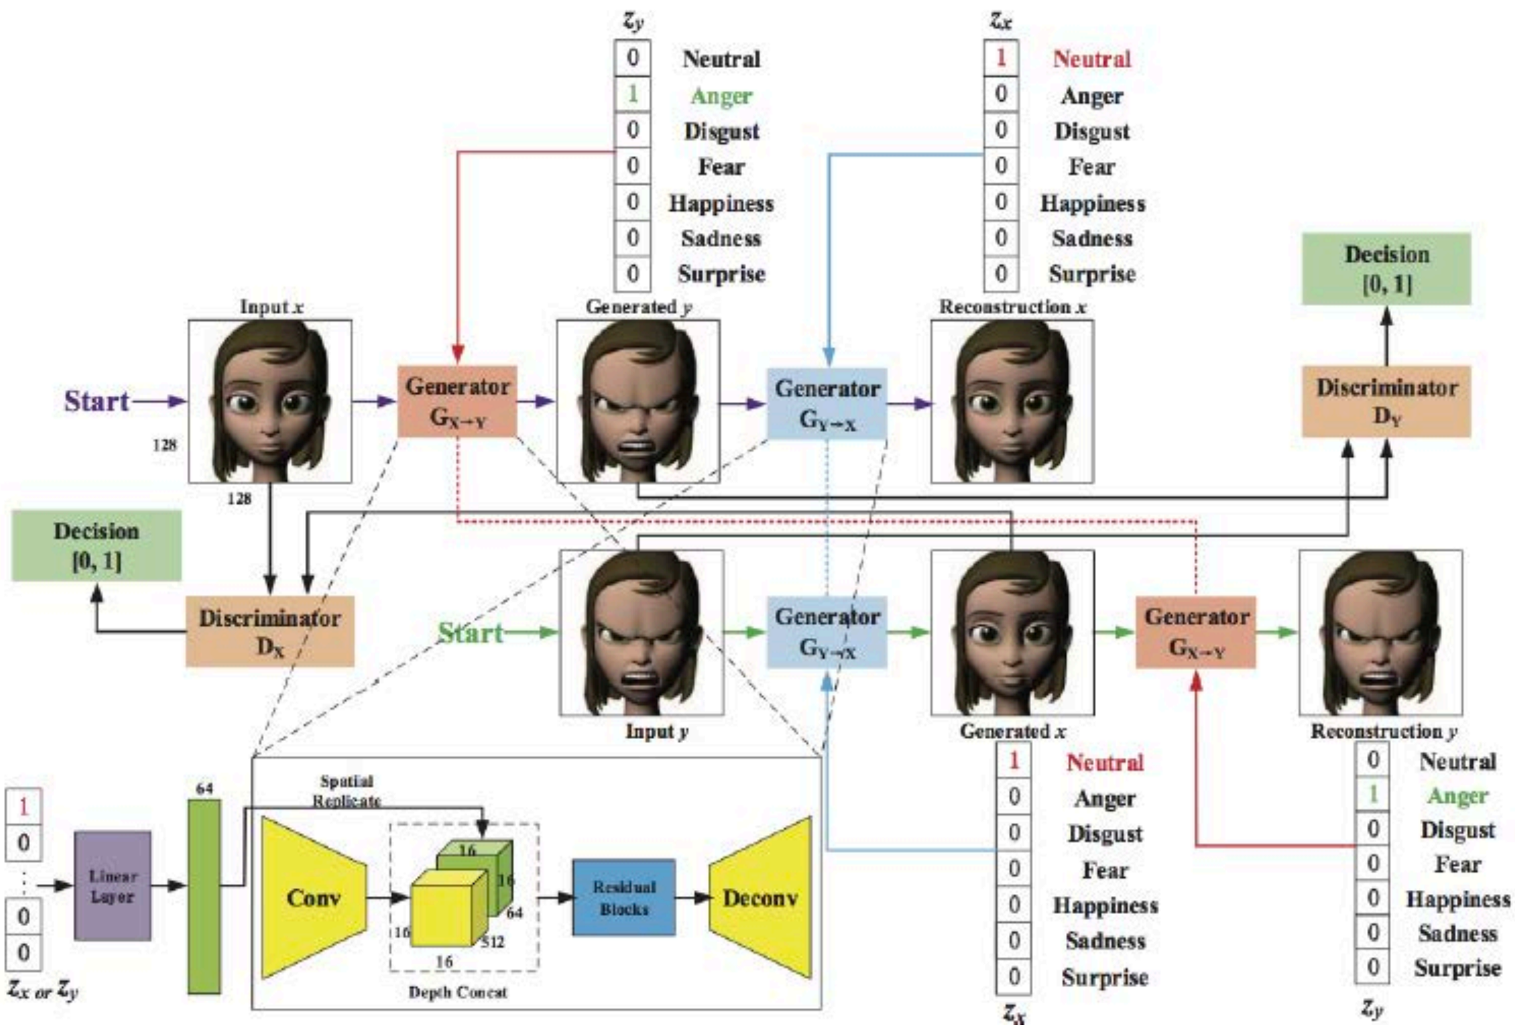

Tang, H., Wang, W., Wu, S., Chen, X., Xu, D., Sebe, N., & Yan, Y. (2019). “Expression Conditional GAN for Facial Expression-to-Expression Translation”. arXiv preprint arXiv:1905.05416, ICIP’19.

Learning: Analysis by synthesis

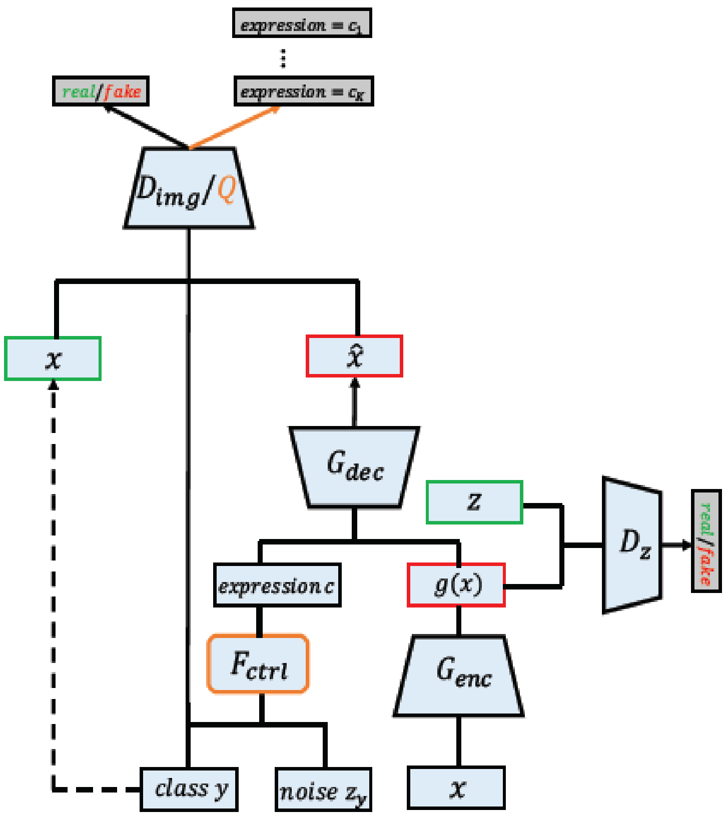

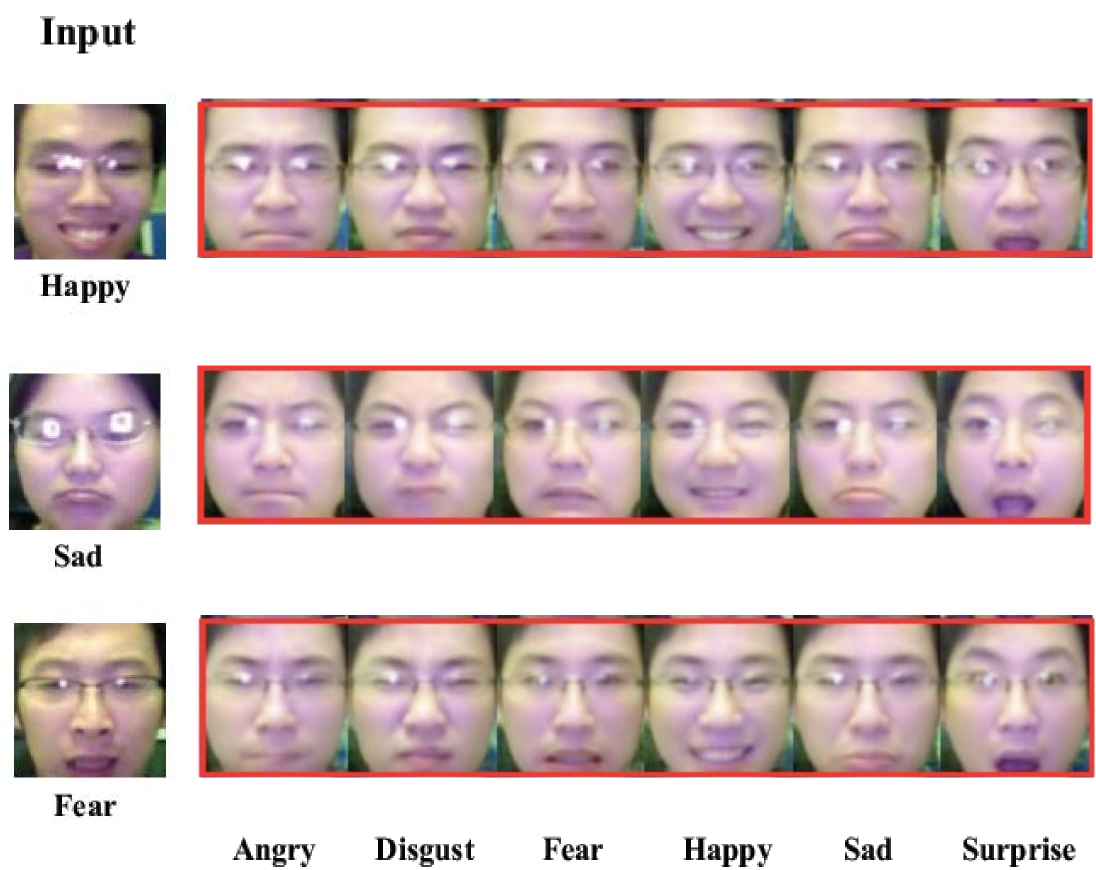

Ding, H., Sricharan, K., & Chellappa, R. (2018, April). “ExprGAN: Facial expression editing with controllable expression intensity”. In *Thirty-Second AAAI Conference on Artificial Intelligence*.

Dynamics: Phases of an expression

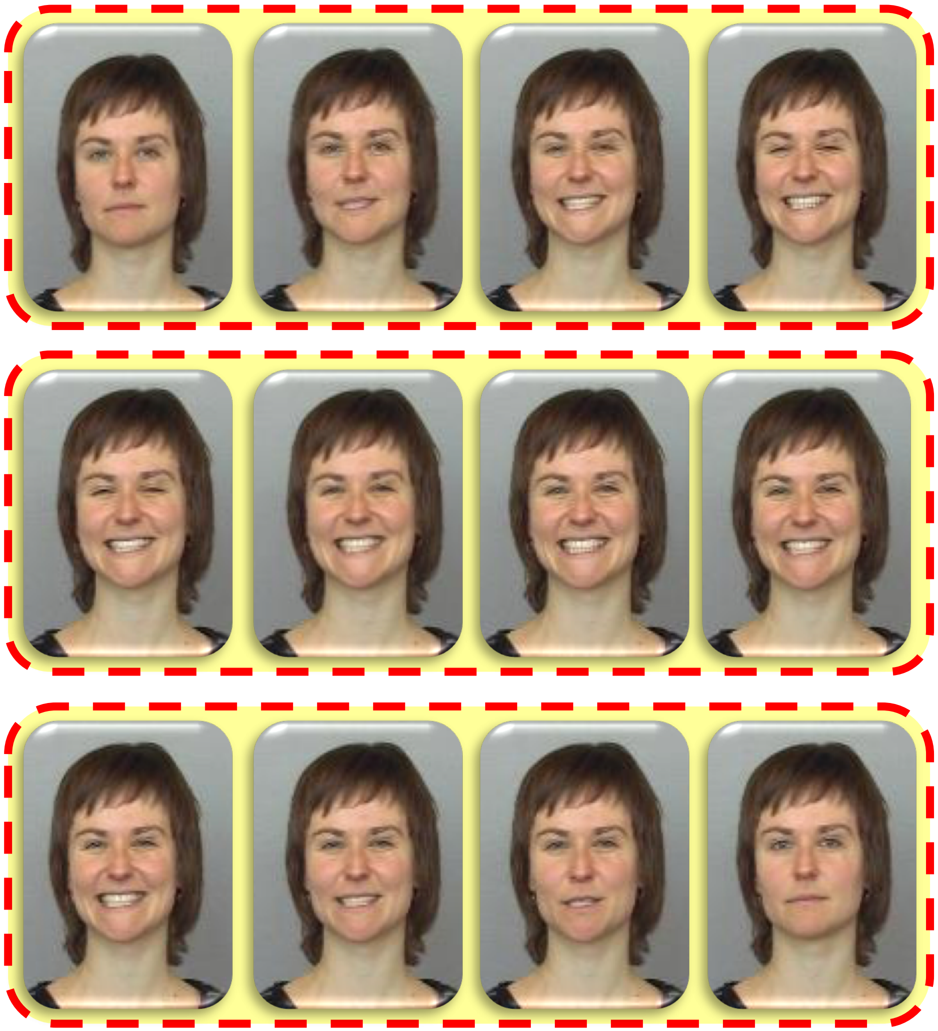

- A facial expression is composed of three main phases:
  - **Onset**: Neutral state to expressive face
  - **Apex**: Stable period of the expressive face
  - **Offset**: Expressive state to neutral face

Dynamics: From static to dynamic

- Extract spatio-temporal features and classify
- Classify at the frame level and combine later
- Model dynamics and do sequence level classification

Dynamics: Spatiotemporal features

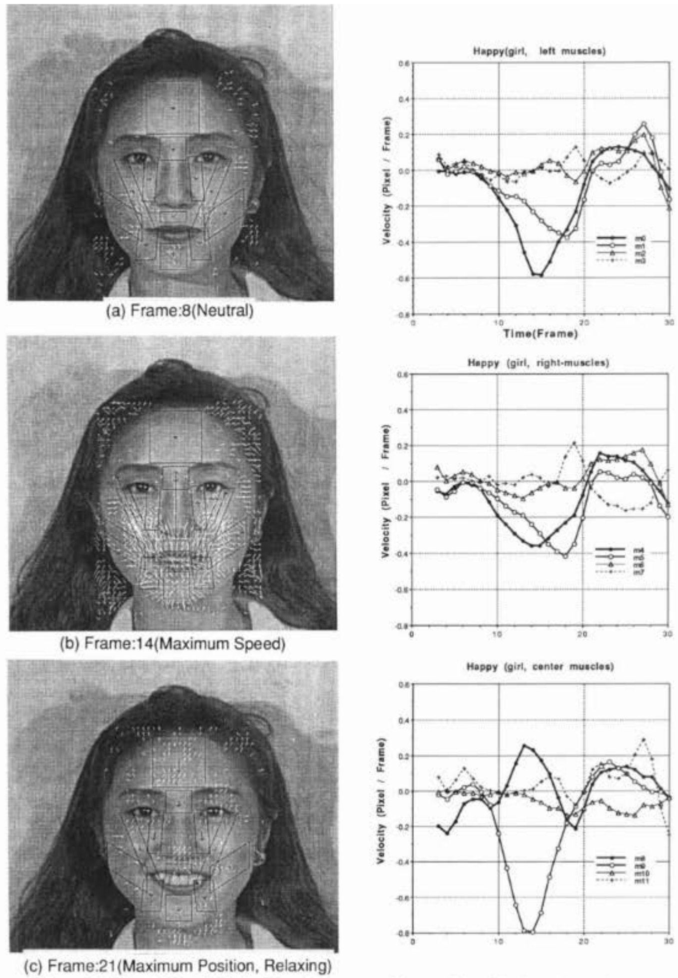

Figure 3: Estimated Muscle Motion(Happy) Figure 4: Plot of Estimated Motion(Happy)

- Optical flow based features
- 15-dimensional vector
- Focus on the apex of the expression
- No variances

Kenji Mase, “An Application of Optical Flow - Extraction of Facial Expression”, Proceedings of IAPR Workshop on Machine Vision Applications, November 28-30, 1990, Kokubunji, Tokyo, Japan

Dynamics: Space-time tubes

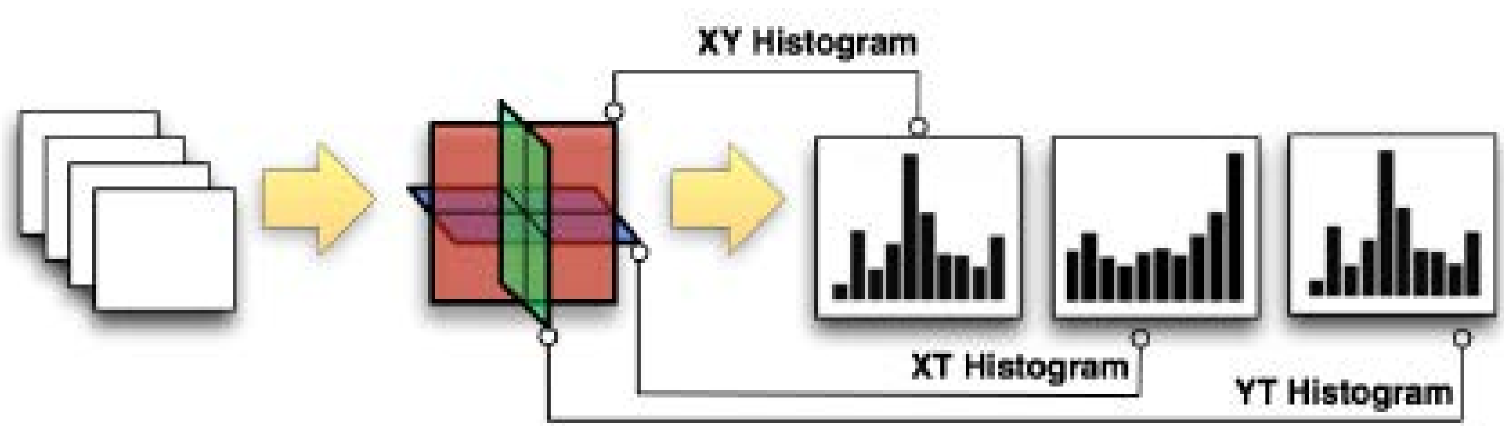

Almaev, T. R., & Valstar, M. F. (2013). “Local gabor binary patterns from three orthogonal planes for automatic facial expression recognition”. In *Humaine Association Conference on Affective Computing and Intelligent Interaction* (pp. 356-361). IEEE.

Dynamics: Combining frame-level outputs

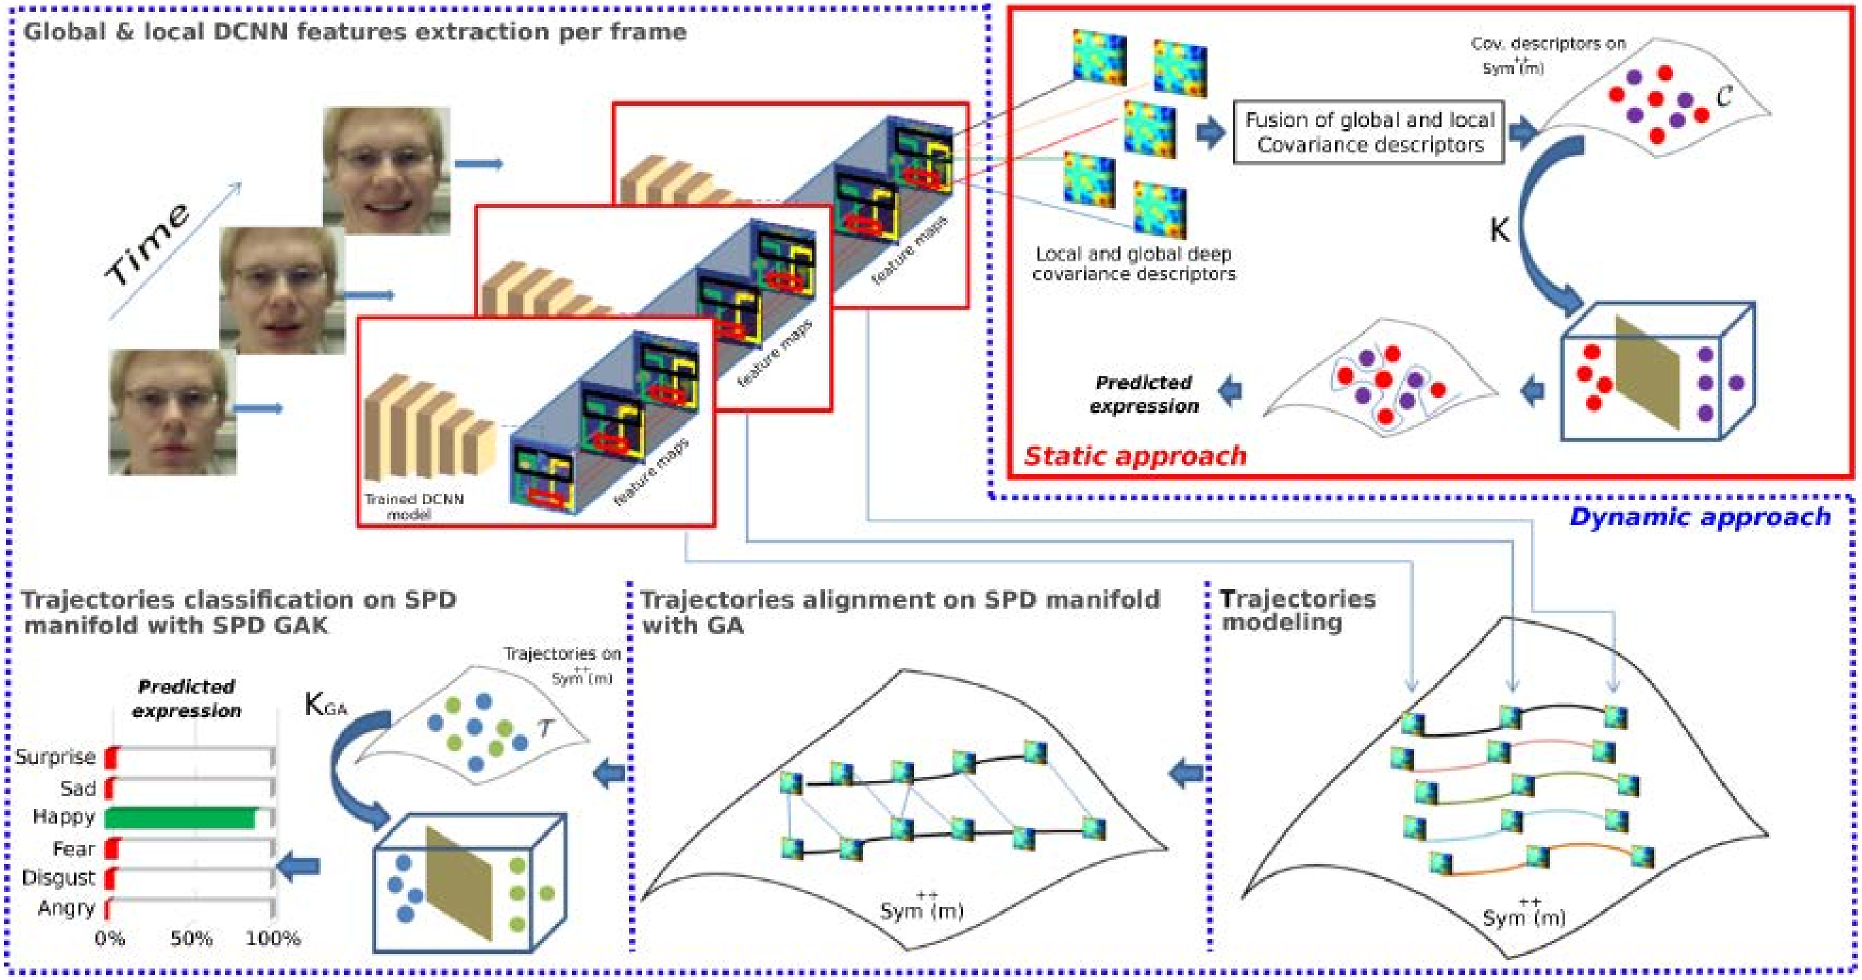

Otberdout, N., Kacem, A., Daoudi, M., Ballihi, L., & Berretti, S. (2018). “Automatic Analysis of Facial Expressions Based on Deep Covariance Trajectories”. *arXiv preprint arXiv:1810.11392*.

