Northeastern University

Large-Scale Social Media Analytics

Yun Raymond Fu
Assistant Professor
Electrical and Computer Engineering (ECE), COE
College of Computer and Information Science (CCIS)
Northeastern University

Motivations

Human-Centered Computing

Human-Computer Interaction

- Computer Vision Systems
  LabelRelation, M-face, EAVA, hMouse, Facetransfer, RTM-HAI, Shrug Detector.

Human Centered Computing

Machine Learning

- Manifold/Subspace Learning
- Transfer Learning
- Low-Rank Matrix Analytics
- Sparse Representation
- Large-Scale Optimization

Social Media Analytics

- Demographic Recognition
- Internet Vision
- Action/Activity/Intention Analysis
- Geolocation from Social Context
- Social Network Analysis

Smart Environments

- Human-Centered Cyber-Physical Systems
- Health Care
- Intelligent Systems

Motivation 1: Smart Environment

Wikipedia.com: conceptually a physical world that is richly and invisibly interwoven with sensors, actuators, displays, and computational elements, embedded seamlessly in the everyday objects of our lives, and connected through a continuous network…

The image is from http://sedl.kaist.ac.kr/images/smart_architecture_spaces.jpg

Motivation 2: Social Media in the Cloud

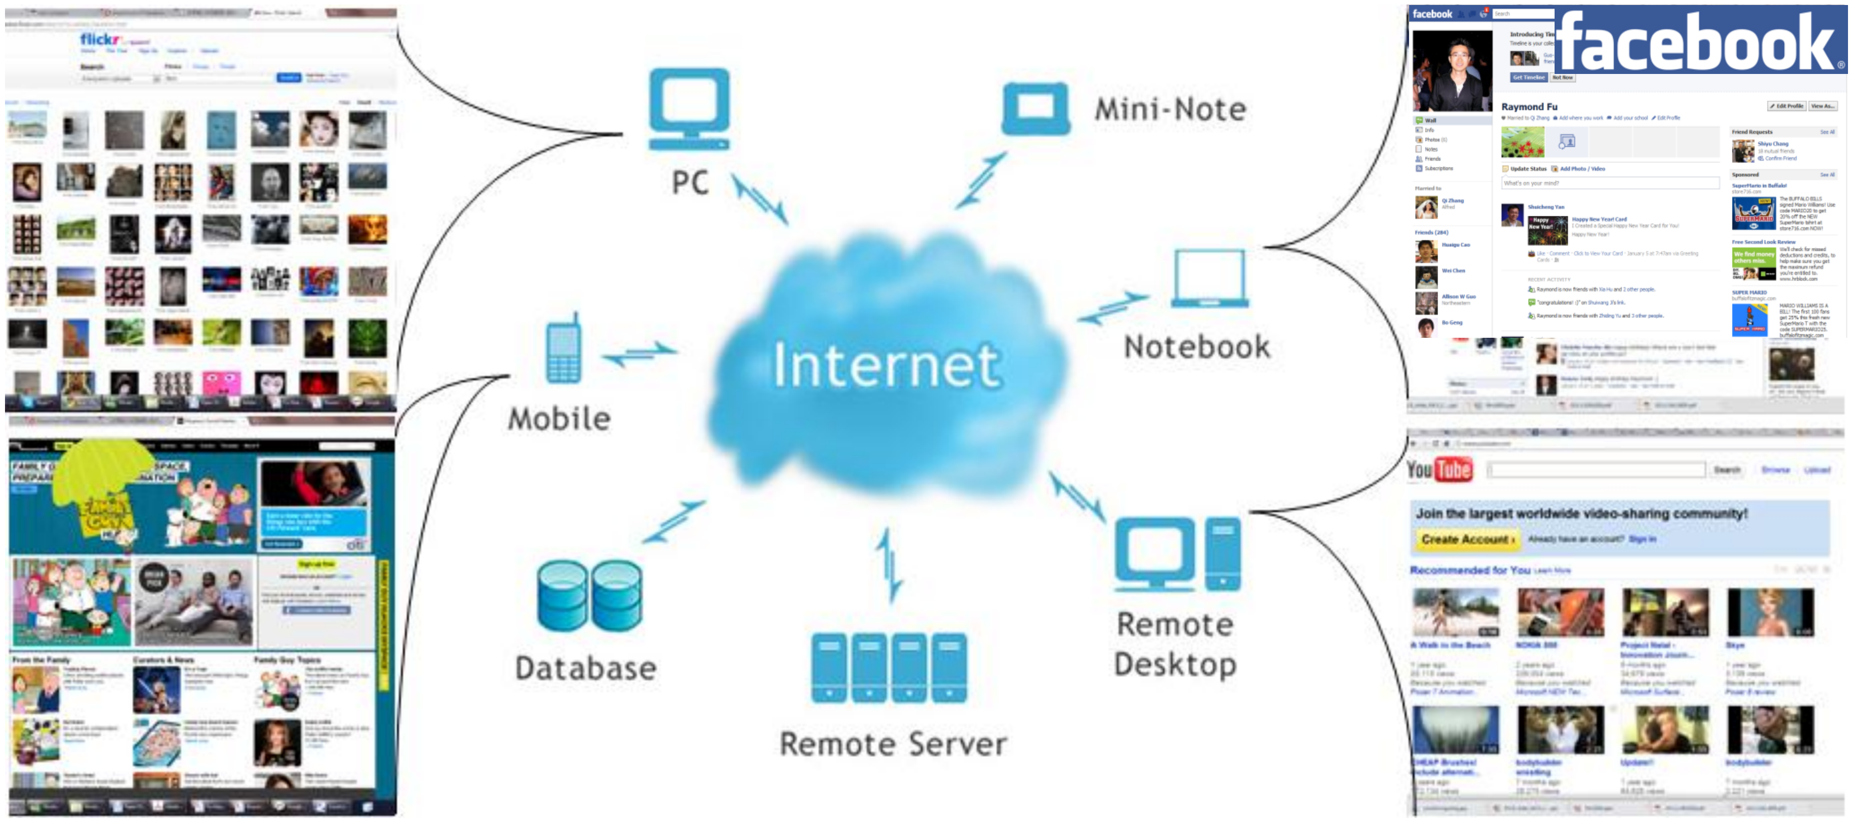

- How to model the multi-label, multi-instance, and multi-task characteristics?
- How to effectively infer meaningful user information from large scale visual data?
- How to provide targeted services through human-computer interactions?