Dynamics: Sequence-level classification

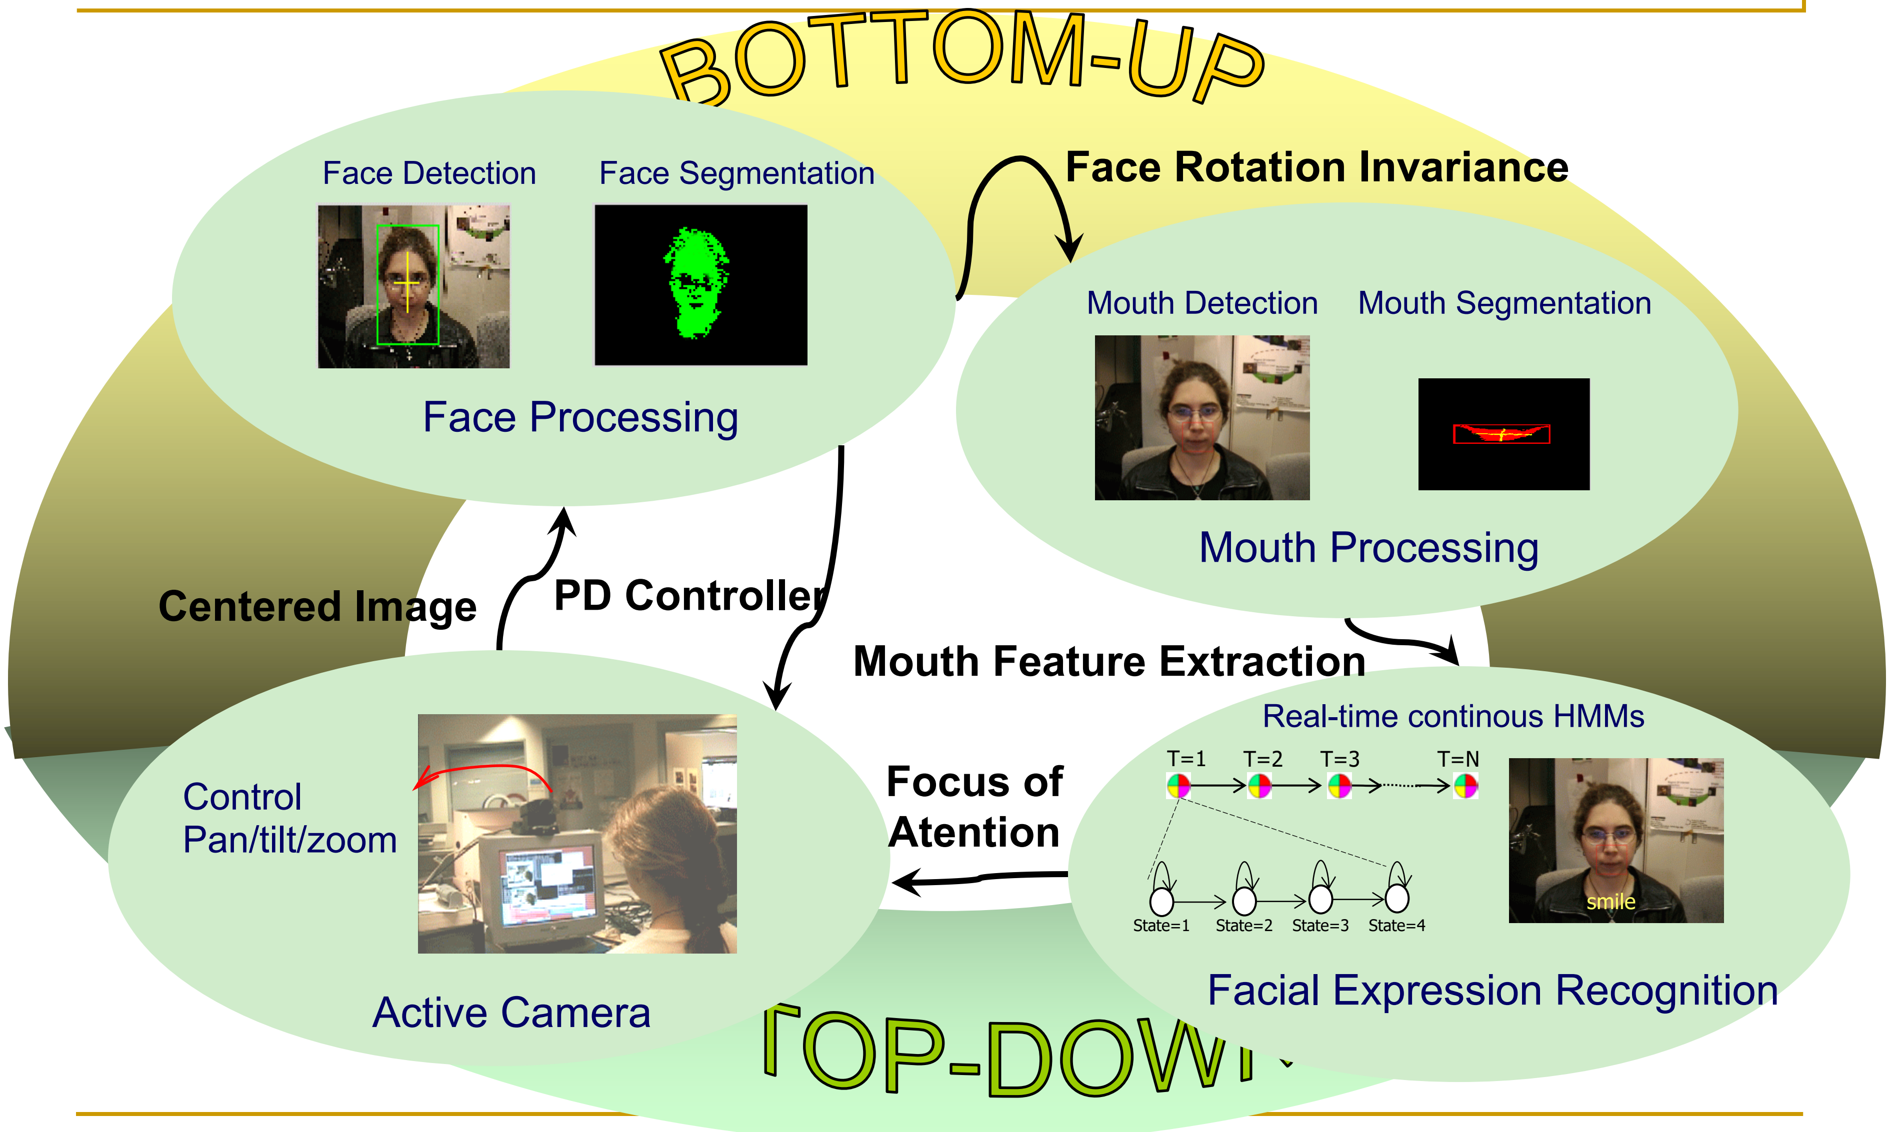

Oliver, N., Pentland, A. and Berard, F., "LAFTER:A Real-time Lips and Face Tracker with Facial Expression Recognition", Proc. CVPR, 1997

Early work: Hidden Markov Models

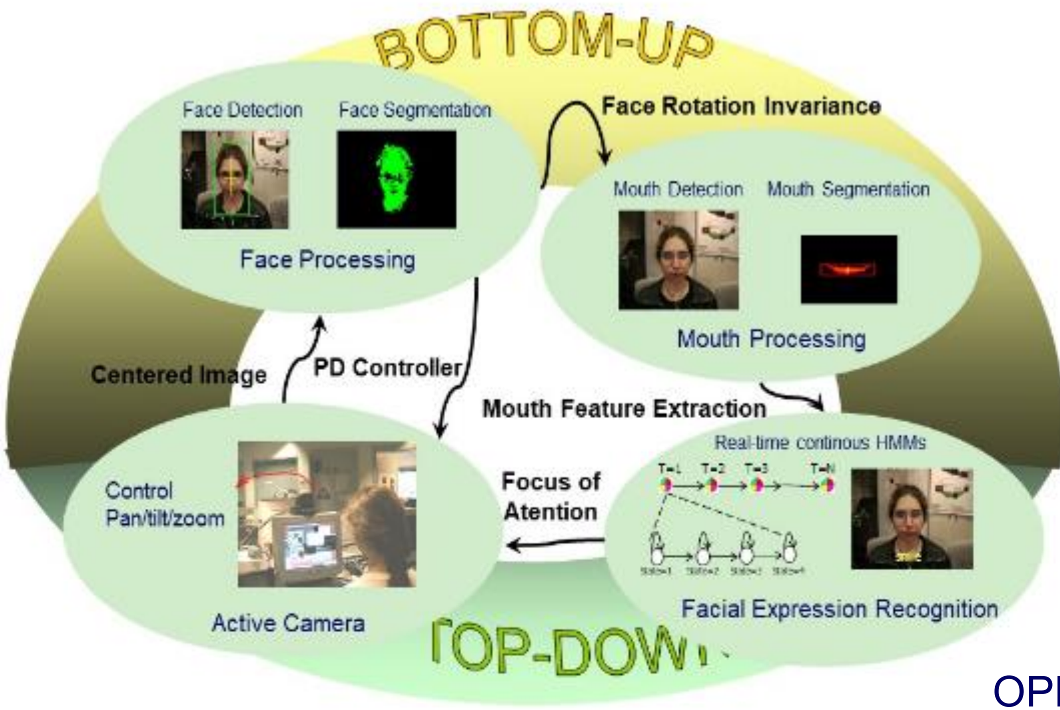

- Basic expressions
- Real-time
- Dynamic modeling
- **~95+%** accuracy

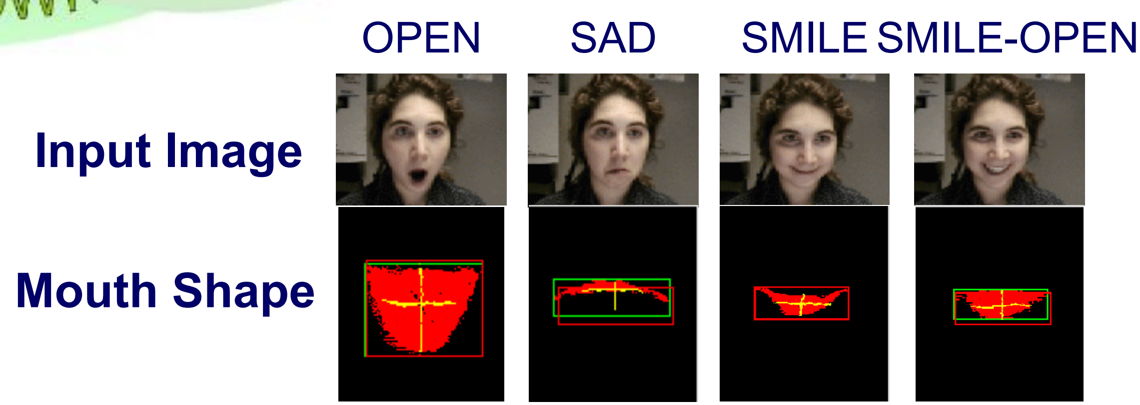

Oliver, N., Pentland, A. and Berard, F., "LAFTER:A Real-time Lips and Face Tracker with Facial Expression Recognition", Proc. CVPR, 1997 and Pattern Recognition 2000

Dynamics: Recurrent neural networks

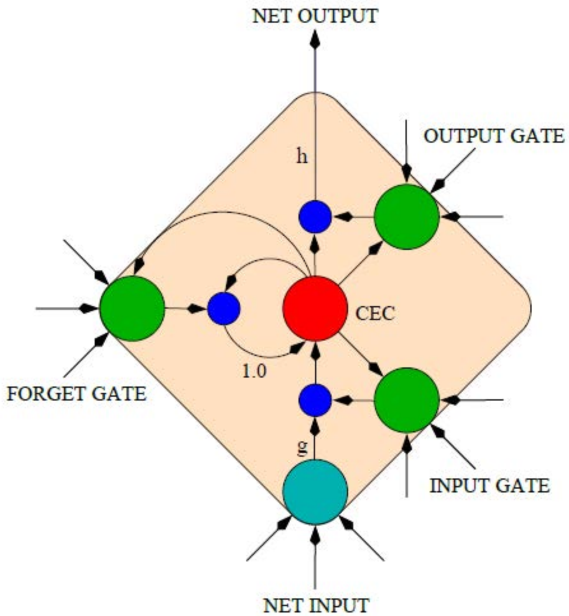

An LSTM node

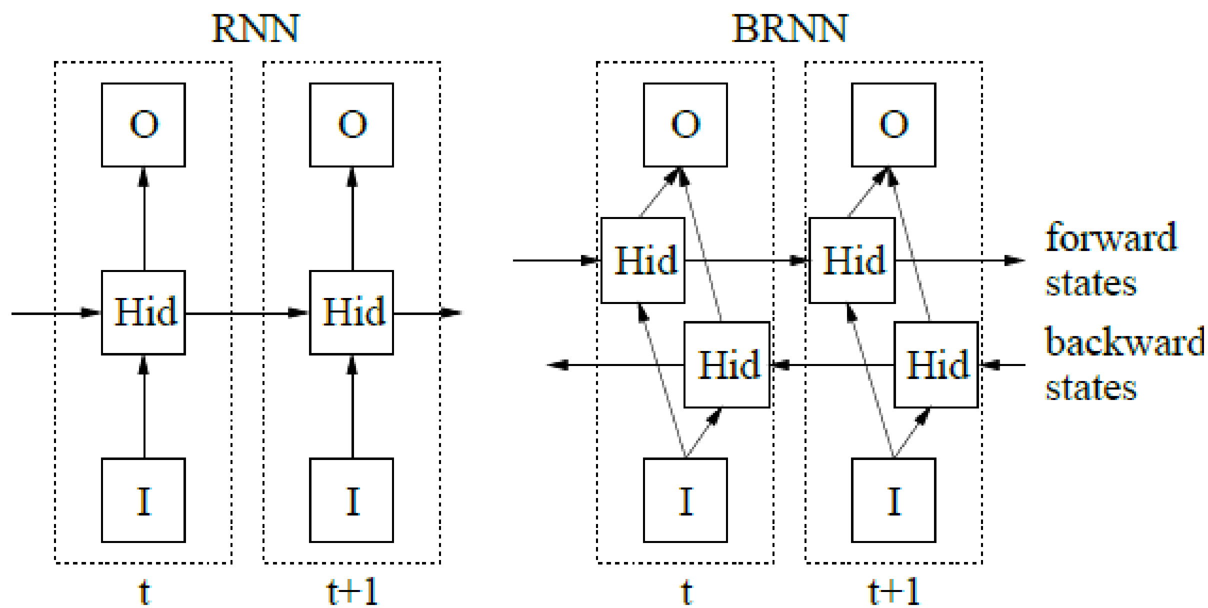

Bidirectional RNN for offline recognition

Graves, A., Mayer, C., Wimmer, M., Schmidhuber, J., & Radig, B. (2008). "Facial expression recognition with recurrent neural networks". In *Proceedings of the International Workshop on Cognition for Technical Systems*.

Dynamics: Recurrent neural networks

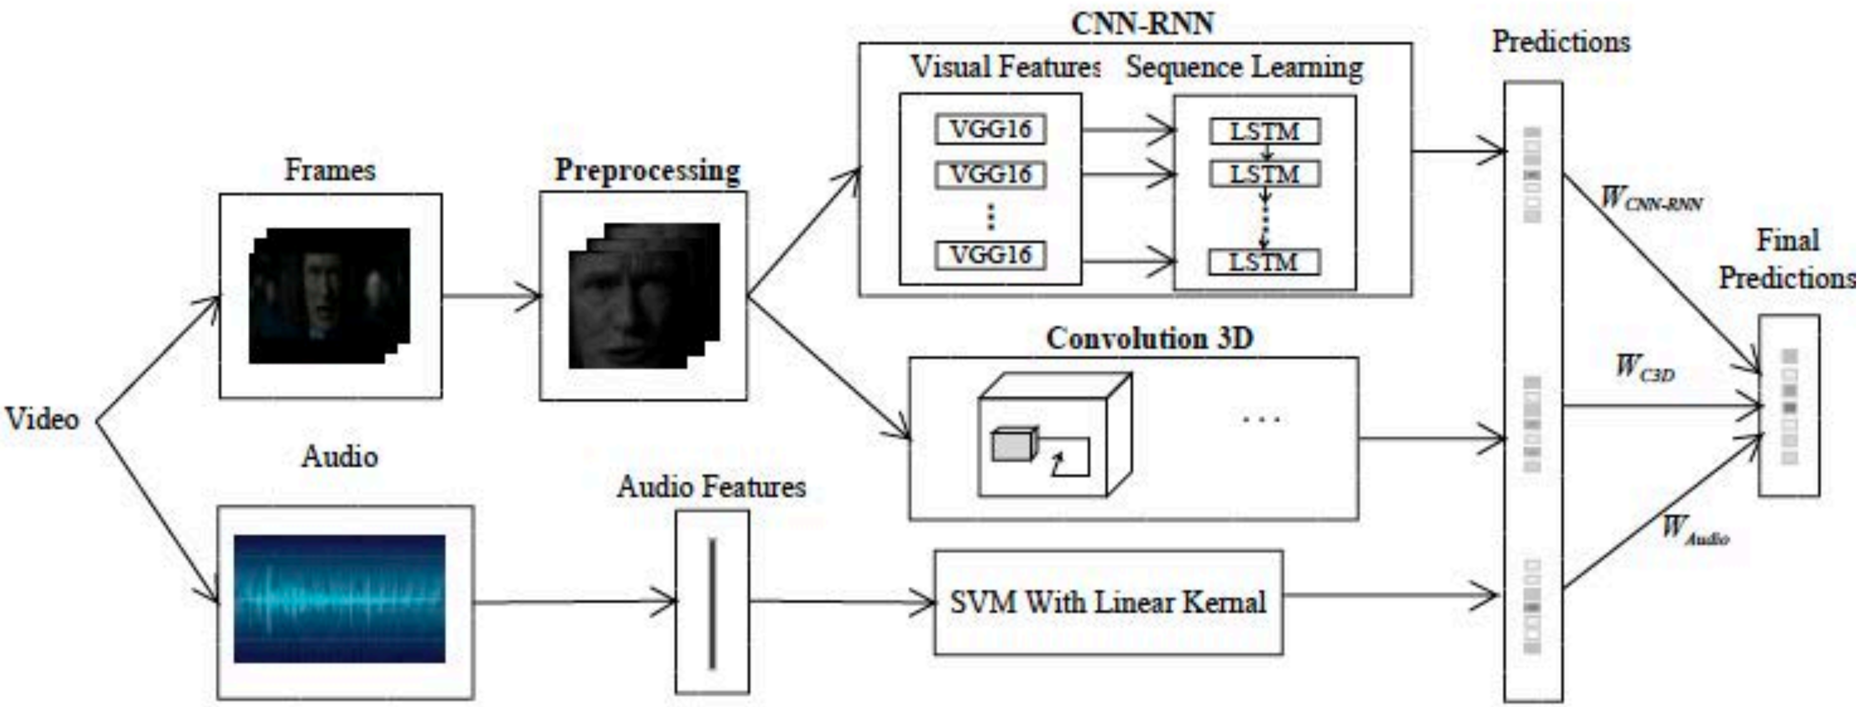

Fan, Y., Lu, X., Li, D., & Liu, Y. (2016). “Video-based emotion recognition using CNN-RNN and C3D hybrid networks”. In Proc. 18th ACM ICMI (pp. 445-450). ACM.

Dynamics: Benchmark Datasets

Dataset	Footage	Year	size	Facial expressions
Cohn-Kanade	Posed video	2000	210 adults, 480 videos	6 basic emotions + AUs (FACS)
MMI	Frontal/profile videos, induced emotion	2005	11 children, 18 adults, 1250 videos	6 basic emotions + AUs (FACS) and observer judgments
RU-FACS	Subjects under interview, Audiovisual	2005	100 adults	AUs (FACS)
UT-Dallas	Video, induced emotion	2006	229 adults	6 basic emotions, puzzle, laughter, boredom, disbelief
BU-4DFE	4D range data	2006	101 adults, 606 seq.	6 basic emotions, 4 levels of intensity
FABO	Facial exp. and body gesture jointly, posed	2006	23 adults, 210 videos	6 basic emotions, neutral, uncertainty, anxiety, boredom
UvA-NEMO	Induced emotion vs. posed, 50 fps video	2012	400 subj., 1240 videos	Spontaneous vs. posed smile
AM-FED	Induced emotions	2013	242 videos	14 AUs (FACS)
DISFA+	Induced emotions	2016	27 adults, 120K frames	Posed and spontaneous, 12 AUs (FACS)
AFEW	Video clips from films	2017	600 clips	Continuous annotations

Summary: Individual behaviors

- Definition, representation and annotation issues for ‘classes’ and their boundaries
- Spatio-temporal alignment of the data
  - Robust preprocessing
  - Modeling dynamics
  - ‘Spotting’ behavior boundaries
- Modeling variances
  - Illumination-, pose-, sex-, identity-related variances
  - More subtle behavioral influences
- Dataset (and annotation) availability

HUMAN INTERACTIONS

Fundamental questions

- How do we **represent** the behavior?

Social signal processing

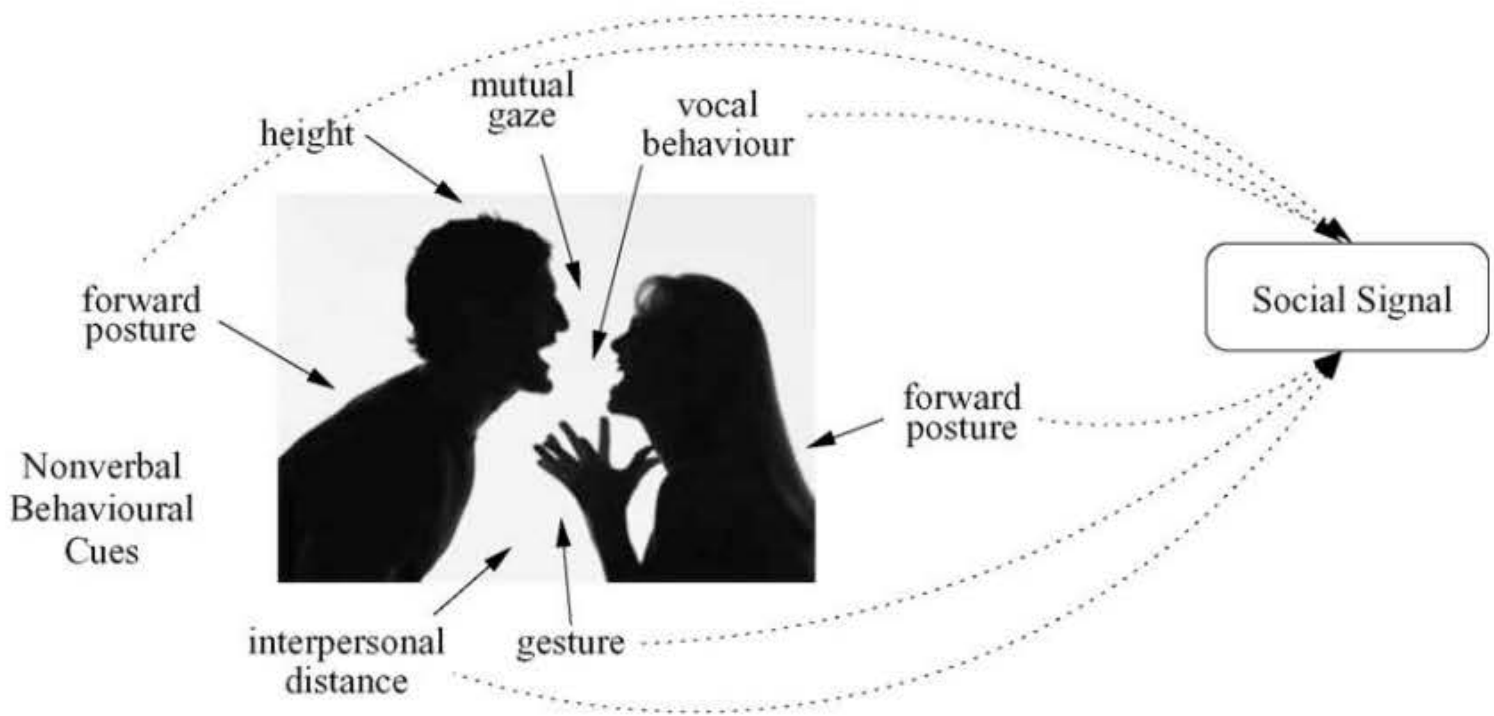

Vinciarelli, A., Pantic, M., & Bourlard, H. (2009). “Social signal processing: Survey of an emerging domain”. *Image and vision computing*, *27*(12), 1743-1759.

Salah, A. A., Pantic, M., & Vinciarelli, A. (2011). “Recent developments in social signal processing”. In *IEEE Int. Conf. Systems, Man, and Cybernetics (SMC)*

Simple pipeline for SSP

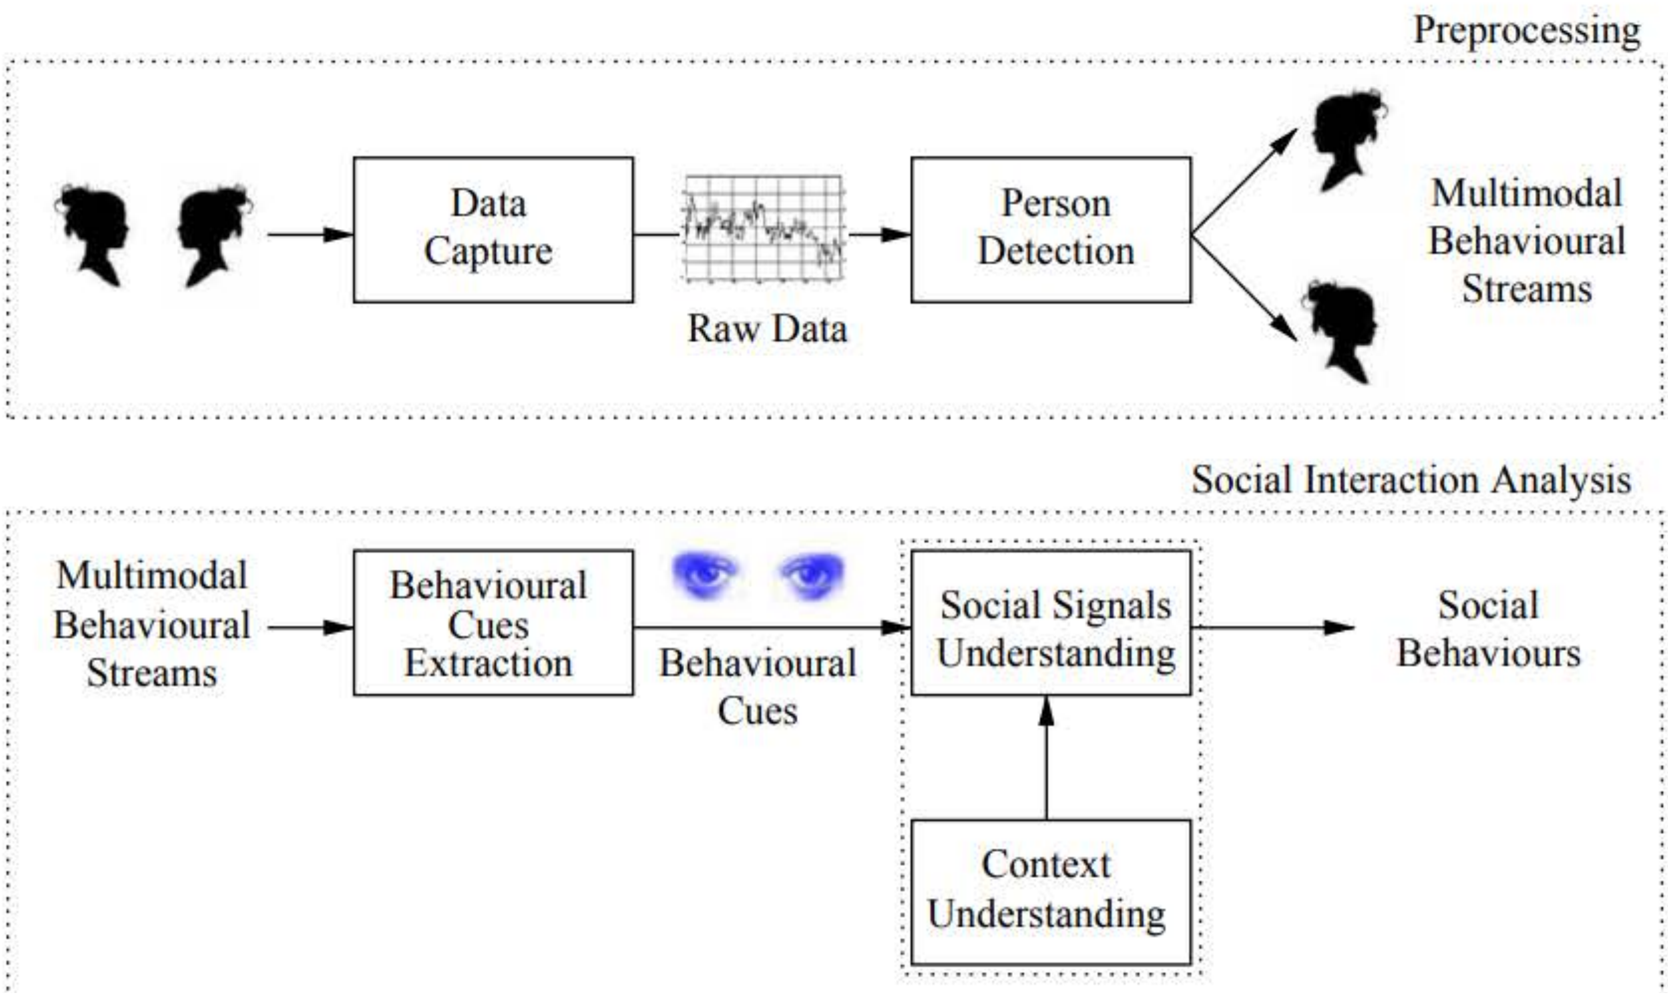

Vinciarelli, A., Pantic, M., & Bourlard, H. (2009). "Social signal processing: Survey of an emerging domain". *Image and vision computing*, *27*(12), 1743-1759.

Fundamental questions

- How do we represent the behavior?
- How do we establish **ground truth**?

Establishing the ground truth

- Before **2005** there were datasets for specific scenarios, such as meeting recordings
- For other types of scenarios researchers had to generate their own datasets
- Behavioral data streams:
  - Video data
  - In some occasions, multi-modal data (audio, wearable sensors)
- Manually annotated

Human-to-human interactions datasets

Dataset	Footage	Year	#Cl	#pp	Human to Human Int.
CASIA action	Outside recordings	2007	7	24	Rob, fight, follow, meet, meet and gather, overtake
BEHAVE	Outside recordings	2009	6	9	Approach, meet, fight, follow, run together, split, ignore, chase, walk together
i3DPost Multiview	Multiview images, 3D mesh models	2009	13 / 2	8	Handwave, handshake
UT-Interaction	Outside recordings (10s)	2009	6	8	Human to human inter.
Collective Activ	Outside recordings	2009	5+2	20+	Crossing, waiting, talking, walking, queuing + dancing, jogging
Hollywood2	Movies (10s)	2009	12 / 4	100 +	Handshake, hug, kiss, fight
TV Human Inter.	TV Shows (1-5s)	2010	4	100 +	Handshake, kiss, hug, high-five
HMDB51	Movies, YouTube, Google	2011	7	100 +	Fencing, hug, kick, kiss, punch, handshake, sword fight
BIT Interaction	Outside recordings (10s)	2012	8	8	Bow, box, handshake, hug, high-five, kick, pat, push
SBU Kinect	Lab recordings (1-5'') (color image, depth map,	2012	8	7	Approach, depart, push, hug, handshake, kick, punch, pass object

Human-to-human interactions datasets

Dataset	Footage	Year	#Classes	#people	Human to Human Int.
ChaLearn	Outside recordings	2015	235 / 5	14	Wave, point, handshake, hug, kiss, fight
CMU Panoptic	Lab recordings (10-15’) (images, 3D skeletons)	2015	6+	16	3 Social games, dance, toddler, office
SALSA	Inside recordings (30’) (multimodal)	2015	2	18	Poster presentation, cocktail party (F-formation)
ShakeFive2	Lab recordings (5-10’’)	2016	8	33	Handshake, hug, pass object
YouTube8M	YouTube videos	2016/ 2018	3862	100+	Few: hug, dance, kick,
AVA	Movies (15’)	2018	80 / 12	100+	Kiss, handshake, push, give object, play, fight, dance, talk, grab
Kinetics	YouTube videos (10’’)	2019	700 / few	100+	Punch person, kiss, hug, pass object, massage
Moments in time	YouTube videos (1-5’’)	2019	340 / few	100+	Hug, fencing
HACS	YouTube videos (2’’)	2019	201 / few	100+	Dance, getting tattoo

Singh, T., Vishwakarma, D.K., “Video benchmarks of human action datasets: a review”, Artificial Intelligence Review, Springer Nature, 2018

Stergiou, A., Poppe, R., “Analyzing human-interactions: a survey”, arXiv 2019

Fundamental questions

- How do we represent the behavior?
- How do we establish ground truth?
- What **machine learning models** are suitable to model interactions?

Interactions among few people

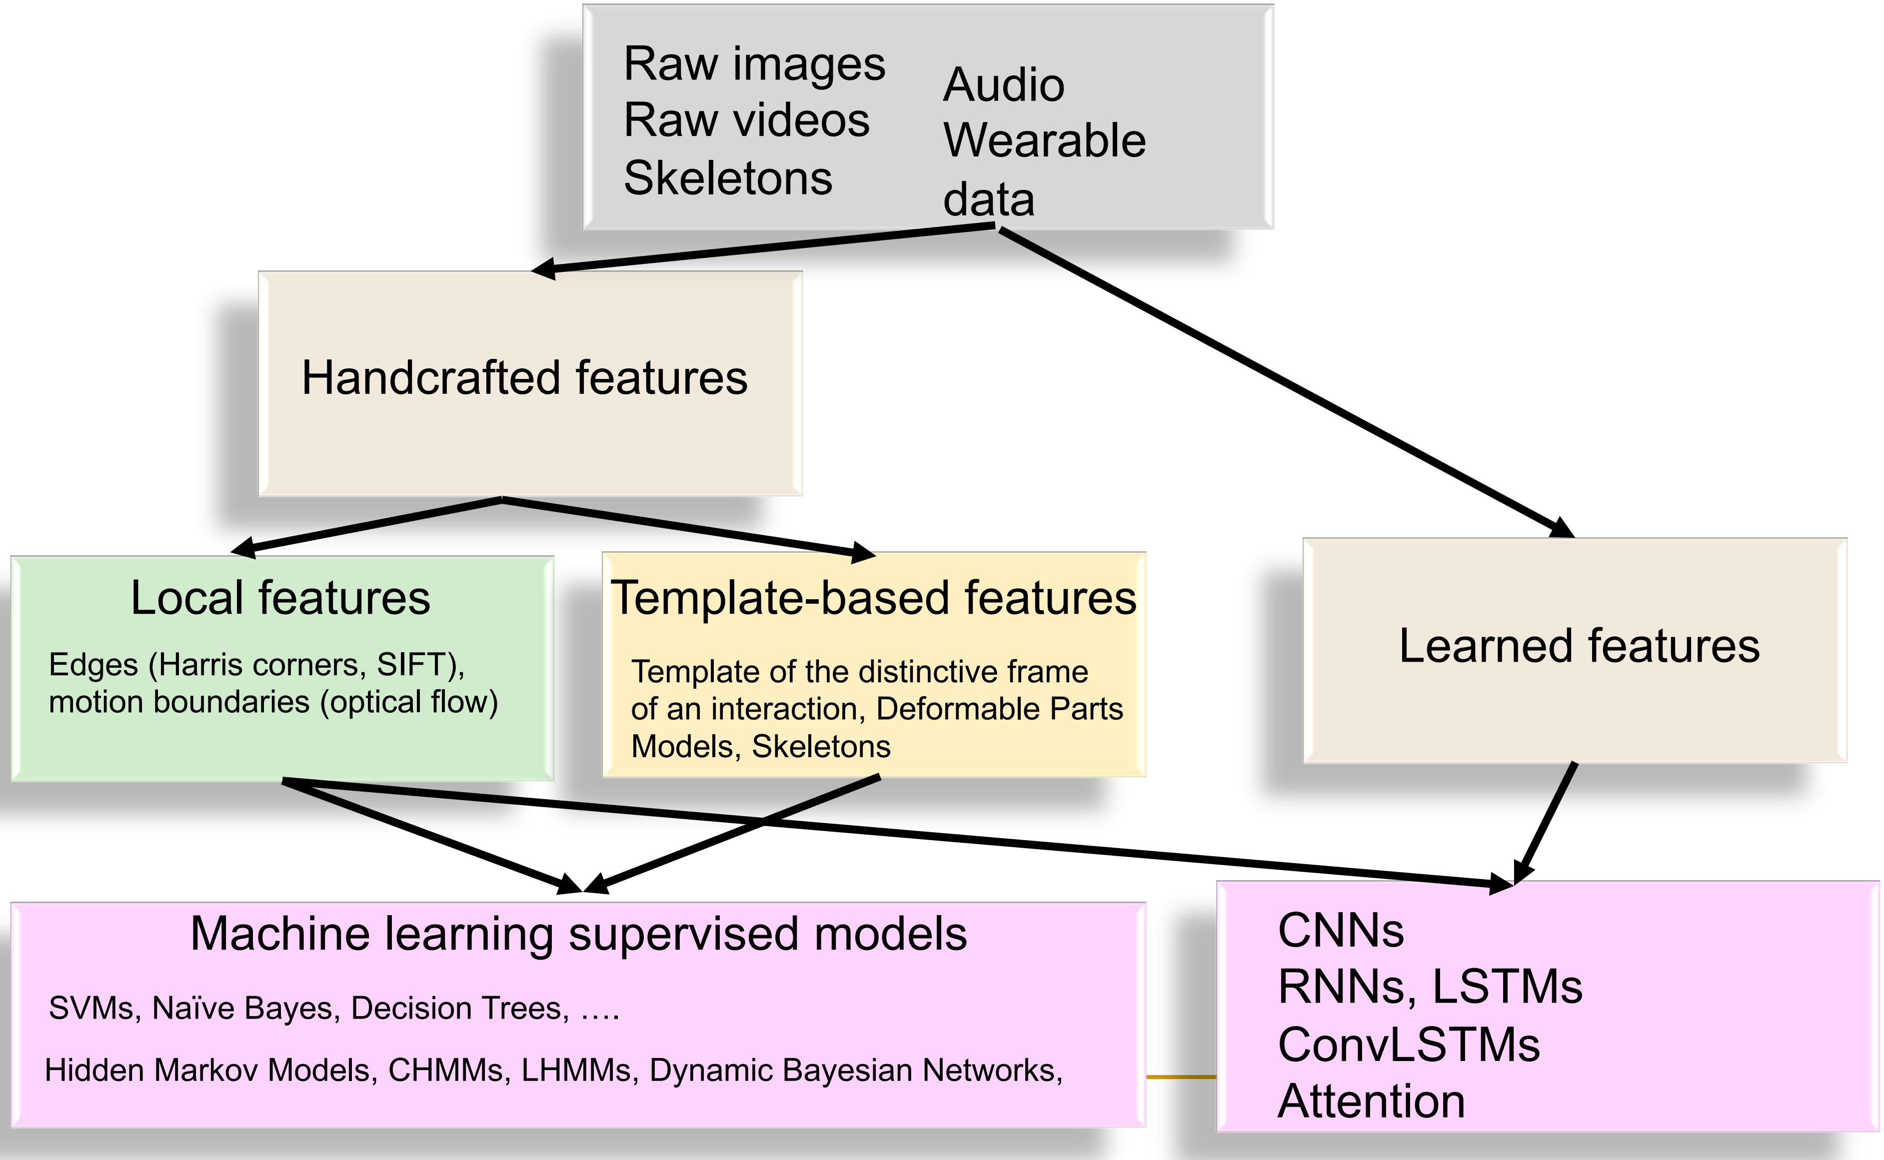

Early work on Visual Surveillance

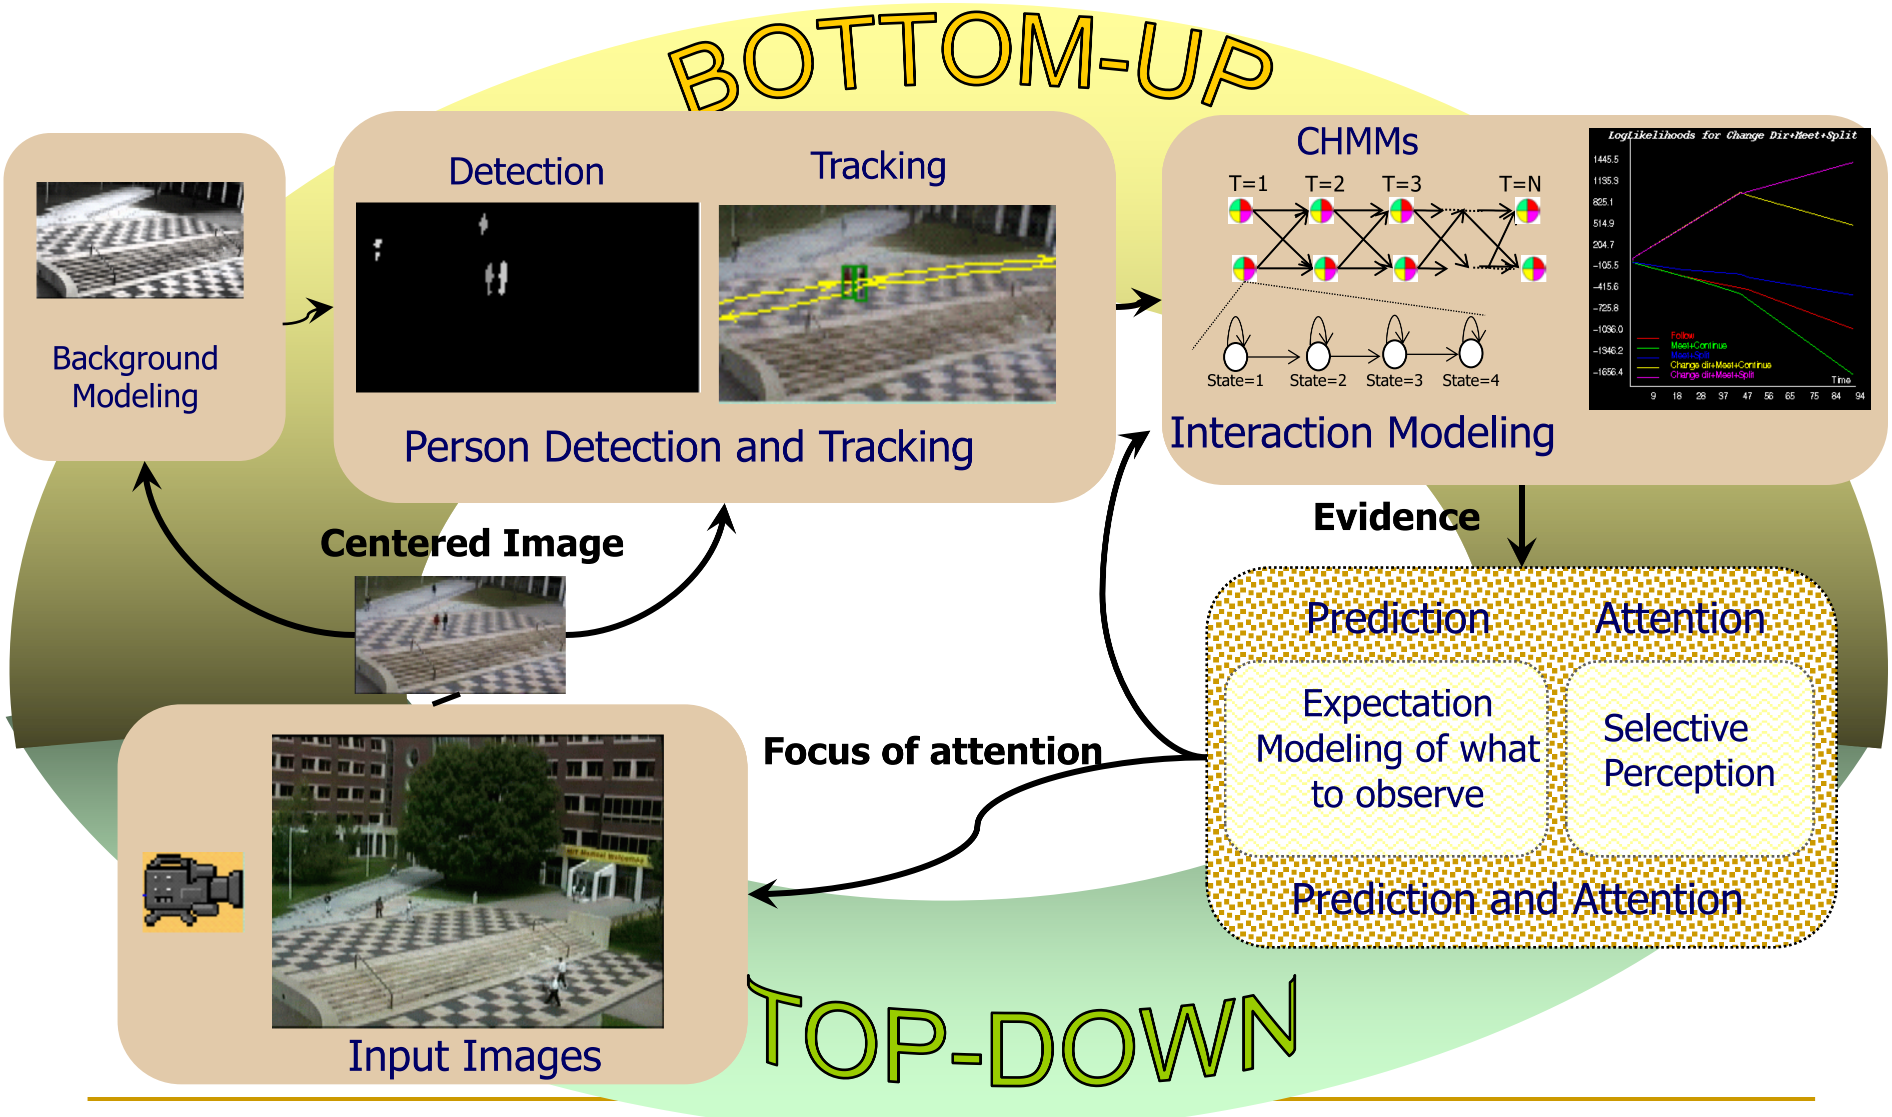

Oliver, N., Rosario, B. and Pentland, A. (2000) "Graphical Models for Recognizing Human Interactions" , Proceedings of Intl. Conf. on Neural Information and Processing Systems **NIPS98**. Also in **IEEE TPAMI**, 2000