Motivation 3: Multi-label Social Media

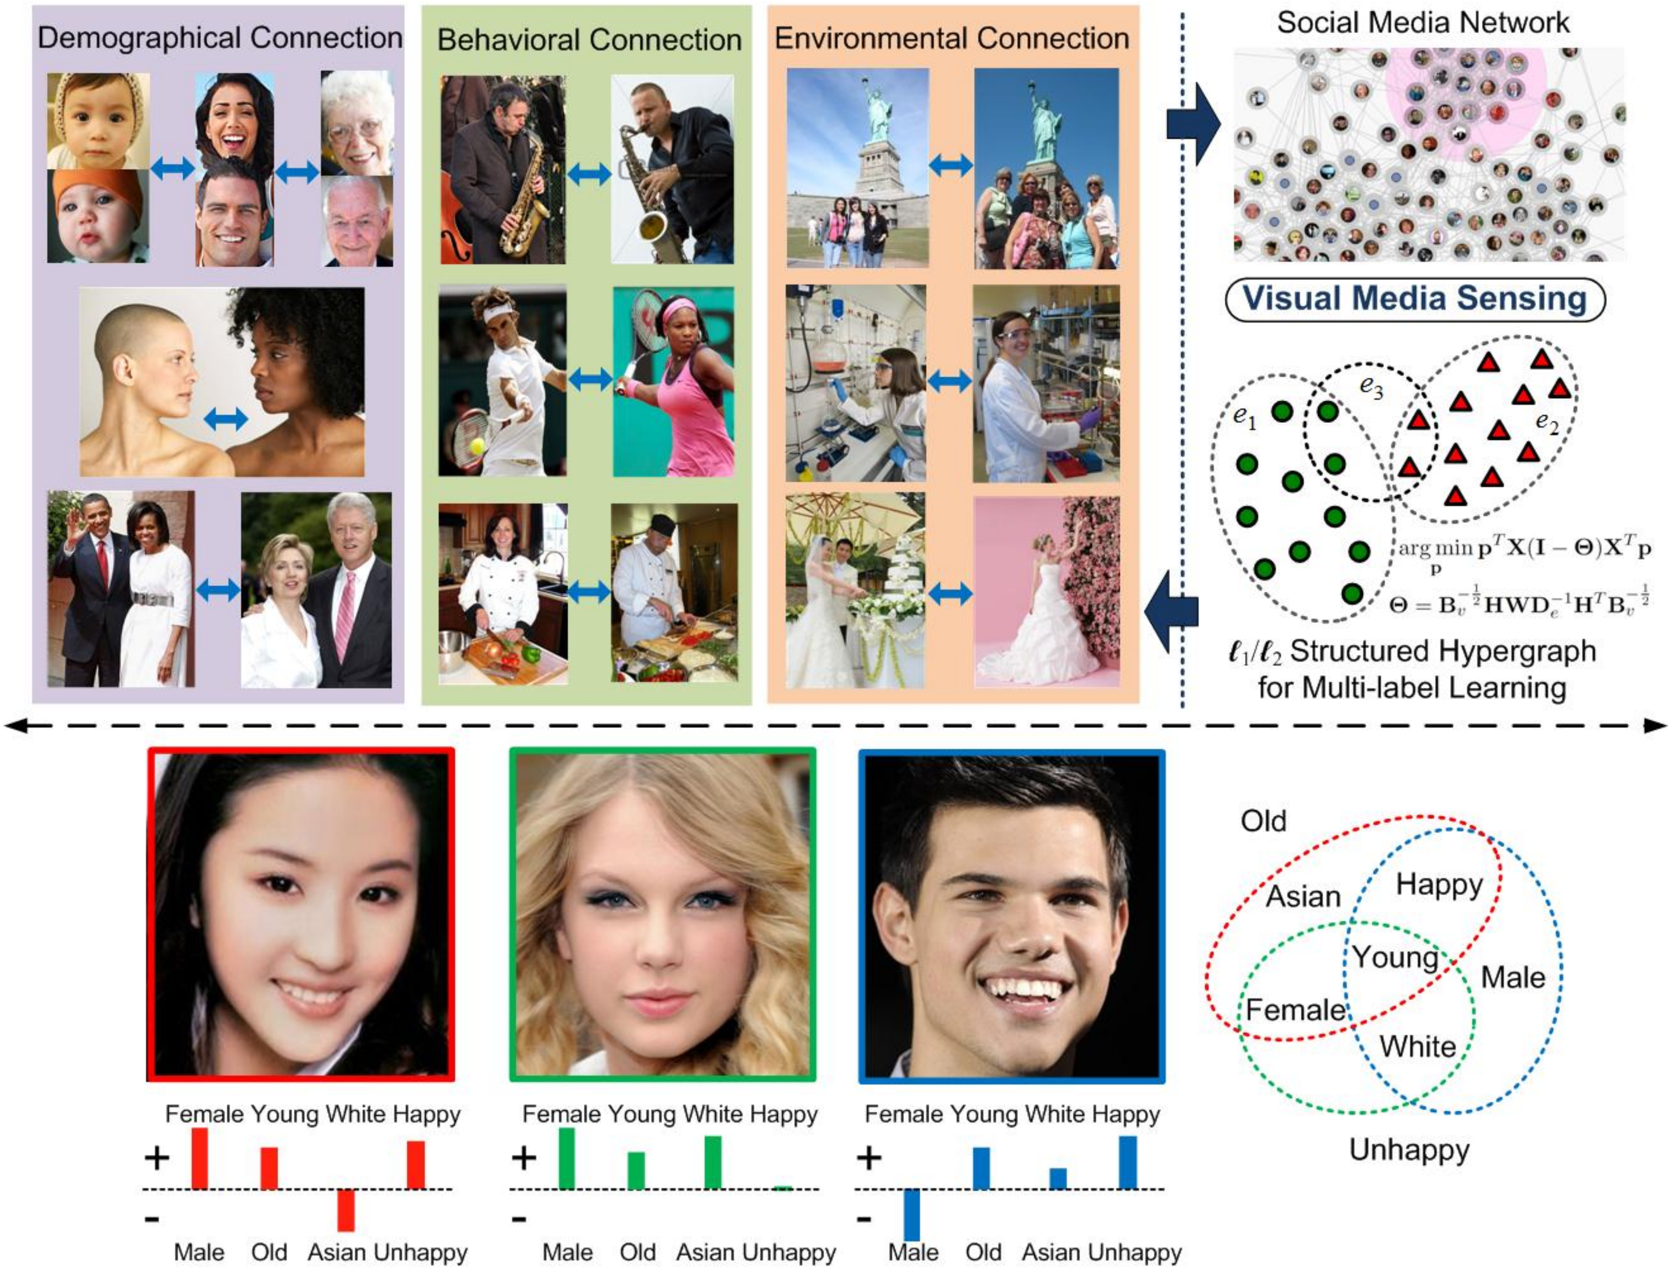

It Is All About Data!