CHMMs: Coupled Hidden Markov Models

4-state CHMM

Graphical Models

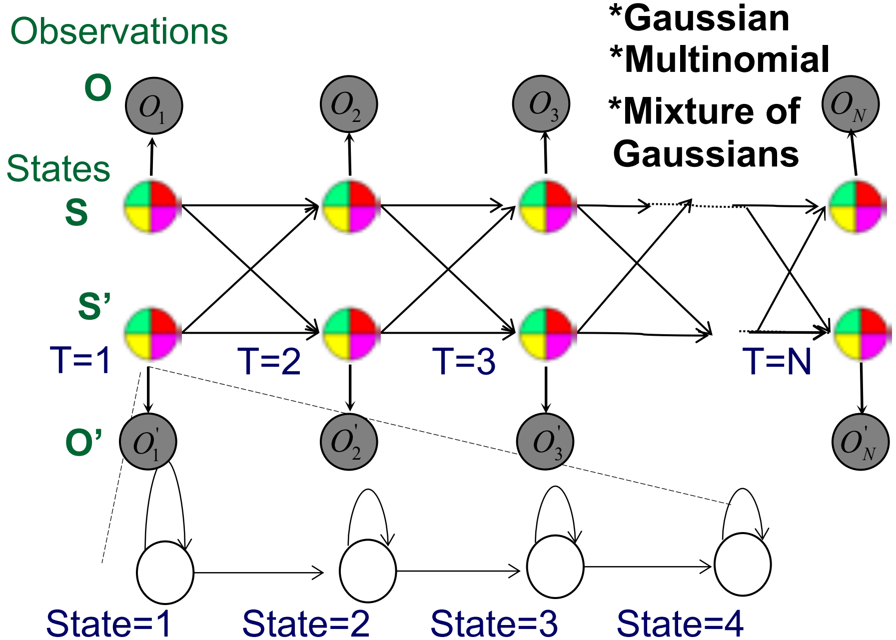

$$P(S|O) \propto P_{s_1} p_{s_1}(o_1) P_{s'_1} p_{s'_1}(o'_1)$$

$$\prod_{t=2}^{T} P_{s_t|s_{t-1}} P_{s_t|s'_{t-1}} P_{s'_t|s'_{t-1}} P_{s'_t|s_{t-1}} p_{s_t}(o_t) p_{s'_1}(o'_t)$$

State Trellis

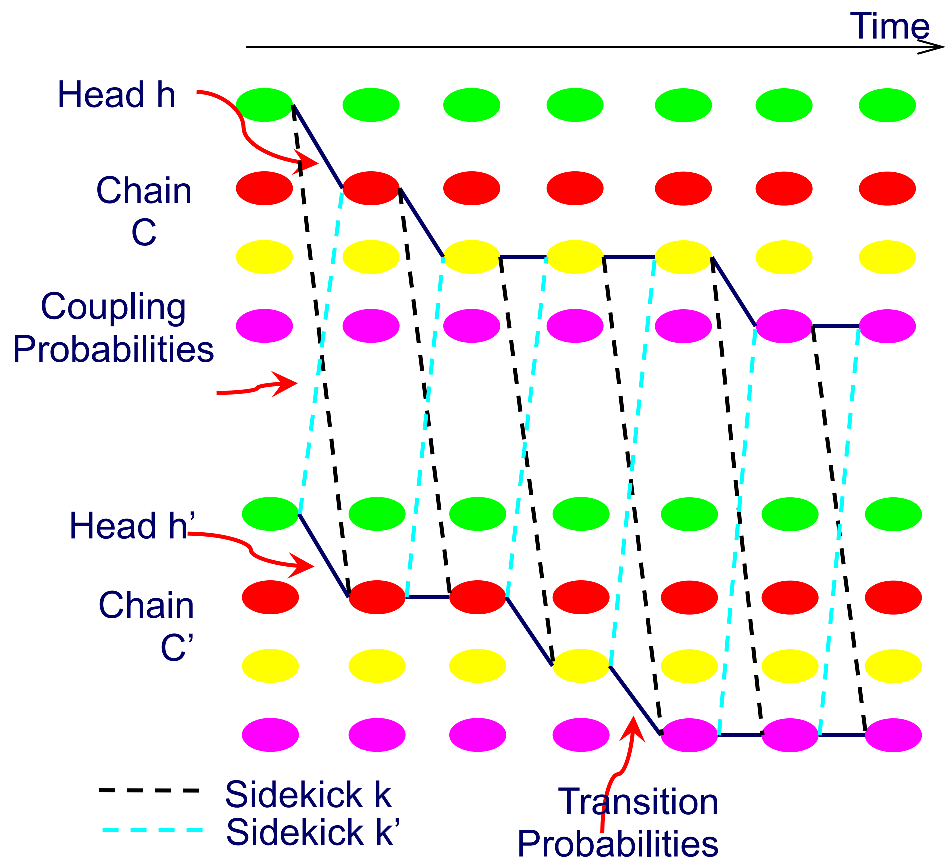

Oliver, N., Rosario, B. and Pentland, A. (2000) "Graphical Models for Recognizing Human Interactions" , Proceedings of Intl. Conf. on Neural Information and Processing Systems NIPS98. Also in IEEE TPAMI, 2000

Generating synthetic ground truth

- **Problem:** Very little real data
- Generate behavioral data from interacting synthetic agents
  - Testbed for behavioral graphical models previous to real data
  - Ground truth is known
  - Model and recognize different interactions
  - Model interaction vs. non-interactive behavior

Discovering Interactions

- **Problem:** There might be multiple people in a scene who might be interacting with each other or not
- **Solution:** Dynamic Graphical Model that detects interactions

Learning: Discovering Interactions

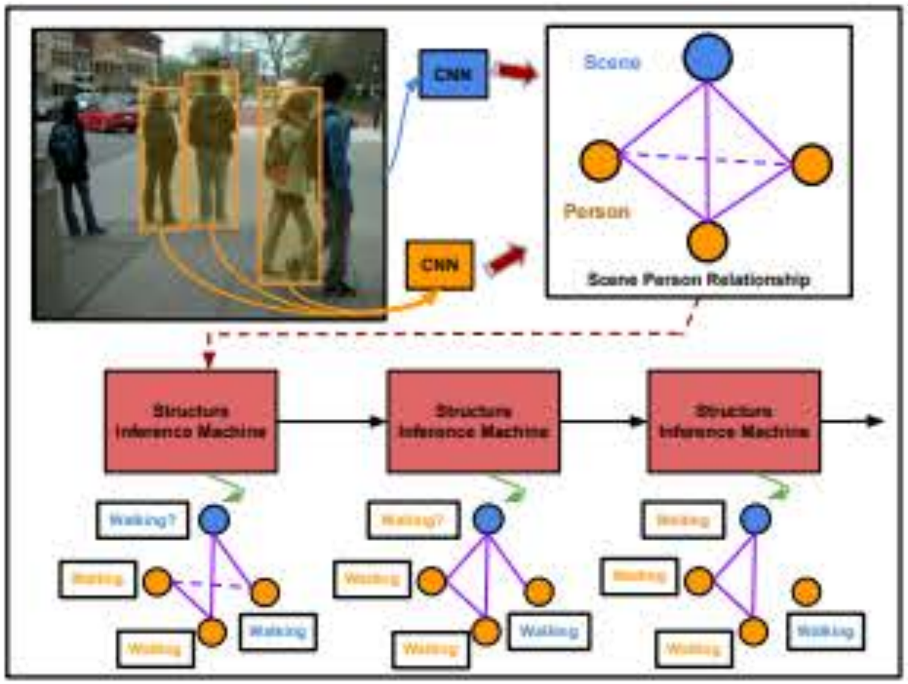

Structure Inference Machine iteratively reasons about which people in a scene are interacting and which are involved in group activity

- **Structure Inference Machine:**
  - **RNN** aggregates cues about the actions of other people in the scene by passing messages that refine estimates of an individual's action;
  - **Trainable gating functions** that can turn on/off connections between individuals in the scene depending on whether they are interacting

Deng, Z., Vahdat, A., Hu, H., Mori, G. (2016). "Structure inference machines: Recurrent neural networks for analyzing relations in group activity recognition", CVPR 2016, pp. 4772–4781.

Learning: Discovering Interactions

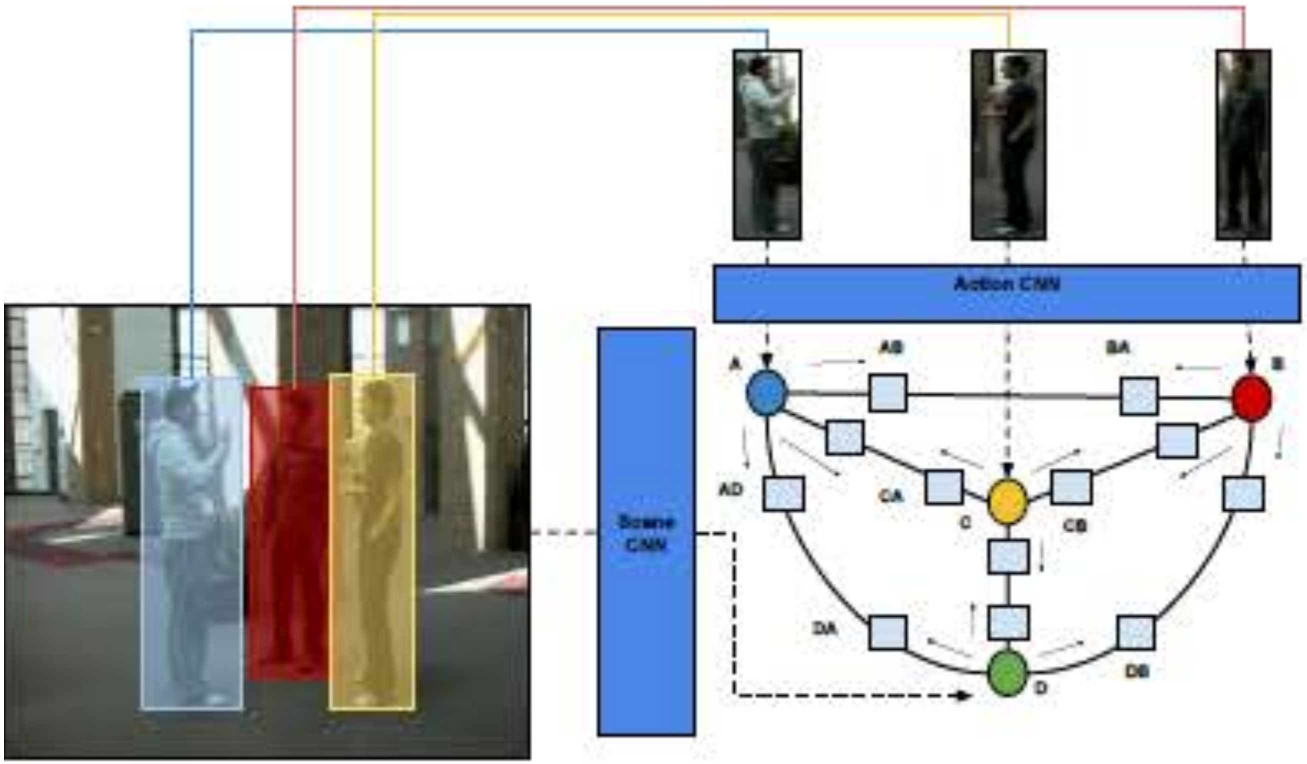

Group activity represented as a graphical model. Estimates of individual person actions and group activity are refined via message passing. The squares are messages. The message units carry information from the source node to the target node

Deng, Z., Vahdat, A., Hu, H., Mori, G., (2016). “Structure inference machines: Recurrent neural networks for analyzing relations in group activity recognition”,CVPR 2016, pp. 4772–4781.

Learning: Discovering Interactions

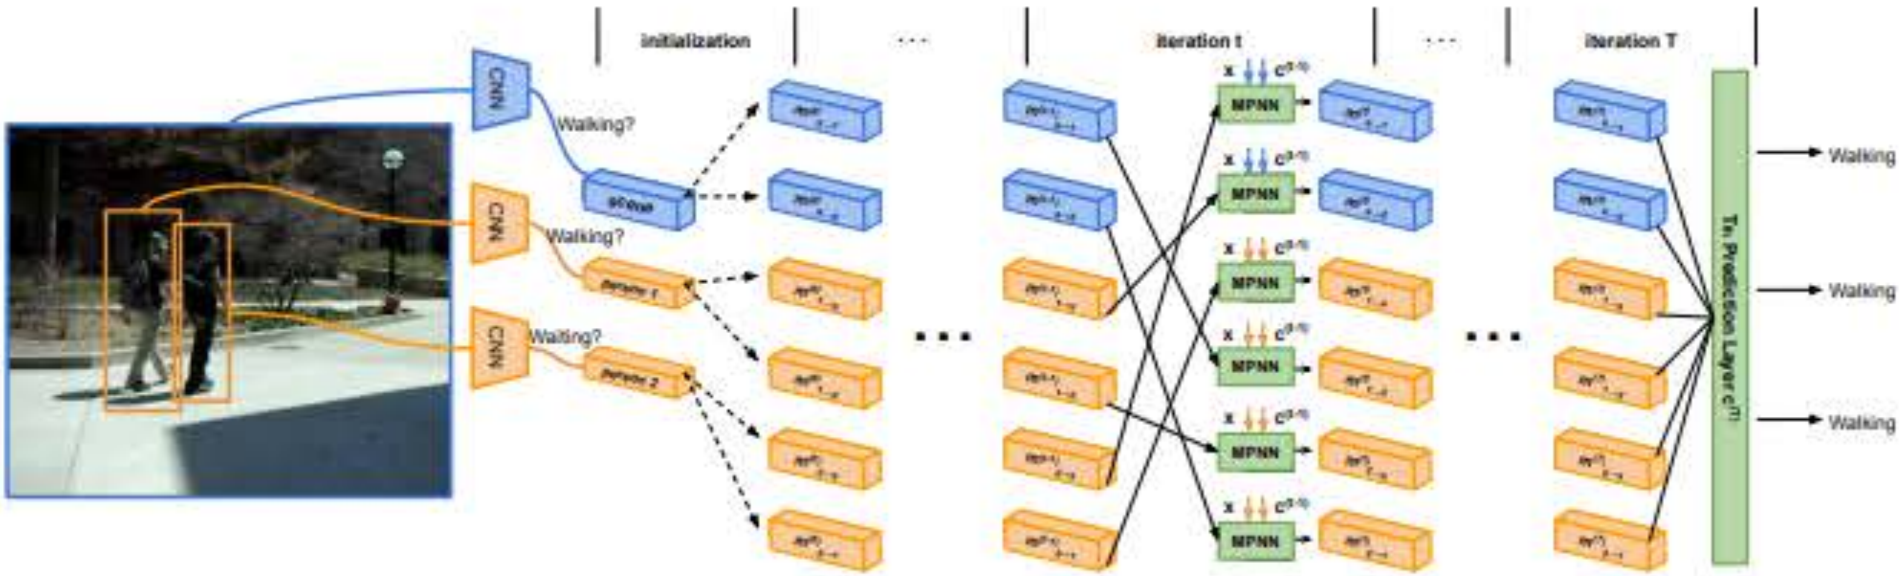

Every iteration new messages are computed using unary scores, related message units and output predictions from the previous timestep. For each timestep, a prediction layer outputs predictions. In training receives loss as in a standard RNN

- **Collective Action Dataset: 81% accuracy (five classes)**

Deng, Z., Vahdat, A., Hu, H., Mori, G., (2016). "Structure inference machines: Recurrent neural networks for analyzing relations in group activity recognition",CVPR 2016, pp. 4772–4781.

Multiple Levels of Abstraction

- **Problem:** Interactive behavior entails modeling the behavior at different levels of abstraction, from individual actions to group interactions
- **Solution:** Hierarchical architectures