- Goal: Interpret given human images in terms of demographic and behavioral attributes (Expression, Age, Gender, Occupation, Kinship, Action, Pose, and Intention, etc.).

- Challenge
  - Dimensionality redundancy
  - Large scale (big data)
  - Unknown distribution
  - Large attributes variations
  - Multimodality , multi-source, multi-label data
  - Noise and outliers

Methodologies
for
Social Media Computing

Background: Existing Methods

- **Global and Local Learning Methods**
  - Local Learning vs. Global Learning, K. Huang, H. Yang, I. King, and M. R. Lyu; Global Versus Local Methods in Nonlinear Dimensionality Reduction, V. de Silva and J. Tenenbaum; Generalized principal component analysis (GPCA), Y. Ma, et. al.; Globally-Coordinated Locally-Linear Modeling, C.-B. Liu.
- **Localized Subspace Learning Methods**
  - Locally Embedded Linear Subspaces, Z. Li, L. Gao, and A. K. Katsaggelos; Locally Adaptive Subspace, Y. Fu, Z. Li, T.S. Huang, A.K. Katsaggelos.
- **Patches/Parts Based Methods**
  - Flexible X-Y Patches, M. Liu, S.C. Yan, Y. Fu, and T. S. Huang; Patch-based Image Correlation, G-D. Guo and C. Dyer.
- **Feature Extraction Methods**
  - Local Binary Pattern (LBP), T. Ojala, M. Pietikainen, and T. Maenpaa; Histogram of Oriented Gradient descriptor (HOG), N. Dalai and B. Triggs.
- **Nonlinear Graph Embedding Methods**
  - Locally Linear Embedding (LLE), S.T. Roweis & L.K. Saul; Isomap, J.B. Tenenbaum, V.de Silva, J.C. Langford; Laplacian Eigenmaps (LE), M. Belkin & P. Niyogi
- **Linear Subspace Learning Methods**
  - Principal Component Analysis (PCA), M.A. Turk & A.P. Pentland; Multidimensional Scaling (MDS), T.F. Cox and M.A.A. Cox; Locality Preserving Projections (LPP), X.F. He, S.C. Yan, Y.X. Hu
- **Fisher Graph Methods**
  - Linear Discriminant Analysis (LDA), R.A. Fisher; Marginal Fisher Analysis (MFA), S.C. Yan, et al.; Local Discriminant Embedding (LDE), H.-T. Chen, et al.
- **Tensor Subspace Learning Methods**
  - Two-dimensional PCA (TPCA), J. Yang, et.al.; Two-dimensional LDA (TLDA), J. Ye, et.al.; Tensor subspace analysis (TSA), X. He, et al.; Tensor LDE (TLDE), J. Xia, et al.; Rank-r approximation, H. Wang.
- **Correlation-based Subspace Learnng Methods**
  - Discriminative Canonical Correlation (DCC), T.-K. Kim, et al.; Correlation Discriminant Analysis (CDA), Y. Ma, et al.

Graph Embedded Multilabel Learning

Machine Learning Framework

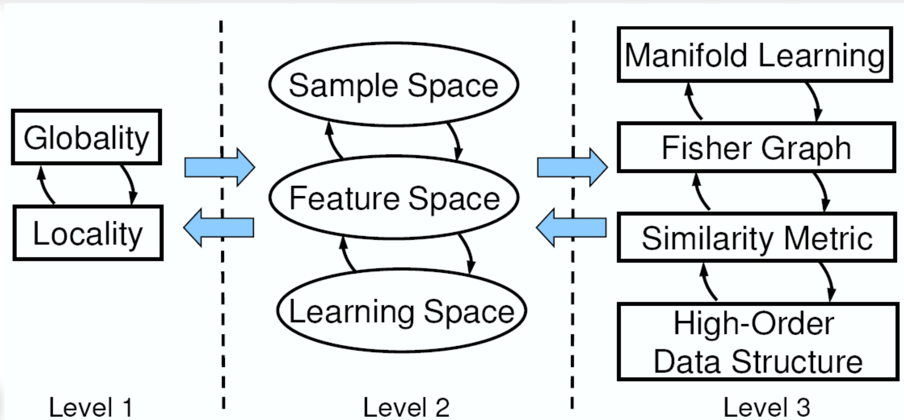

Subspace Learning

Demographic Recognition

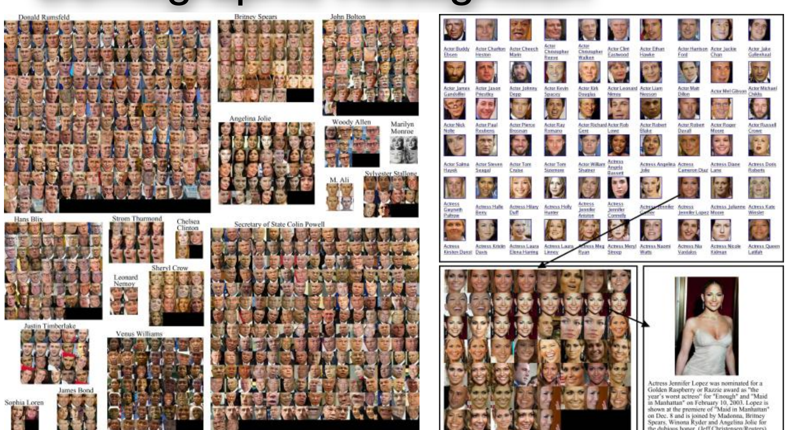

Courtesy of Tamara Berg

Inference

Emotion/Expression Analysis

Age/Gender Estimation

Ethnic Group Recognition

Kinship Recognition

Occupation Recognition

Human-Centered Computing

Level 3: Manifold Learning

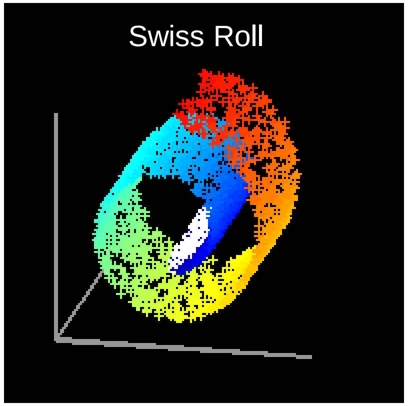

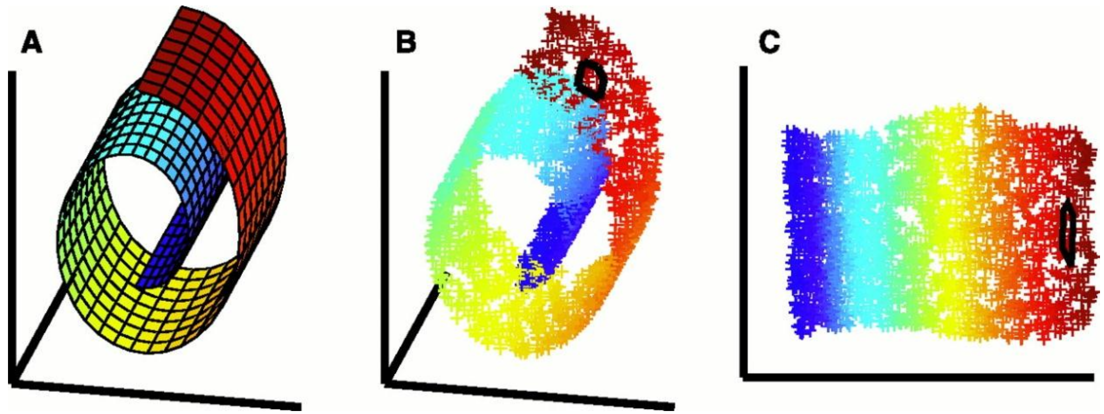

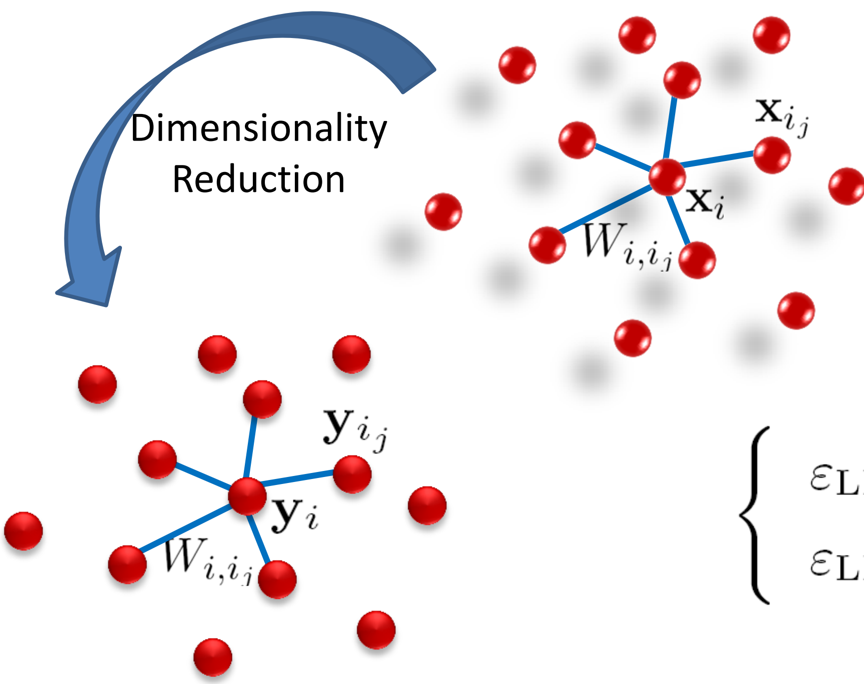

$$\varepsilon(W) = \sum_{i=1}^{n} \left\| \mathbf{x}_i - \sum_{j=1}^{k} W_{i,i_j} \mathbf{x}_{i_j} \right\|^2,$$

$$\begin{cases} \varepsilon_{\mathrm{LLE}} = \sum_{i=1}^{n} \left\| \mathbf{y}_i - \sum_{j=1}^{k} W_{i,i_j} \mathbf{y}_{i_j} \right\|^2 \\ \varepsilon_{\mathrm{LEA}} = \sum_{i=1}^{n} \left\| \mathbf{P}^T \mathbf{x}_i - \sum_{j=1}^{k} W_{i,i_j} \mathbf{P}^T \mathbf{x}_{i_j} \right\|^2 \end{cases}$$