Hierarchical, Multi-modal Recognition

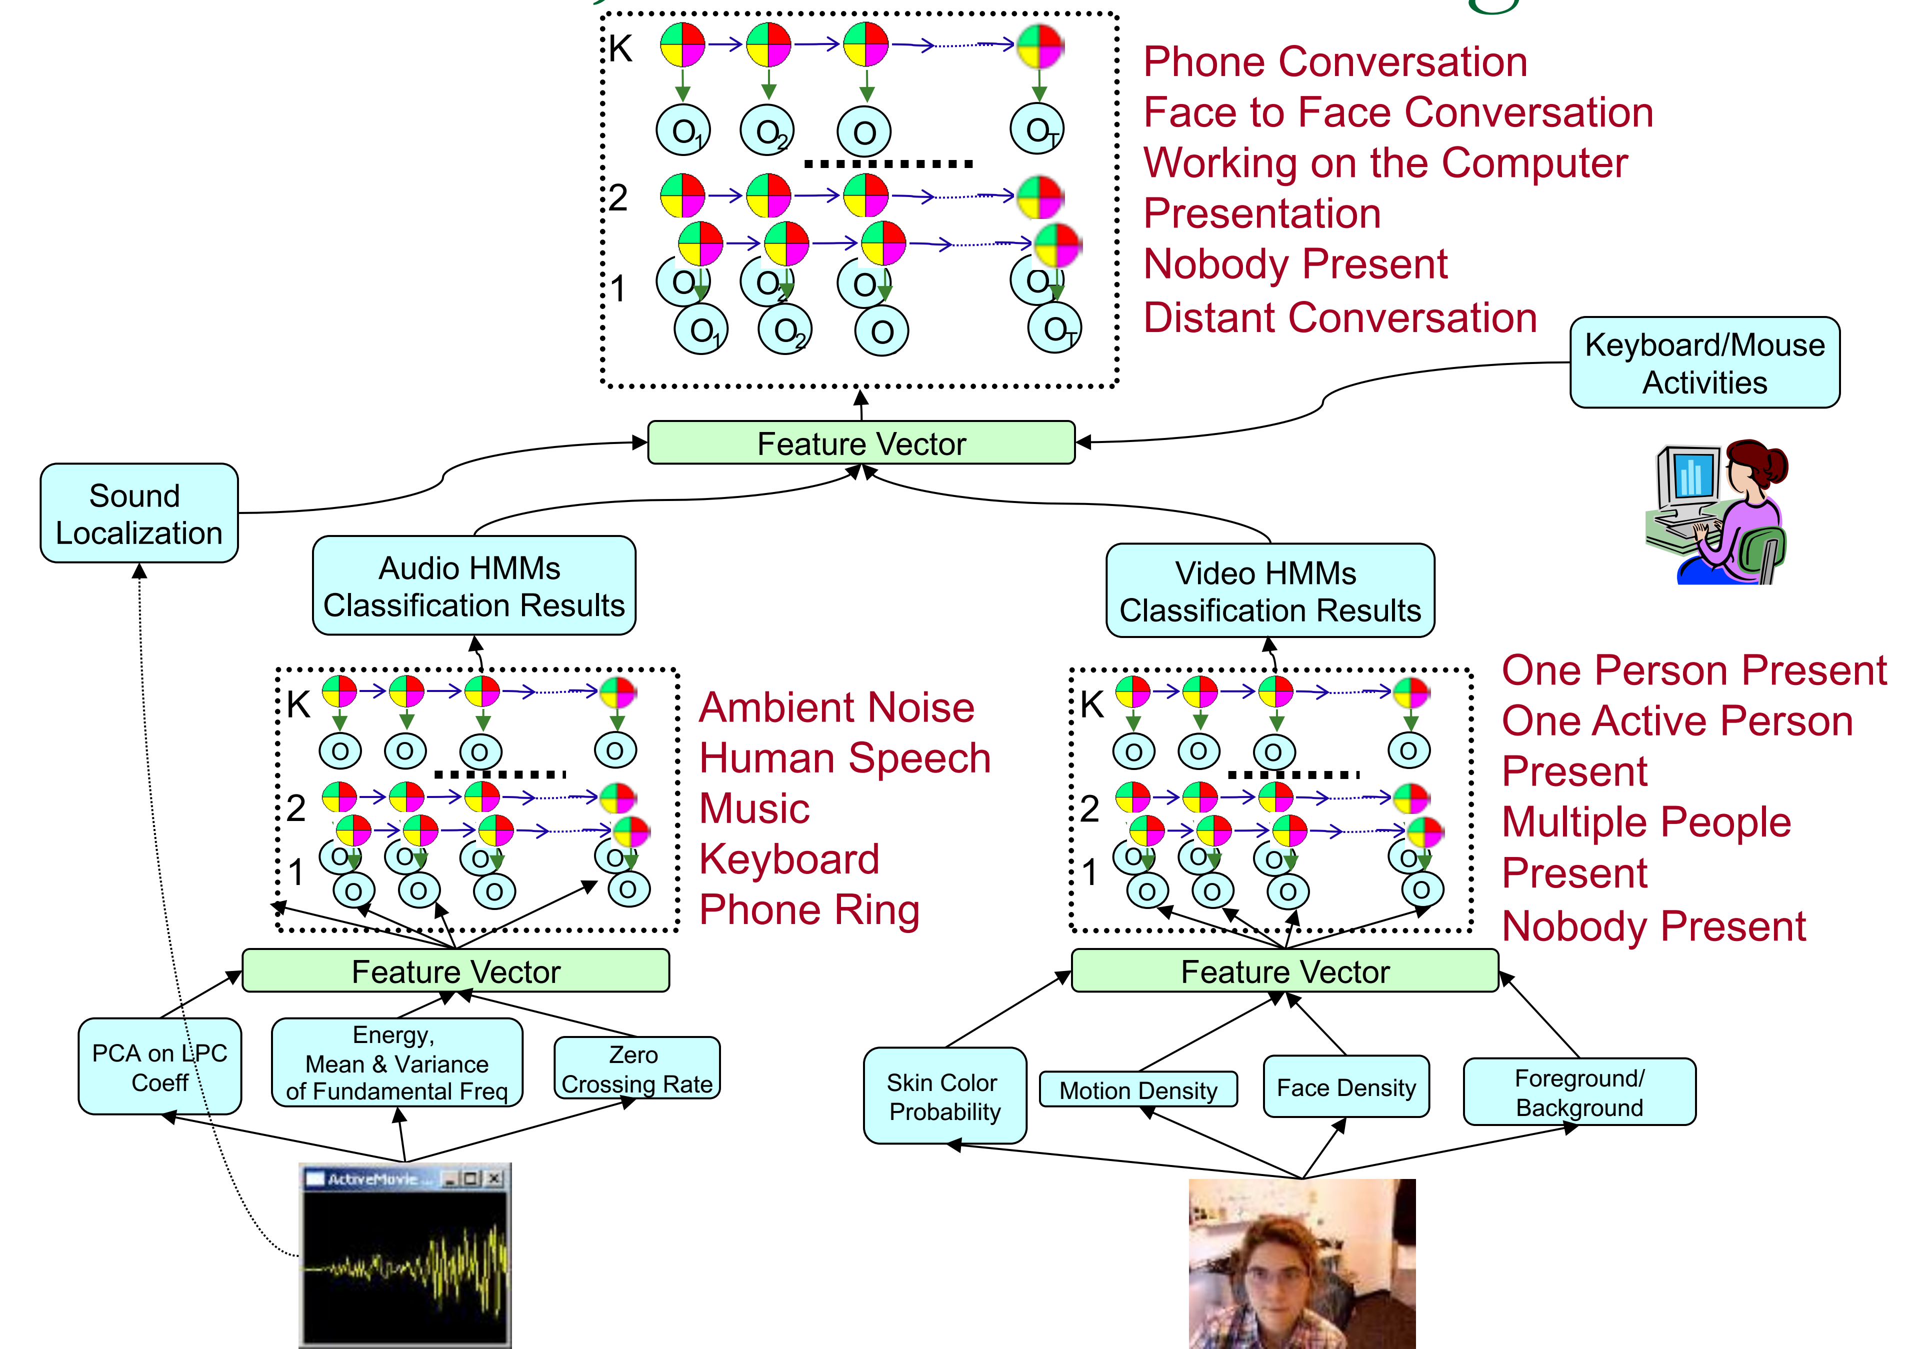

Layered Hidden Markov Models

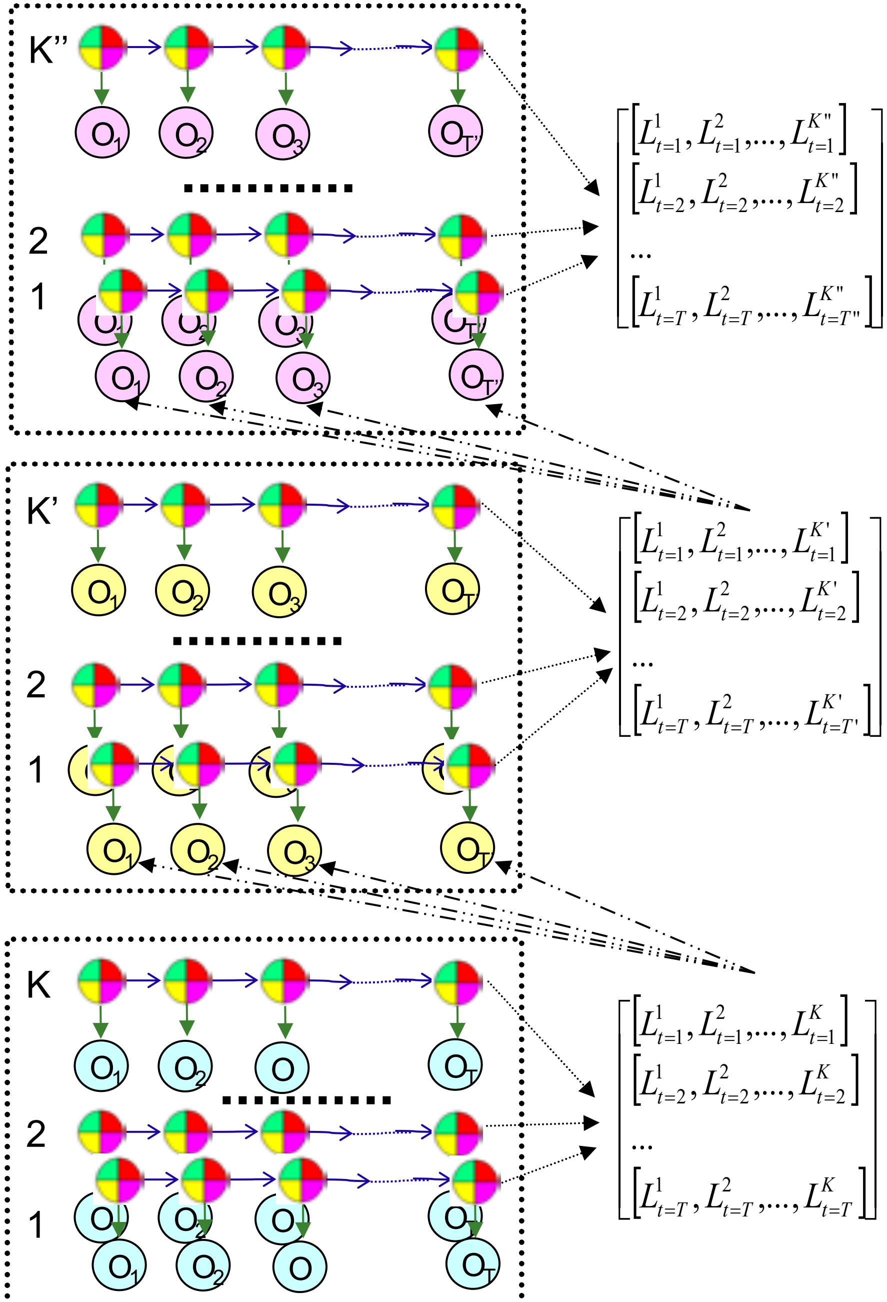

Layer 3

Bank of K''
-Level 3 HMM
Classifiers with
time granularity T''

Layer 2

Bank of K'
-Level 2 HMM
Classifiers with
time granularity T'

Layer 1

Bank of K
-Level 1 HMM
Classifiers with
time granularity T

Hierarchical Deep Temporal Models

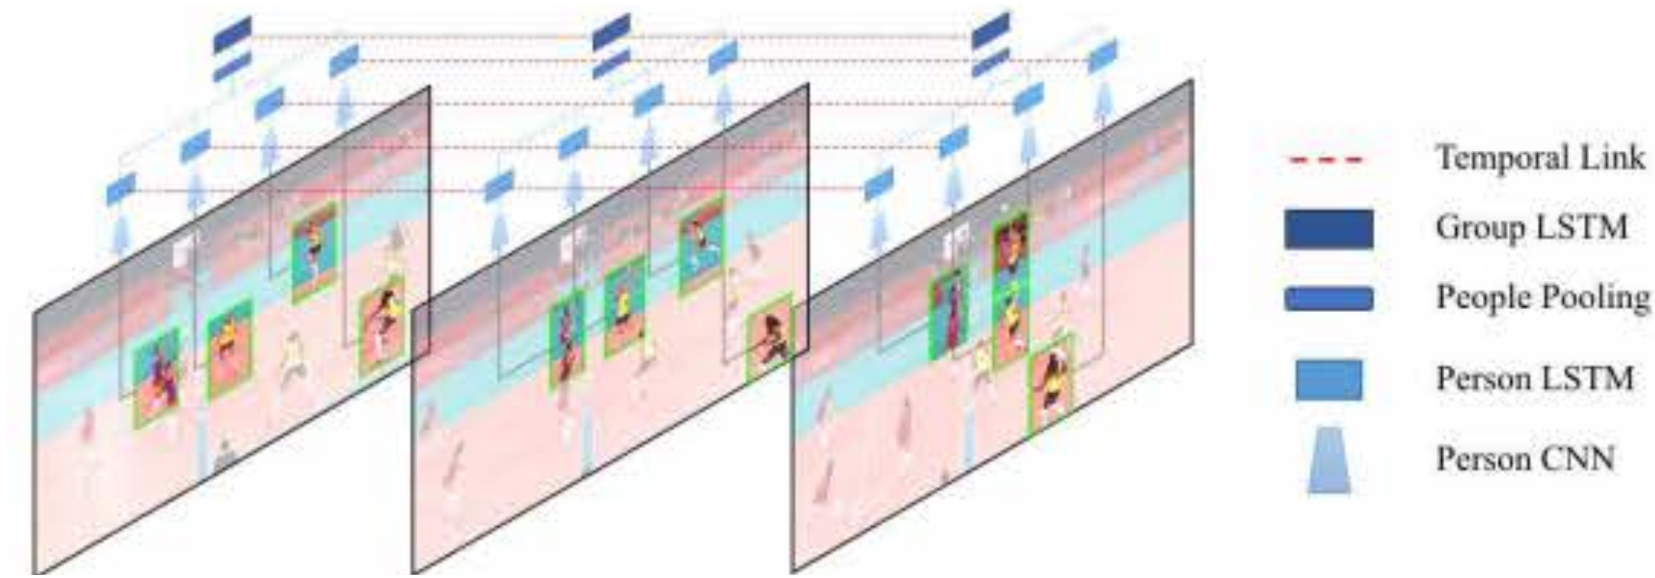

Two-state model for a volleyball match. There is a tracklet for each player which is the input to a CNN, followed by a person LSTM layer to represent each player's action. They pool over all the people's temporal features in the scene. The output of the pooling layer is fed to the second LSTM network to identify an entire team's activity

- **Hierarchy of LSTMs** to model interactions
- **Collective Action Dataset: 83% accuracy.** Best recognition: talking (99%); worst: crossing (61%)

Ibrahim, M.S., Muralidharan, S., Deng, Z., Vahdat, A., Mori, G., (2016). "A hierarchical deep temporal model for group activity recognition", CVPR 2016, pp. 1971–1980.

Hierarchical Models

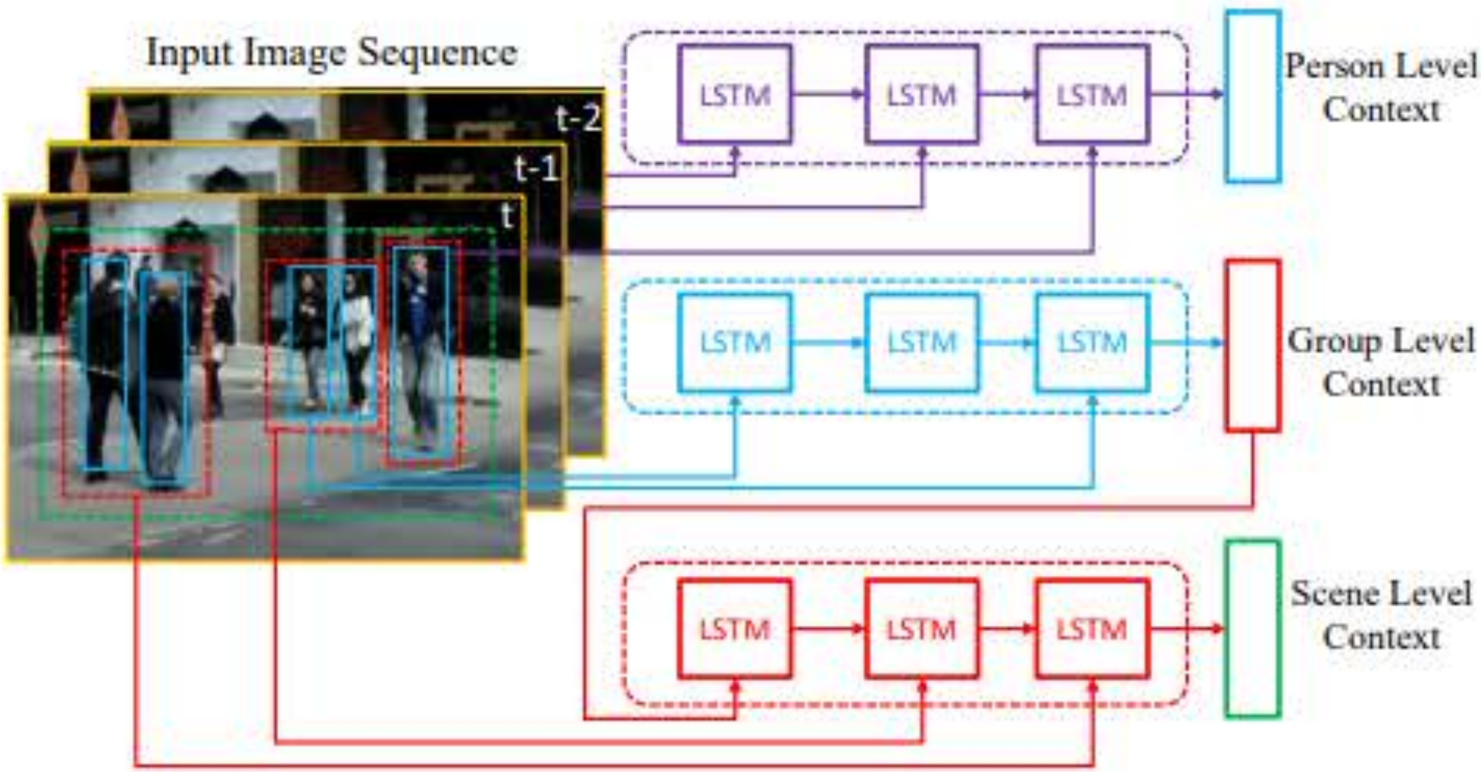

- **Hierarchy of LSTMs** to model intra-group and inter-group interactions

Wang, M., Ni, B., Yang, X. and Jiao, S., 2017, “Recurrent modeling of interaction context for collective activity recognition”, CVPR 2017

Hierarchical Models

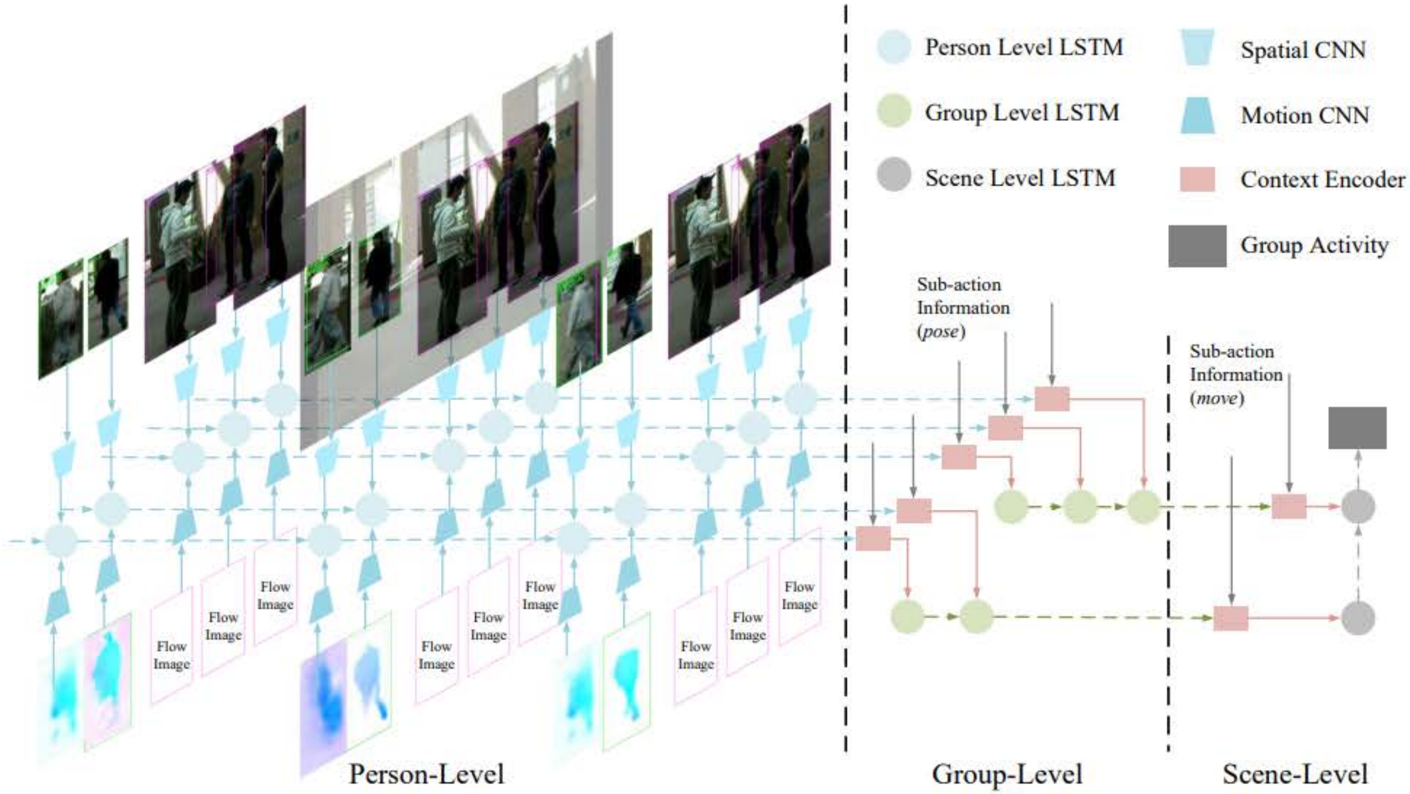

- Evaluated on **two datasets**: Collective Activity Dataset and Choi's new dataset
- Competitive performance on some interactions but limited performance in others due to small amount of examples

Wang, M., Ni, B., Yang, X. and Jiao, S., 2017, "Recurrent modeling of interaction context for collective activity recognition", CVPR 2017

GANs for Social Interaction Modeling

- **Goal: Jointly** reason and predict future trajectories of all agents in a scene
- **Approach:** Socially aware GANs where the human-human interaction is modeled via a Pooling Module
- **Use case:** Autonomous vehicles

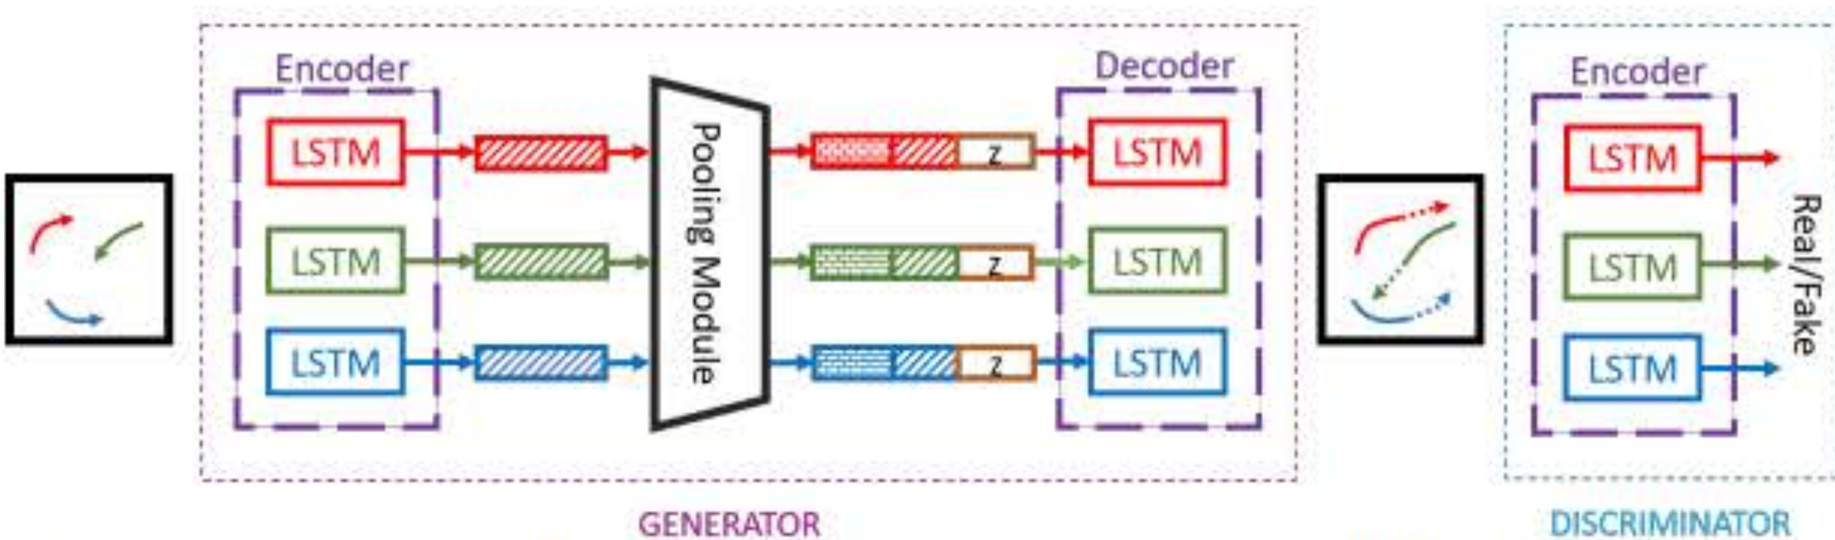

Gupta, A., Johnson, J., Li, F.F., Savarese, S., Alahi, A., 2018. «Social GAN: Socially acceptable trajectories with generative adversarial networks”, CVPR 2018, pp. 2255–2264.

GANs for Social Interaction Modeling

- Proposed method **outperforms LSTMs**
- **Pooling** helps avoid collisions between people when meeting, following each other or avoiding another person

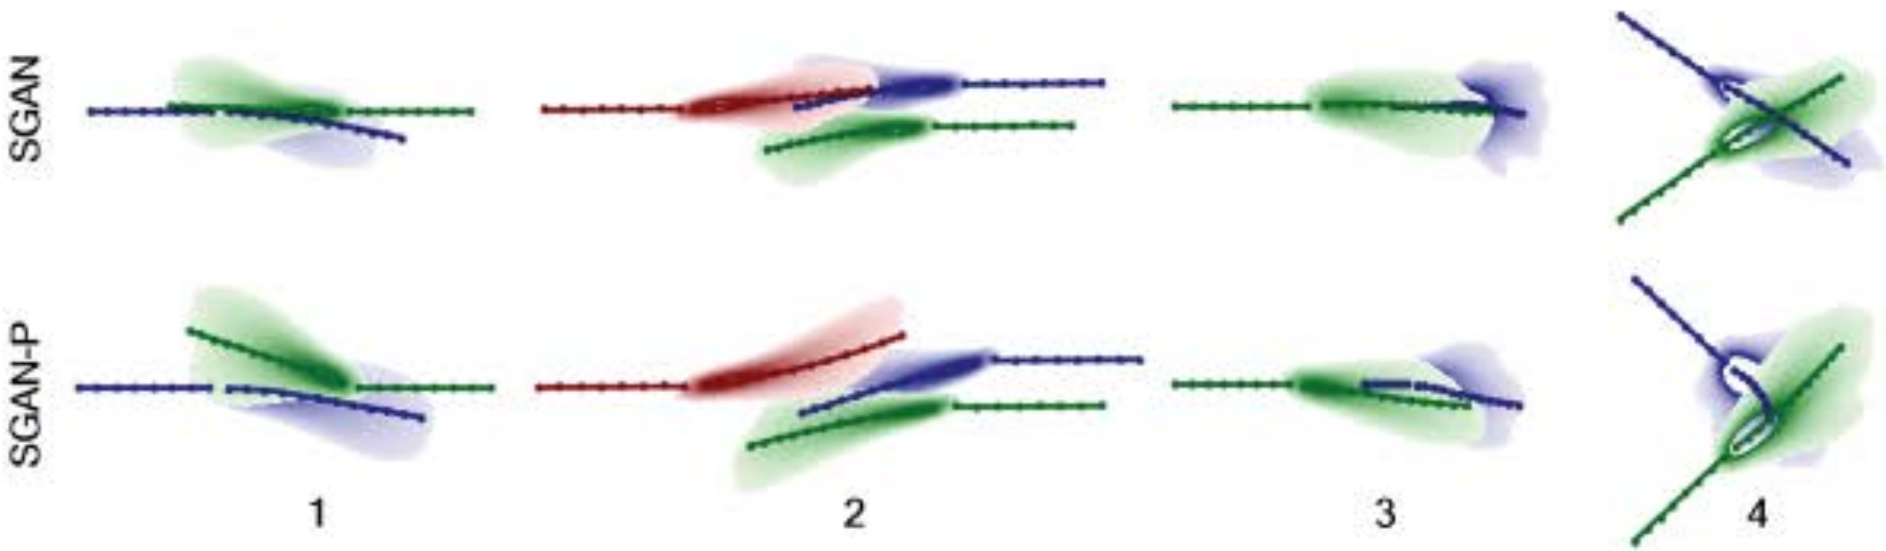

Gupta, A., Johnson, J., Li, F.F., Savarese, S., Alahi, A., 2018. «Social GAN: Socially acceptable trajectories with generative adversarial networks", CVPR 2018, pp. 2255–2264.

Summary: Interactive behaviors

- Much less mature and harder area than individual behavior modeling
- Worse performance in general than in individual action recognition
- Different levels of abstraction → Hierarchical dynamic models perform well
- Sparsity of available ground truth data
  - Synthetic data generated with e.g. GANs can help address this limitation

LARGE SCALE HUMAN BEHAVIOR MODELING: COMPUTATIONAL SOCIAL SCIENCES

Fundamental questions

- How do we represent the behavior?

Fundamental questions

- Why investigate behavior at this scale?
- How do we represent the behavior?

Computational Social Science

The ubiquity of mobile phones enables us to collect and analyze, for the first time in human history, **large-scale aggregated** and anonymized **human behavioral data** of entire cities, countries or even continents

The opportunity is HUGE to help decision making units (governments, UN, Red Cross…) make more informed decisions thanks to the existence of quantitative real-time information about populations

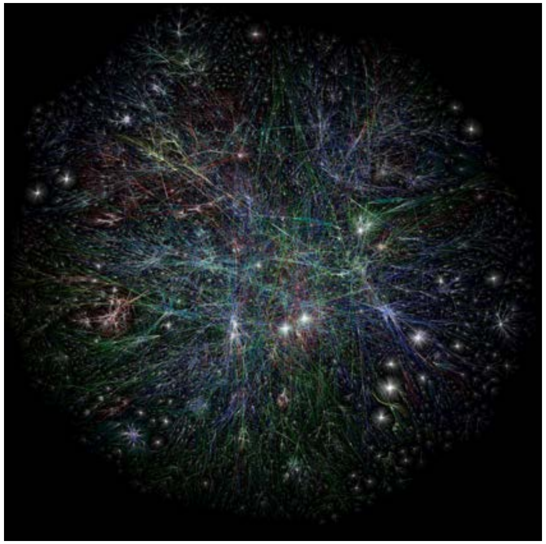

Source: Kaspersky Lab

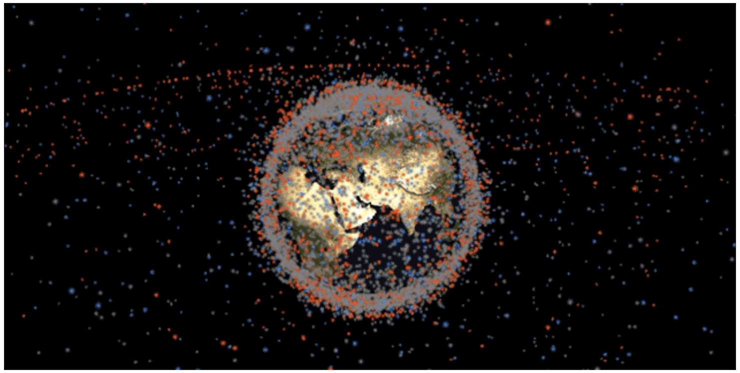

Source: Kaspersky Lab

Lazer D., Pentland A., Adamic L., Aral S., Barabási AL et al, 2009, "Computational Social Science", Science, Vol 323, pp 721-723

High-level Panel ON the Post-2015 Development Agenda

HOME | ABOUT | THE PANEL | THE SECRETARIAT | OUTREACH | THE REPORT | MEDIA

HIGH LEVEL PANEL RELEASES RECOMMENDATIONS FOR WORLD'S NEXT DEVELOPMENT AGENDA

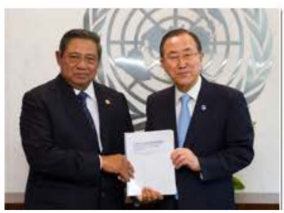

Eminent Persons from Around the World Call for a New Global Partnership to Eradicate Poverty and Transform Economies through Sustainable Development

The High Level Panel on the Post-2015 Development Agenda today released **"A New Global Partnership: Eradicate Poverty and Transform Economies through Sustainable Development,"** a report which sets out a universal agenda to eradicate extreme poverty from the face of the earth by 2030, and deliver on the promise of sustainable development. The report calls upon the world to rally around a new Global Partnership that offers hope and a role to every person in the world.

Wanted: A data revolution

"Data are the lifeblood of decision-making and the raw material for accountability*.*

*Governments, companies, researchers and citizen groups are in a ferment of experimentation, innovation and adaptation to the new world of data****, a world in which data are bigger, faster and more detailed than ever before. This is the data revolution****."*

Data Revolution Report

'A WORLD THAT COUNTS' Presented to Secretary-General

The Secretary-General's Independent Expert Advisory Group on a Data Revolution for Sustainable Development (IEAG) met the Secretary-General today to hand over their culminating report *A World That Counts: Mobilising The Data Revolution for Sustainable Development.*

Download 'A World That Counts'

The IEAG consists of over 20 international experts convened by the Secretary-General Ban Ki-moon to propose ways to improve data for achieving and monitoring sustainable development. The report highlights two big global challenges for the current state of data:

- The challenge of invisibility (gaps in what we know from data, and when we find out)
- The challenge of inequality (gaps between those who with and without information, and what they need to know make their own decisions)

https://www.undatarevolution.org/wp-content/uploads/2014/11/A-World-That-Counts.pdf

Never again should it be possible to say "we didn't know". No one should be invisible. This is the world we want – a world that counts.

https://www.undatarevolution.org/wp-content/uploads/2014/11/A-World-That-Counts.pdf

Sustainable Development Goals

The (Big) Data Revolution and the Sustainable Development Goals

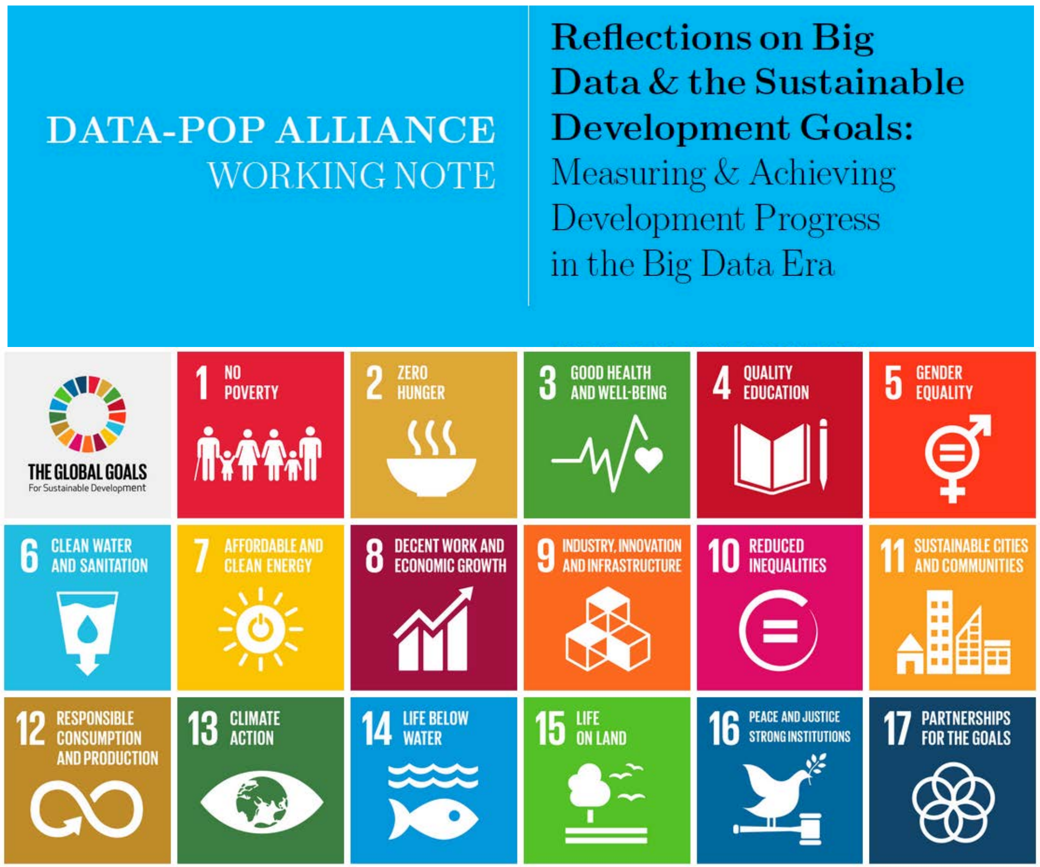

1. **How can (Big) Data help *monitor* the SDGs** by “filling data gaps” with more granular & disaggregated data—*and what does monitoring something do to that something?*

2. **How can (Big) Data help *promote (or impede?)* the SDGS** and their underlying human development vision and objectives—*including towards and through lower (or higher?) inequalities?*

The UN World Data Forum 2018 will be hosted by Federal Competitiveness and Statistics Authority, of United Arab Emirates from 22 to 24 October 2018, with support from the Statistics Division of the UN Department of Economic and Social Affairs, under the guidance of the United Nations Statistical Commission and the High-level Group for Partnership, Coordination and Capacity-Building for Statistics for the 2030 Agenda for Sustainable Development.

2000 experts from more than **100 countries**, with the aim of building broad consensus on **how to harness the power of data for sustainable development**.

UN News Centre

with breaking news from the UN News Service

News | Radio | Television | Photo | Webcast | Meetings Coverage | Media Accreditation | Secretary-G

Africa | Americas | Asia Pacific | Europe | Middle East | ICYMI

Dubai Declaration

Supporting the Implementation of the Cape Town Global Action Plan for Sustainable Development Data

1. We, the participants gathered here in Dubai, United Arab Emirates, for the 2018 United Nations World Data Forum, from national statistical offices and other parts of the national statistical systems, other data communities, government institutions, private sector, civil society, academia, and media.
2. *Stressing* that the full ambition of the 2030 Agenda for Sustainable Development (2030 agenda) cannot be realized without quality, timely, relevant, open and disaggregated data to ensure that no one is left behind.
3. *Recognizing* that the 2030 agenda requires that national statistical systems transform and develop to be agile and responsive to meet the increased demands of data users, including for the full implementation of the 2030 agenda.
4. *Recognising* that the Cape Town Global Action Plan for Sustainable Development Data (CTGAP), launched at the first United Nations World Data Forum held in Cape Town, South Africa in January 2017 and adopted by the UN Statistical Commission, and welcomed by member states in the General Assembly resolution 71/313[1], guides the implementation of programmes and activities to respond to the data needs of the 2030 Agenda.
5. *Stressing* the need to support fundamental data collection programmes, such as the 2020 population and housing census round.
6. *Stressing* the importance of coordination across the statistical system, including better use and integration of administrative data sources.
7. *Acknowledging* that the data demands for the 2030 Agenda require urgent new, standards-based and interoperable solutions that leverage the power of new data sources and technologies through partnerships between national statistical authorities and the private sector, civil society, and the academia and other research institutions.
8. *Acknowledging* that the 2030 Agenda explicitly calls for enhanced support for strengthening data collection and capacity-building in Member States and the capacity of national statistical

Fundamental questions

- Why investigate behavior at this scale?
- How do we **represent** the behavior?

Representation: ‘Sensors’

- Mobile call data records (CDR), extended data records (XDR)
- Satellite images
- Social media content
- Wearables and smart watches
- Infrastructure usage (e.g. railways)
- … any data trace on systems used by thousands of people

Representation: Mobile CDR

5.1 billion mobile phone subscribers worldwide
~70% of world's population (ITU)

Mobile penetration of 120% to 89% of population (ITU)

More time spent on our phones than watching TV or with our with our partner (US and UK)

Emerging and developed regions

Representation: Mobile CDR

5.1 billion mobile phone subscribers worldwide
~70% of world's population (ITU)

Mobile penetration of 120% to 89% of population (ITU)

More time spent on our phones than watching TV or with our with our partner (US and UK)

Emerging and developed regions

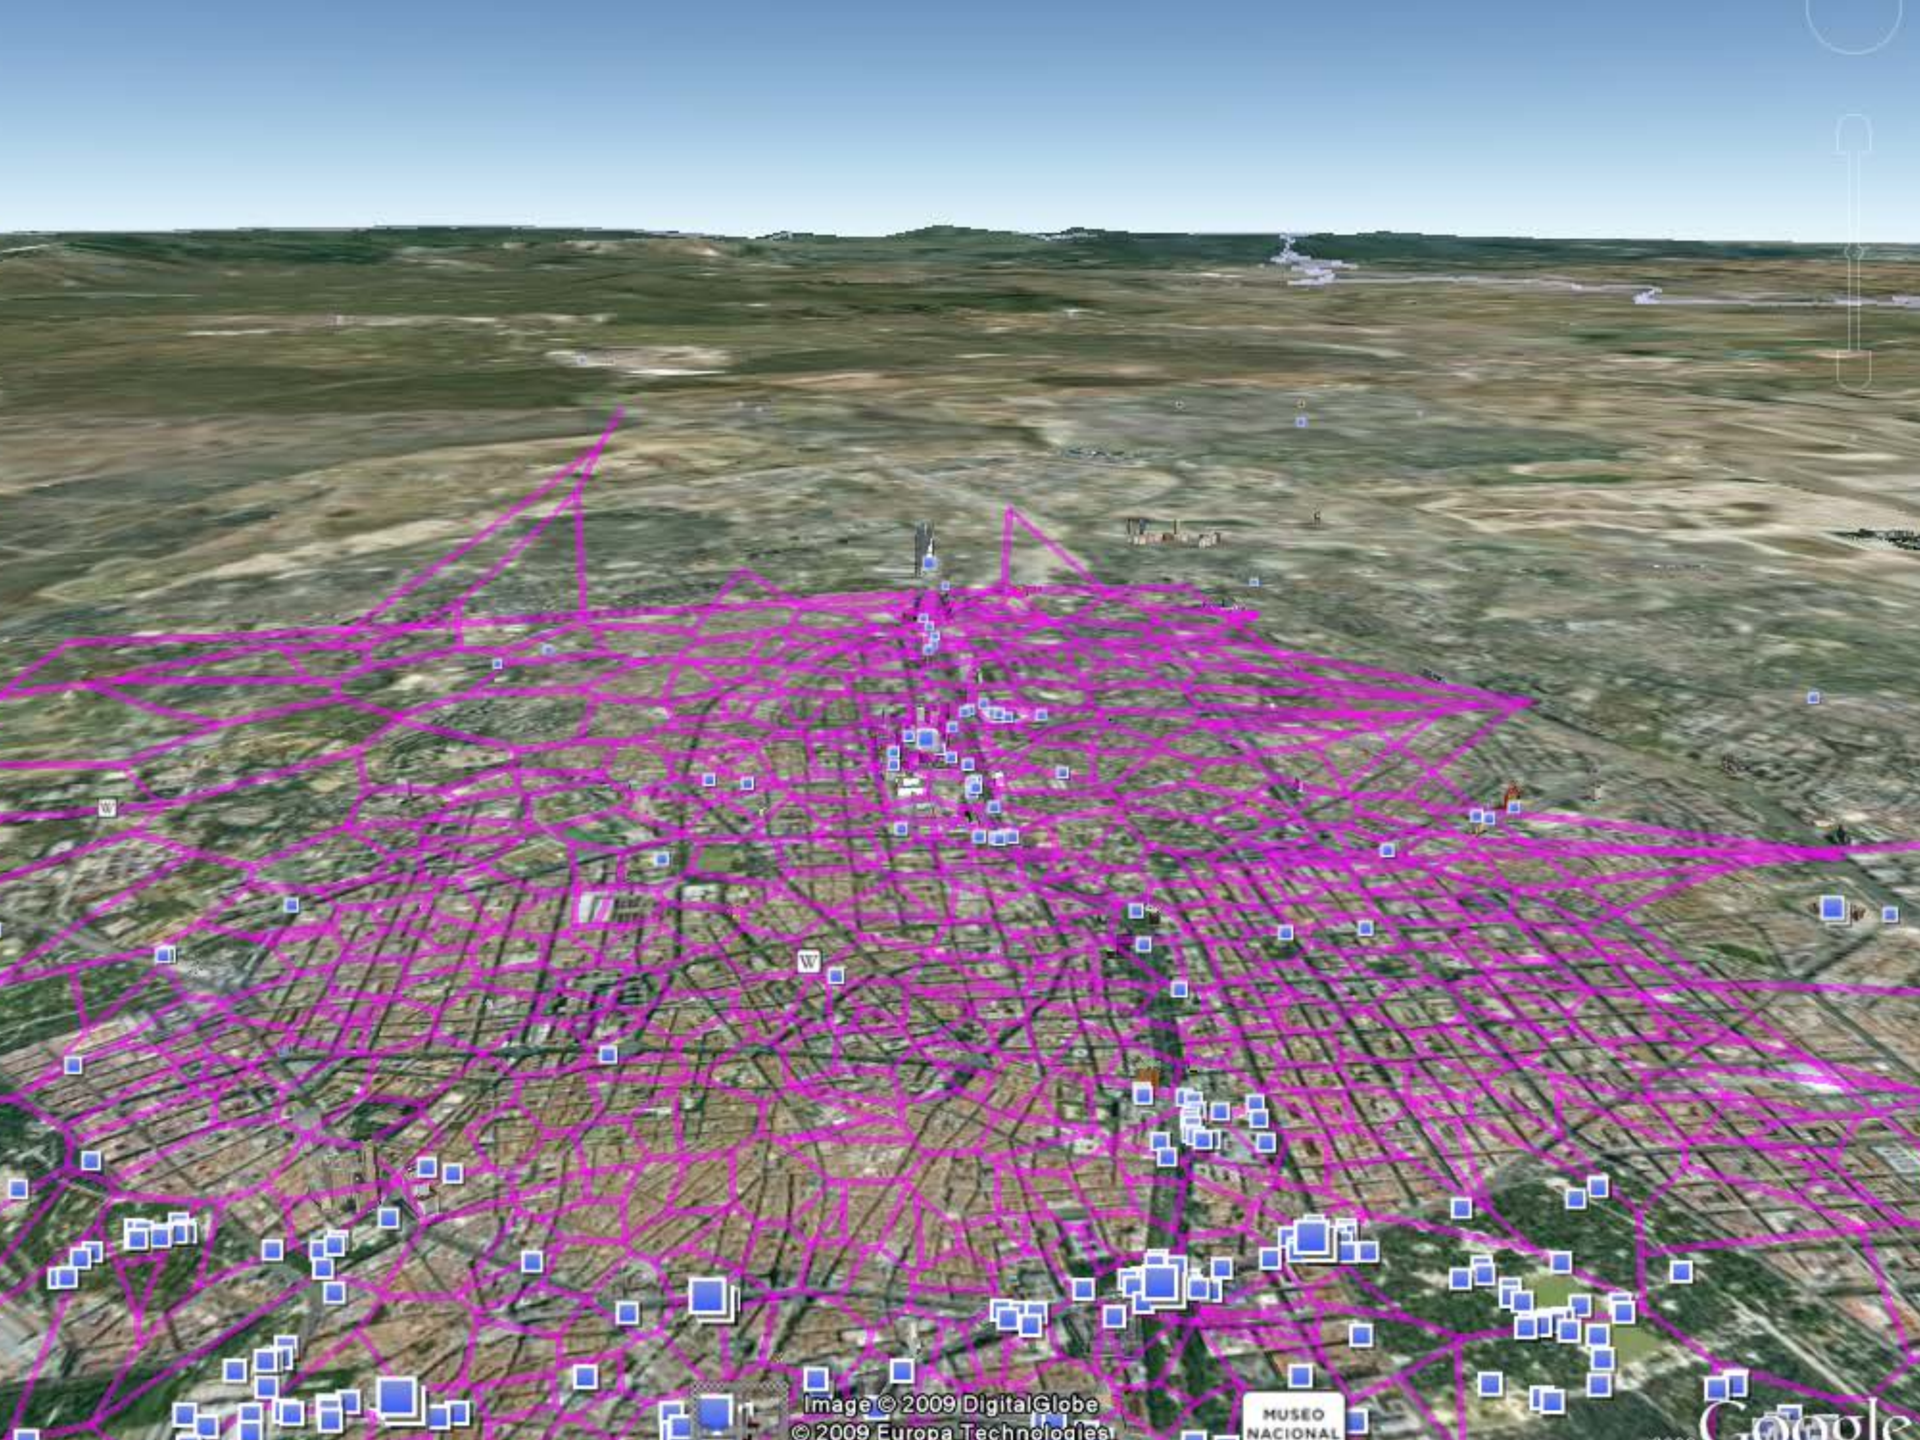
Image © 2009 DigitalGlobe
© 2009 Europa Technologies
MUSEO NACIONAL
Google

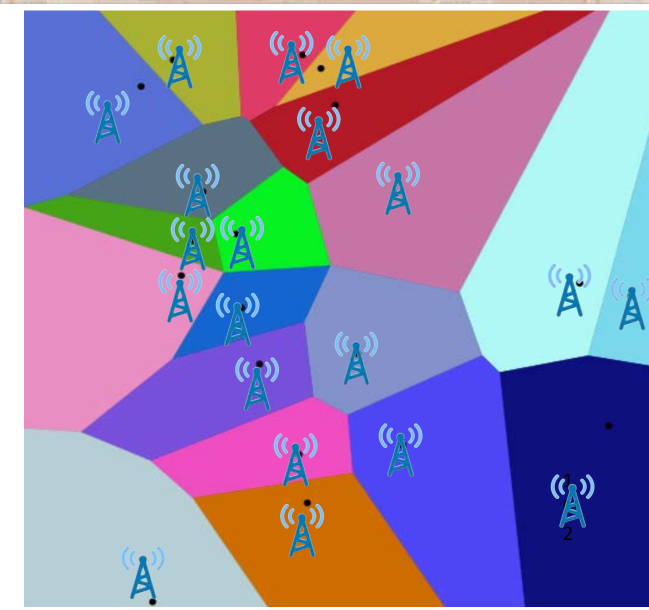

1
2

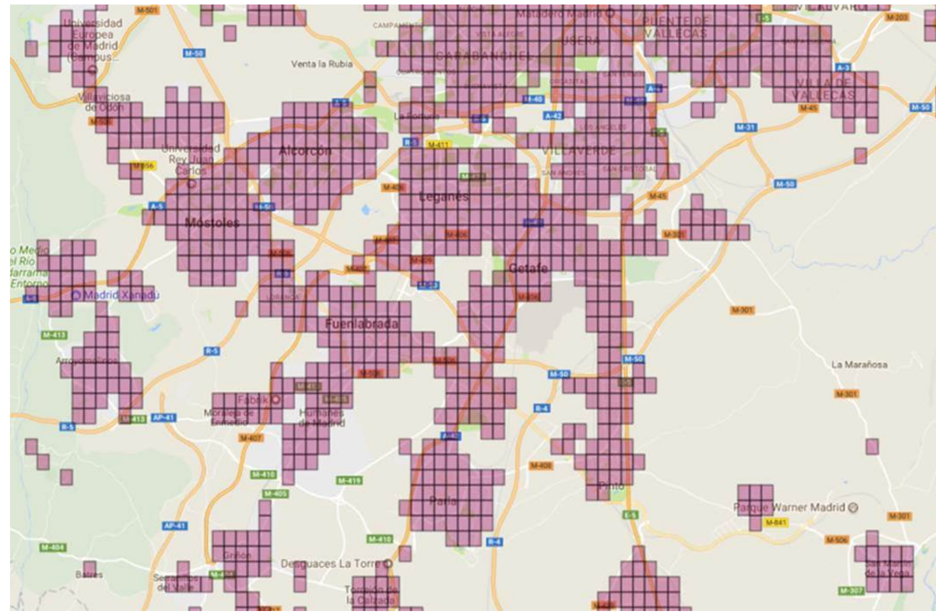

Venta la Rubia
Fuenlabrada
Getafe
Leganes
Parla
La Marañosa
Parque Warner Madrid
Desguaces La Torre

Representation: Mobile CDR

CDR (voice)

HR_ORG	TLFN_A	TLFN_B	CD_GEO_A	CD_GEO_B	DT_ORG	CD_SNTD	CD_ERB	CD_CCC	QT_DUR
20:05:31	XXX	YYY	3	11	20140519	2	1562	568	33
…	…	…	…	…	…	…	…	…	…

CDR (SMS)

HR_ORG	TLFN_A	TLFN_B	CD_GEO_A	CD_GEO_B	DT_ORG	CD_SNTD	QT_TRFG
15:53:54	XXX	ZZZ	3	25	20140506	2	1
…	…	…	…	…	…	…	…

Consumption	Social Network	Mobility
Call duration	In/Out Degree	Radius of gyration
N. Events	Delta w.r.t time window	Travelled distance
Lapse between events	Unique Calls per day	Rate of popular antennas
Reciprocated events	Unique SMS per day	Regularity of popular antennas



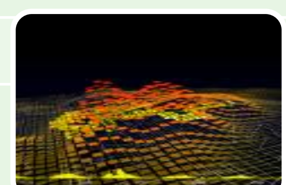

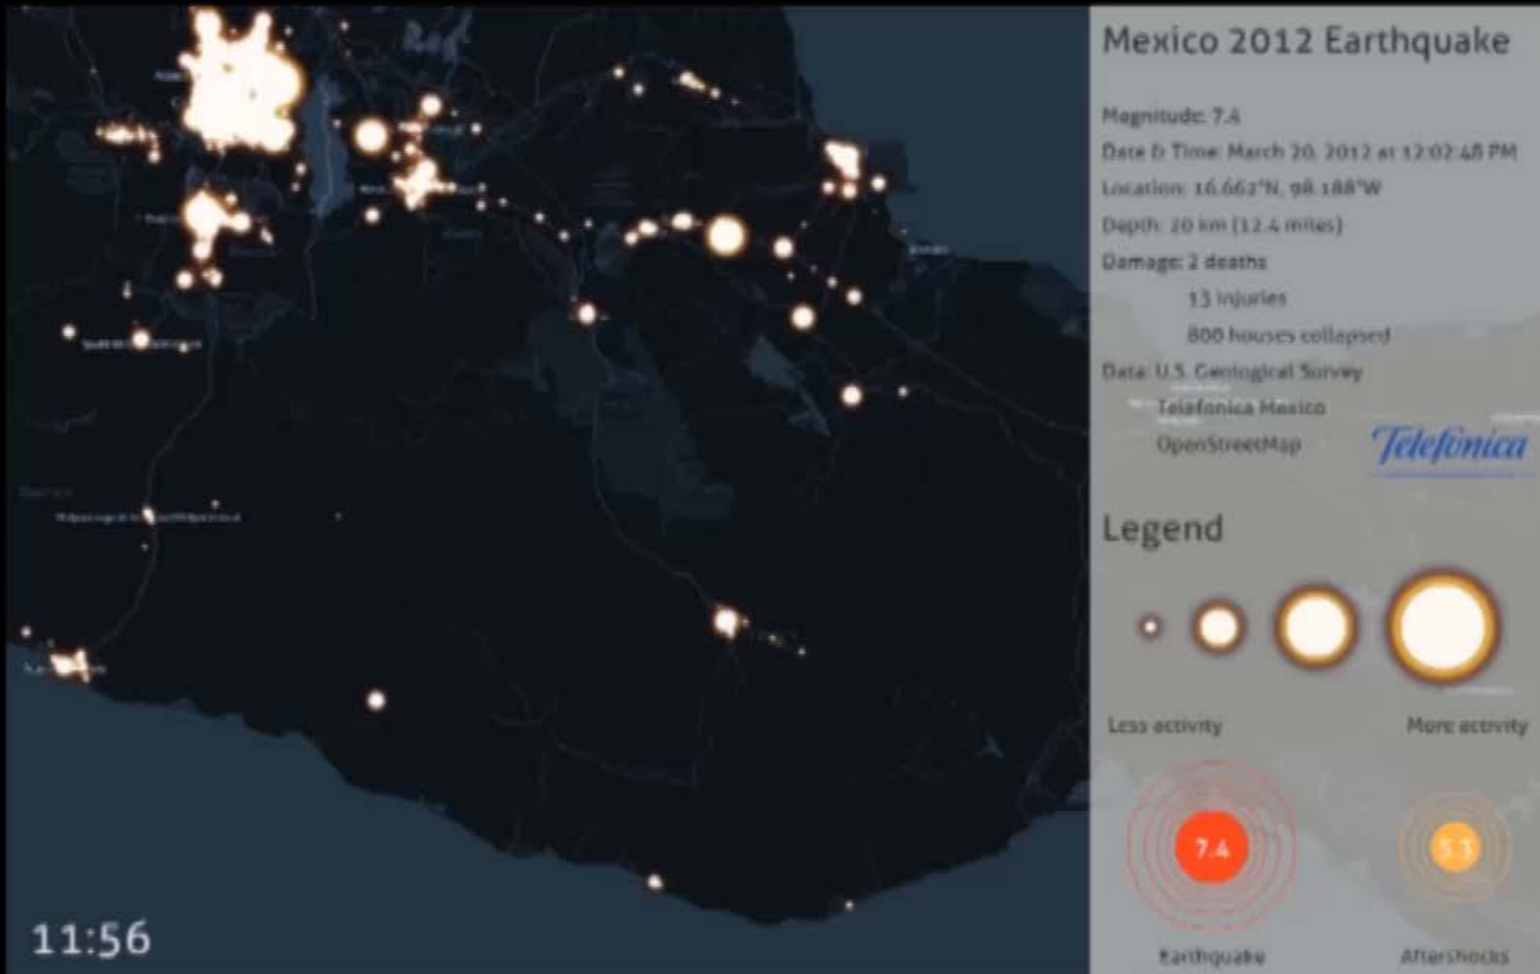
Mexico 2012 Earthquake
Magnitude: 7.4
Date & Time: March 20, 2012 at 12:02:48 PM
Location: 16.662°N, 98.188°W
Depth: 20 km (12.4 miles)
Damage: 2 deaths
13 injuries
800 houses collapsed
Data: U.S. Geological Survey
Telefonica Mexico
OpenStreetMap
Legend
Less activity
More activity
7.4
Earthquake
Aftershocks
11:56

Areas of impact


How can we help the refugees with data analytics and AI?
D4R
YOUR VOICE, YOUR DATA, YOUR FUTURE.
DATA FOR REFUGEES TURKEY IS
A BIG DATA CHALLENGE BY TURK TELEKOM

Example: Data for Refugees Challenge

- Collected from 200K refugees and 800K non-refugee residents, over the entire country, over one year.

- Includes:

  1- Cell tower locations, lists of cell towers for each prefecture

  2- Site-to-site antenna traffic on an hourly basis.

  Total number and duration of calls given, separated into “originating from refugees” and “not originating from refugees”.

  3- Fine grained mobility of a small subset of (anonymous) users. Only 15 days for each user, at cell tower level.

  4- Coarse grained mobility of a small subset of (anonymous) users. For the entire data collection period, but provided at prefecture level.

Salah, A. A., Pentland, A., Lepri, B., Letouzé, E., Vinck, P., de Montjoye, Y. A., Dong, X., Dağdelen, Ö. (2018). Data for Refugees: The D4R Challenge on Mobility of Syrian Refugees in Turkey. arXiv preprint arXiv:1807.00523.

Example: Data for Refugees Challenge

- **Safety & Security**
  Violence, theft, illegal trafficking
- **Health**
  Access to resources, spread of diseases, vaccination
- **Education**
  Access to education, language learning, schools
- **Unemployment**
  Movement due to unemployment, skill and resource management
- **Social Integration**
  Events and institutions for social integration, segregated/mixing patterns of behavior

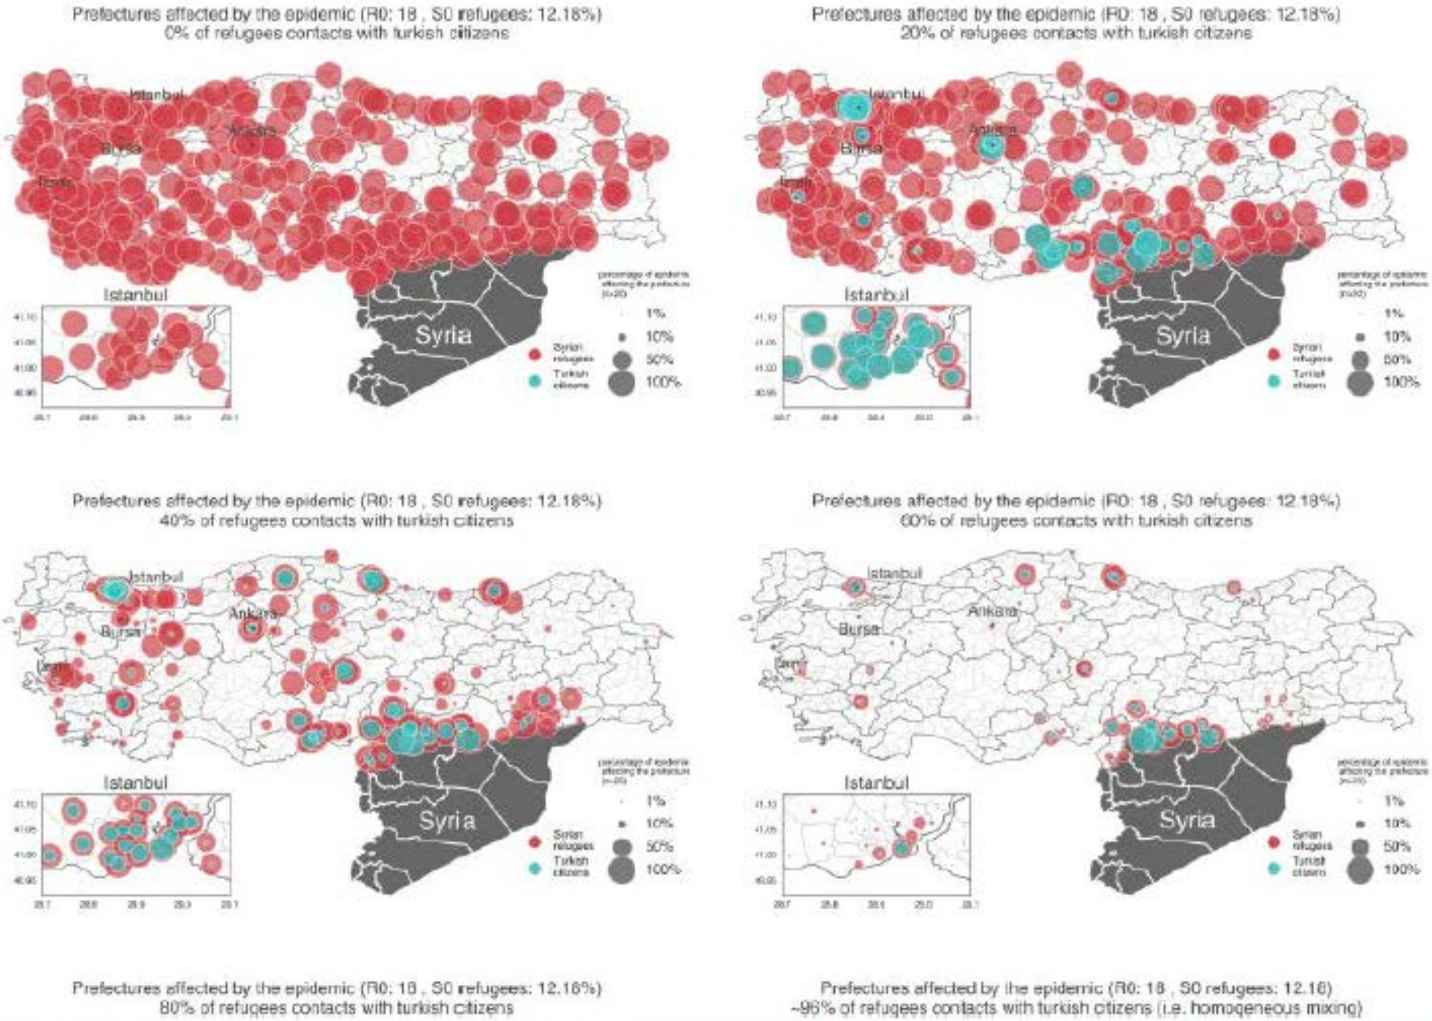

Bosetti, P., Poletti, P., Stella, M., Lepri, B., Merler, S., & De Domenico, M. (2019). Reducing measles risk in Turkey through social integration of Syrian refugees. In Data for Refugees Challenge Workshop.

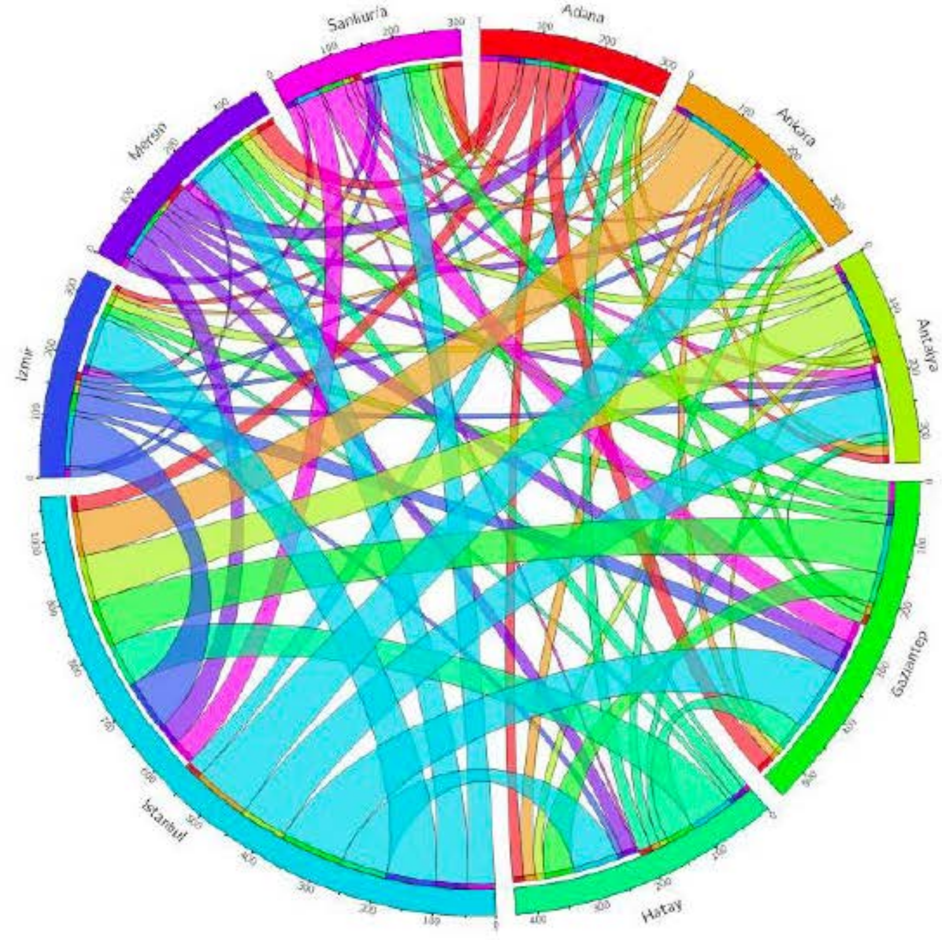

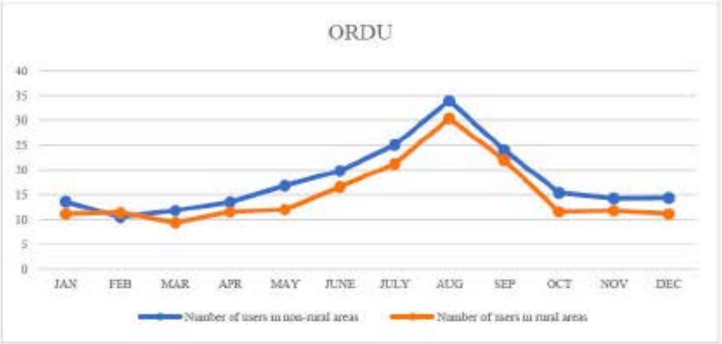

Alısık ST, Aksel DB, Yantaç AE, Baruh L, Salman S, Kayı I, Içduygu A, Bensason I (2019). UDMIT: an urban deep map for integration in Turkey. In Data for Refugees Challenge Workshop

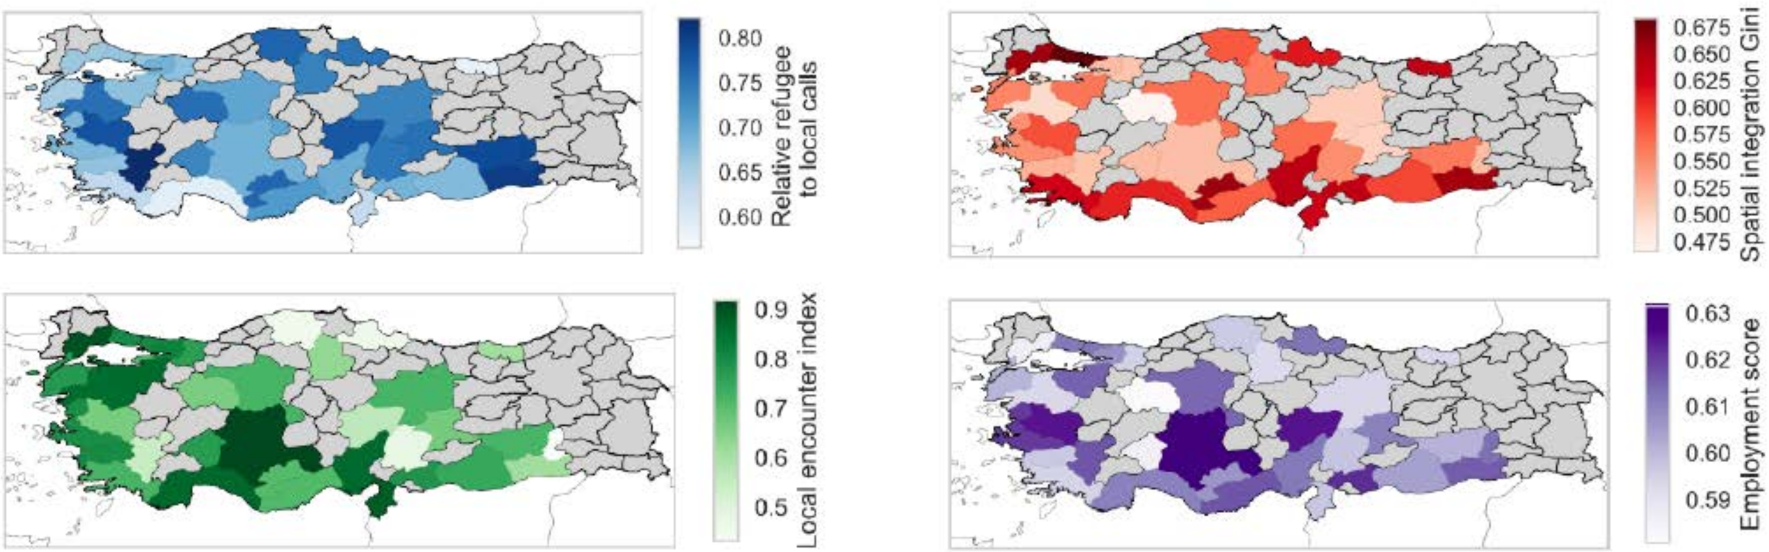

Bakker M, Piracha D, Lu P, Bejgo K, Bahrami M, Leng Y, Balsa-Barreiro J, Ricard J, Morales A, Singh V, Bozkaya B, Balcısoy S, Pentland A (2019) Measuring fine-grained multidimensional integration using mobile phone metadata: the case of Syrian refugees in Turkey. In: Data for Refugees Challenge Workshop

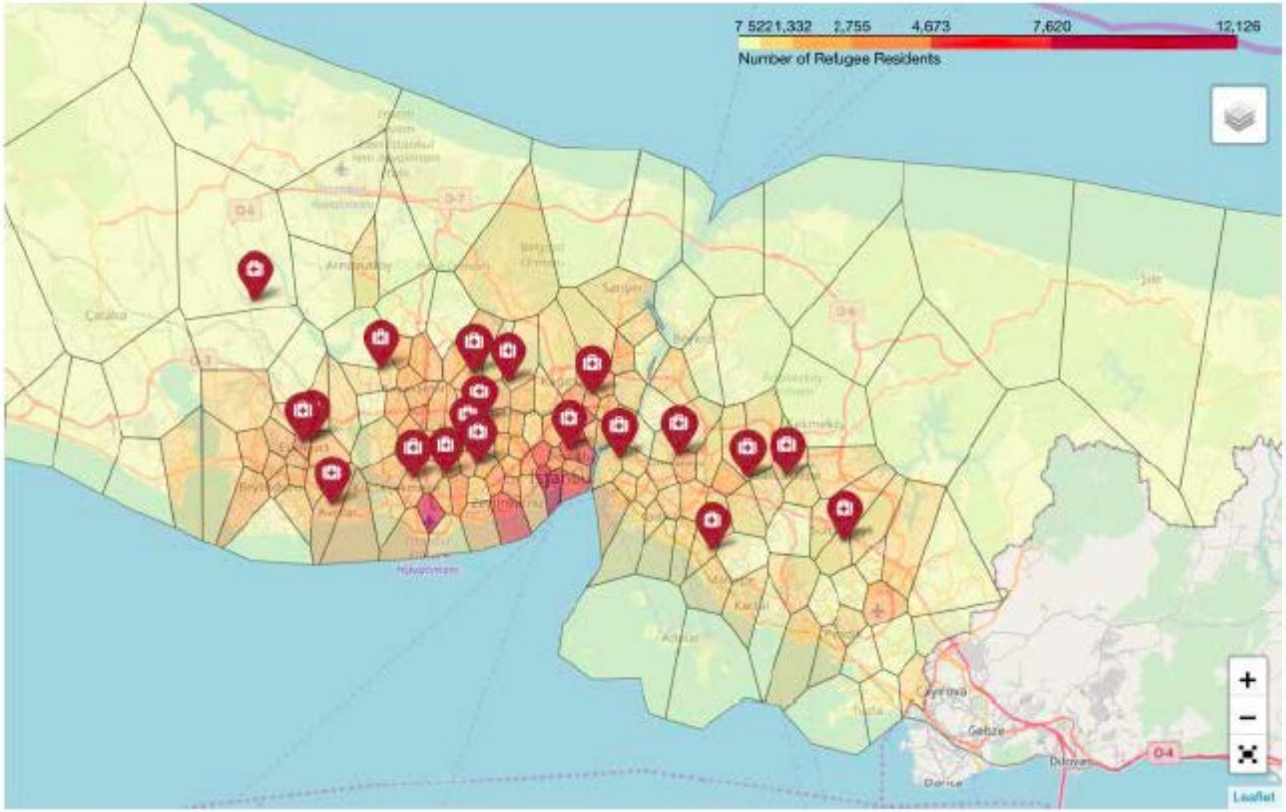

Fig. 2. Map with red pins showing the locations of current MHCs over the choropleth which illustrates the density of refugee residents at different districts in Istanbul.

Altuncu T, Sevencan N, Kaptaner AS (2019) Optimizing the access to healthcare services in dense refugee hosting urban areas: a case for Istanbul. In: Data for Refugees Challenge Workshop

Example: Data for Refugees Challenge

The 'Data for Refugees Challenge' Study

- Provides evidence-based insights into issues of refugee health, education, unemployment, social integration, safety, and security
- Serves as a sourcebook for refugee policy interventions based on big data analysis
- Describes best practices for ethically processing sensitive data on refugee mobility
- Presents results from the first big data challenge on refugees, offering insights into the dynamics of the Syrian refugee population in Turkey, currently the world's largest refugee population

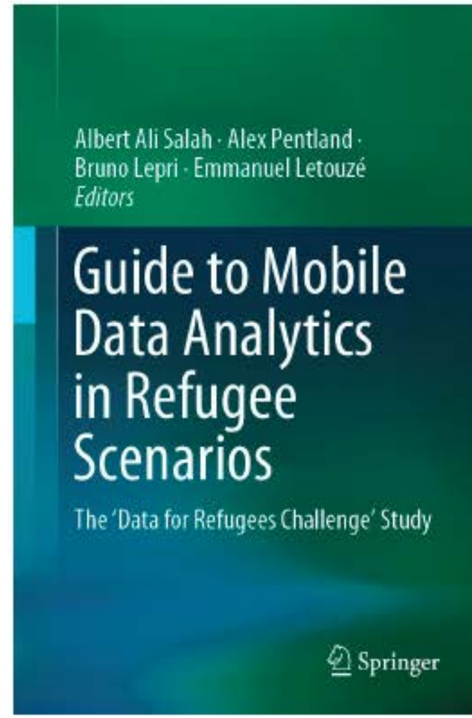

http://d4r.turktelekom.com.tr/Content/Documents/d4rproceedings.pdf

Large-scale Datasets: CDR

Dataset	Data	Year	#pp	Observations
D4D Ivory Coast	CDRs	2011-2012 (1 year)	5 million	Antenna-to-antenna traffic (hourly), 50k individual trajectories at antenna level; 500k individual trajectories at prefecture level; 5k comm graphs
D4D Senegal	CDRs	2013 (1 year)	9+million	Antenna-to-antenna traffic, 300k users mobility
Telecom Italia	CDRs, electricity, weather, rain, news, geolocated tweets	2014	300+k	Milan and Trento
Telefonica	Smart steps, hospital admission, transportation, Twitter, crime	2013 (3 weeks)	500+k	London metropolitan area
D4R Turk Telekom	CDRs, labeled as refugee & non-refugee	2018	1 million	Syrian refugee movements in Turkey
HummingBird Turkcell	CDRs	Est. 2021	-	To study migration from Turkey into Europe

Large-scale Datasets: Smartphone

Dataset	Data	Year	#pp	Observations
MIT Reality Mining	Nokia smartphones: Bluetooth devices, locations, call and SMS logs	2004-2005 academic year	100	MIT Students and faculty
Friends and Family	Smartphones: location labels, calls/SMS, BT proximity, periodic surveys	2010-2011	140	Mobile Territorial Lab
Mobile Data Challenge	Nokia smartphones: Bluetooth devices, locations, calls, SMS, apps and media usage, battery status, acoustic information	2009-2011	185	Young individuals
LiveLab	iPhone data: calls, SMS, web history, accelerometer, battery, display, app usage, cell tower/wifi ID	2010	25	College students
Device Analyzer	Android smartphones: apps, WiFi networks, battery, calls	2014 (no longer avail)	20,000	175+ countries
PhoneLab	Android smartphones: location, battery, WiFi, cell tower	2015	199+ 288	
Sensible DTU	FB data, school performance, smartphone Android app: WiFi, calls, SMS, BT proximity	2012-2013	1000	Students at TU Denmark
Copenhagen Networks Study	Android smartphones: BT proximity, calls, SMS, FB friendships	4 weeks, 2019	700+	Students at TU Denmark

Fundamental questions

- Why investigate behavior at this scale?
- How do we represent the behavior?
- What **machine learning models** are suitable to model large-scale behavior?

Learning: Approaches

Type of modeling	Approach	Papers
Supervised	Decision trees & Random Forests	Monreale A, Pinelli F, Trasarti R, Giannotti F (2009) WhereNext: a location predictor on trajectory pattern mining. Proc 15th ACM SIGKDD pp 637–646 Krumm J, Horvitz E (2006) Predestination: inferring destinations from partial trajectories. UbiComp 2006, Springer, pp 243–260 Etter V, Kafsi M, Kazemi E (2012) Been there, done that: What your mobility traces reveal about your behavior. Mobile data challenge by Nokia Workshop Khoroshevsky F, Lerner B (2017) Human mobility-pattern discovery and next-place prediction from GPS data. Schwenker F, Scherer S (eds) Multimodal pattern recognition of social signals in human computer interaction (MPRSS). Lecture notes in computer science, vol 10183. Springer, Berlin

Learning: Approaches

Type of modeling	Approach	Papers
Supervised	SVMs	Sohn T,, Varshavsky A, LaMarca A, Chen MY, Choudhury T, Smith I, Consolvo S, Hightower J, Griswold WG, De Lara E (2006) Mobility detection using everyday GSM traces. UbiComp 2006, Springer, pp 212–224 Li B, Zhang D, Sun L, Chen C, Li S, Qi G, Yang Q (2011) Hunting or waiting? Discovering Passenger finding strategies from a large-scale real-world taxi dataset. IEEE PERCOM, pp 63–68 Wang J, Prabhala B (2012) Periodicity based next place prediction. Proc. of the Nokia mobile data challenge workshop
	Neural Networks	Etter V, Kafsi M, Kazemi E (2012) Been there, done that: What your mobility traces reveal about your behavior. Mobile data challenge by Nokia Workshop
		Ben Zion, E. and Lerner, B. (2018), “Identifying and Predicting social lifestyles in People’s trajectories by neural Networks”, EPJ Data Science Feng, J., Li, Y., Zhang, C., Sun, F., Meng, F., Guo, A., Jin, D., “DeepMove: Predicting Human Mobility with Attentional Recurrent Networks”, WWW 2018, Lyon, France Jiang, R., Song, X., Huang, D., Song, X., Xia, T., Cai, Z., Wang, Z., Kim, K.S., Shibasaki, R., “DeepUrbanEvent: A system for predicting citywide crowd dynamics at big events”, KDD’19, Anchorage, USA

Learning: Approaches

Type of modeling	Approach	Papers
Unsupervised	Clustering methods	Ashbrook D, Starner T (2003) Using GPS to learn significant locations and predict movement across multiple users. Pers Ubiquitous Comput 7(5):275–286 Shoval N et al (2008) The use of advanced tracking technologies for the analysis of mobility in Alzheimer's disease and related cognitive diseases. BMC Geriatr 8(1):7 Andrienko N, Andrienko G, Stange H, Liebig T, Hecker D (2012) Visual analytics for understanding spatial situations from episodic movement data. Künstliche Intell 26(3):241–251 Ying JJ-C, Lee W-C, Tseng VS (2013) Mining geographic-temporal-semantic patterns in trajectories for location prediction. ACM Trans Intell Syst Technol (TIST) 5(1):2
	Other	Eagle, N. and Pentland, A., "Eigenbehaviors: identifying structure in routine", Behavioral ecology and sociobiology, 63 (7), pp 1057-1066, May 2009

Learning: Approaches

Type of modeling	Approach	Papers
Unsupervised	Topic Models (e.g. LDA)	Hariharan R, Toyama K (2004) Project lachesis: parsing and modeling location histories. Egenhofer, MJ et al. (eds) Geographic information science, Springer, pp 106–124 Farrahi K, Gatica-Perez D (2011) Discovering routines from large-scale human locations using probabilistic topic models. ACM Trans Intell Syst Technol (TIST) 2(1):3 Wang H, Fu Y,Wang Q, Yin H, Du C, Xiong H (2017) A location-sentiment-aware recommender system for both home-town and out-of-town users. Proc. 23rd ACM KDD, pp 1135–1143
		Ben-Zion E, Lerner B (2017) Learning human behaviors and lifestyle by capturing temporal relations in mobility patterns. Proc. of the European symposium on artificial networks, computational intelligence and machine learning (ESANN), Bruges

Unsupervised: Eigenbehaviors

- **Reality Mining Dataset**:
  - Call logs, Bluetooth devices in proximity, cell tower IDs, app usage and phone status (charging vs idle)
  - **400,000 h** of location, communication, app usage behavior for 100 people

- **Approach: Principal components analysis** is performed on standard size behavior vectors

- The vectors with **the highest eigenvalues** are considered an individual's primary **eigen-behaviors**

Eagle, N. and Pentland, A., "Eigenbehaviors: identifying structure in routine", Behavioral ecology and sociobiology, 63 (7), pp 1057-1066, May 2009

Unsupervised: Eigenbehaviors

- Top 3 eigenbehaviors for 1 person (columns)
- First eigenbehavior: home vs work
- Second eigenbehavior: weekend behavior
- Third eigenbehavior: no signal (country side/indoors?)

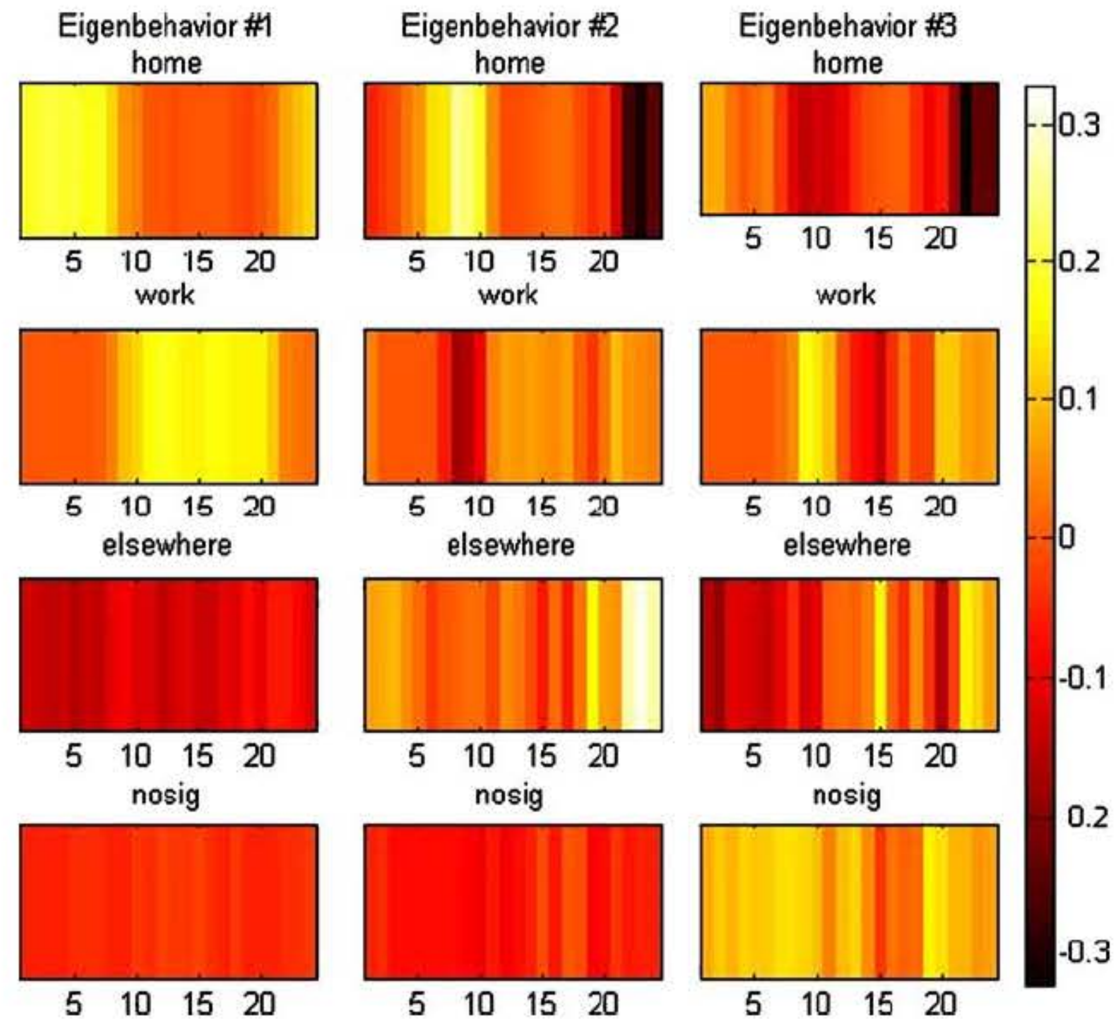

Eagle, N. and Pentland, A., "Eigenbehaviors: identifying structure in routine", Behavioral ecology and sociobiology, 63 (7), pp 1057-1066, May 2009

Unsupervised: Eigenbehaviors

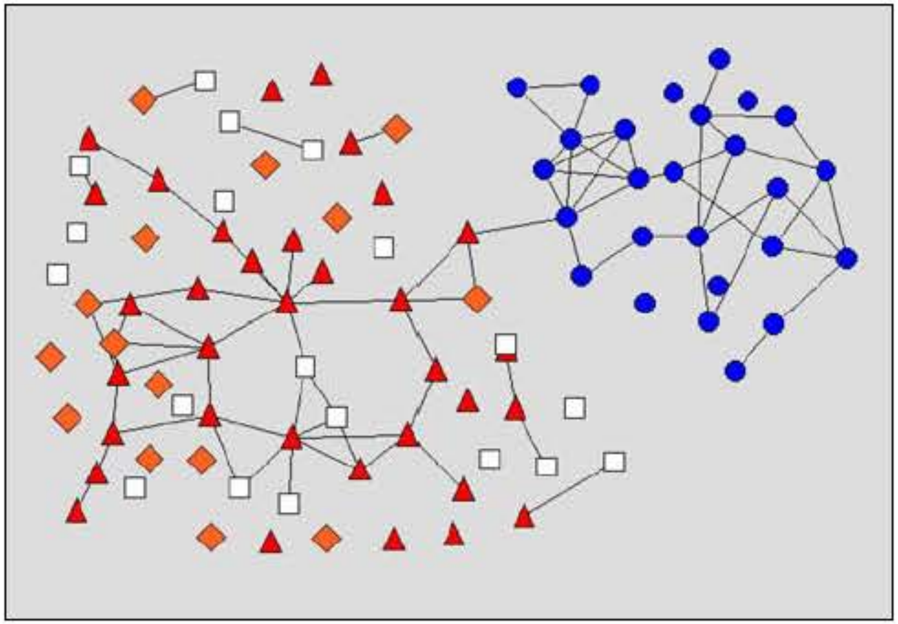

- ● business school students
- ▲ senior lab students
- ◆ incoming students
- □ lab staff/faculty

- Top 3 eigenbehaviors for each group
- Business students: coffee break
  New students: stay later in lab

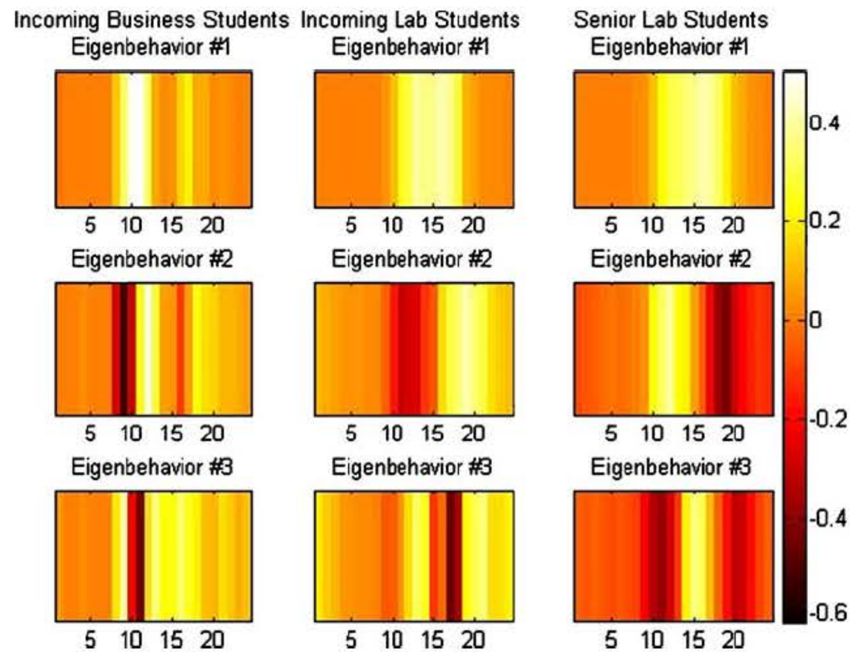

Eagle, N. and Pentland, A., "Eigenbehaviors: identifying structure in routine", Behavioral ecology and sociobiology, 63 (7), pp 1057-1066, May 2009

Unsupervised: Eigenbehaviors

- An **individual's behavior** over a specific day can be approximated by a **weighted sum** of his or her primary eigenbehaviors. When these weights are calculated **halfway through** a day, they can be used to predict the day's remaining behaviors with **79% accuracy** for test subjects

- **Clustering individuals** into a "behavior space" make it possible to determine the behavioral similarity between both individuals and groups, enabling **96% classification** accuracy of community affiliations within the population-level social network

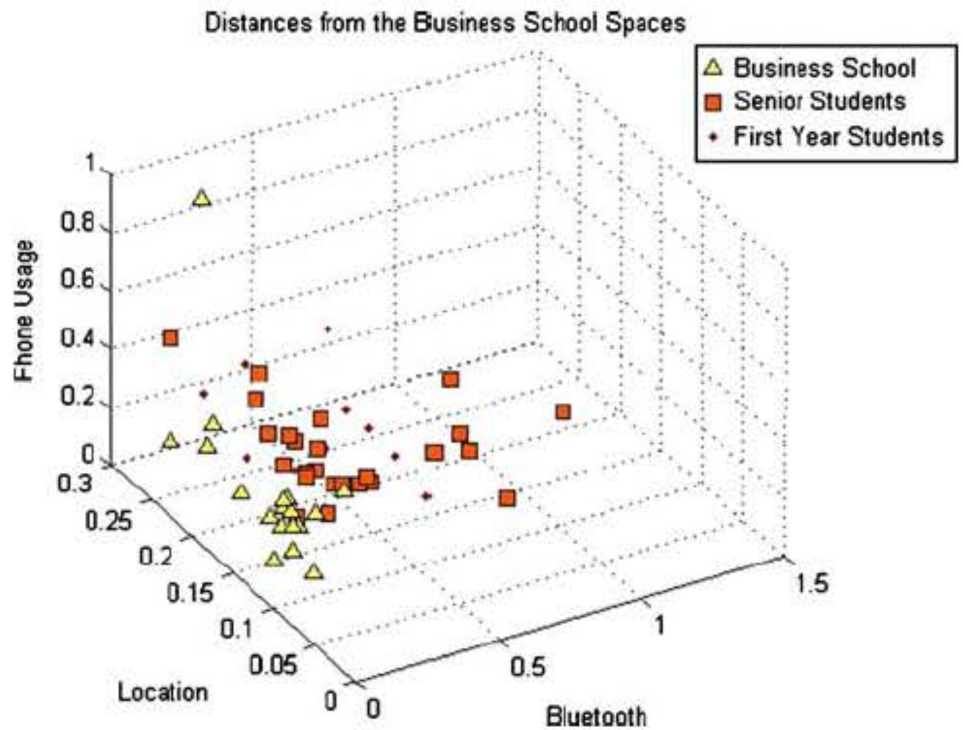

Distance between the three groups of students in the BT, location and phone usage behavior space

Eagle, N. and Pentland, A., "Eigenbehaviors: identifying structure in routine", Behavioral ecology and sociobiology, 63 (7), pp 1057-1066, May 2009

Two Set of Challenges

Technical Challenges

Are we able to automatically interpret and predict complex human behavior using machine learning techniques?

Human(ity) Challenges

What are the social implications and ethical considerations in the deployment and wide-spread use of these tools?

HUMAN(-ITY) CENTRIC CHALLENGES

Additional factors

❖ Human behavior modeling and prediction has left the lab and is part of today's intelligent services and systems, including self-driving cars, personal assistants, smart speakers, recommender systems, camera apps, search engines, visual surveillance, social robots…

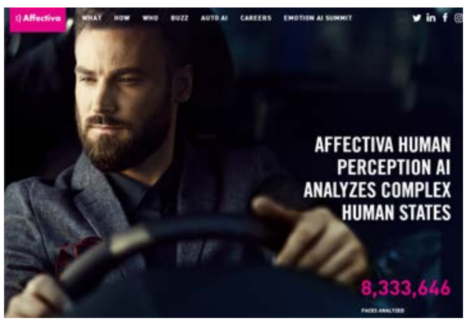




Additional factors

- Human behavior modeling and prediction has the potential to significantly improve people's lives but…

- Important considerations need to be considered beyond the technical factors, namely:

  - Computational violations of privacy
  - Bias, discrimination and social exclusion
  - Asymmetry
  - Opacity
  - Veracity
  - Ethics

Wrapping up…

- Human behavior modeling and prediction via machine learning is an exciting area with a lot of opportunities…but…
- Human behavior is very complex and multi-faceted. No human should be reduced to a 'data point' (or a lot of data points!)
  - Individual
  - Interactive, small groups
  - Aggregate, computational social sciences
- Human-centric approaches are a must
- Societal implications must be considered

Slide credits

- Slides are (mostly) based on N. Oliver, A.A. Salah, “Human Behavior Understanding with Machine Learning: Challenges and Opportunities”, invited tutorial @ NeurIPS 2019.