Courtesy of Sam T. Roweis and Lawrence K. Saul, Sience 2002

Level 3: Fisher Graph

- Graph Embedding (*S. Yan, IEEE TPAMI, 2007*)
  - *G*={*X*, *W*} is an undirected weighted graph.
  - *W* measures the similarity between a pair of vertices.
  - Laplacian matrix $L = D - W, \quad D_{ii} = \sum_{j \neq i} W_{ij}, \quad \forall\, i.$
  - Most manifold learning method can be reformulated as

$$y^* = \arg \min_{y^T By = d} \sum_{i \neq j} \|y_i - y_j\|^2 W_{ij} = \arg \min_{y^T By = d} y^T L y,$$

where *d* is a constant and *B* is the constraint matrix.

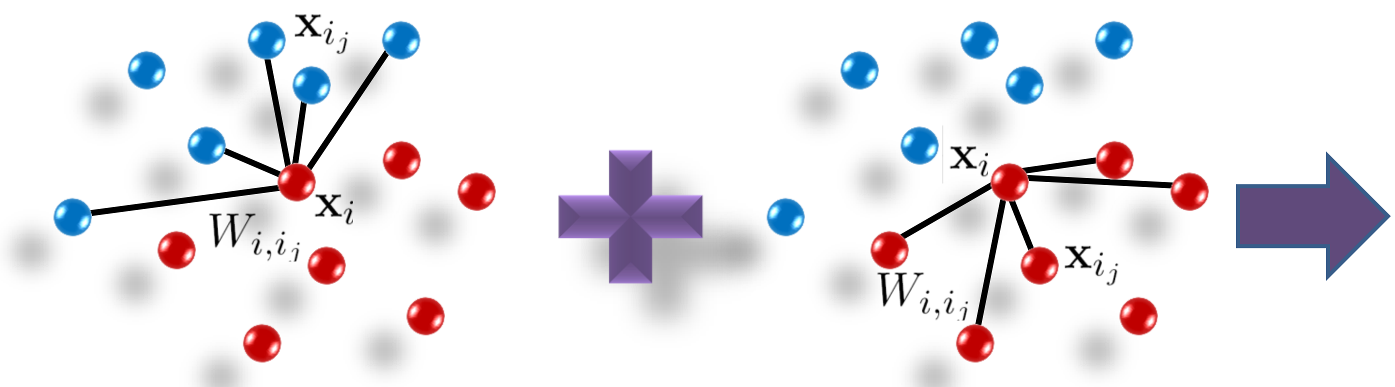

$$\begin{cases} \arg\min \mathbf{y}^T \mathbf{L}_{\mathrm{w}} \mathbf{y} \\ \arg\max \mathbf{y}^T \mathbf{L}_{\mathrm{b}} \mathbf{y} \end{cases}$$

Between-Locality Graph

Within-Locality Graph

Courtesy of Shuicheng Yan

Discriminant Simplex Analysis

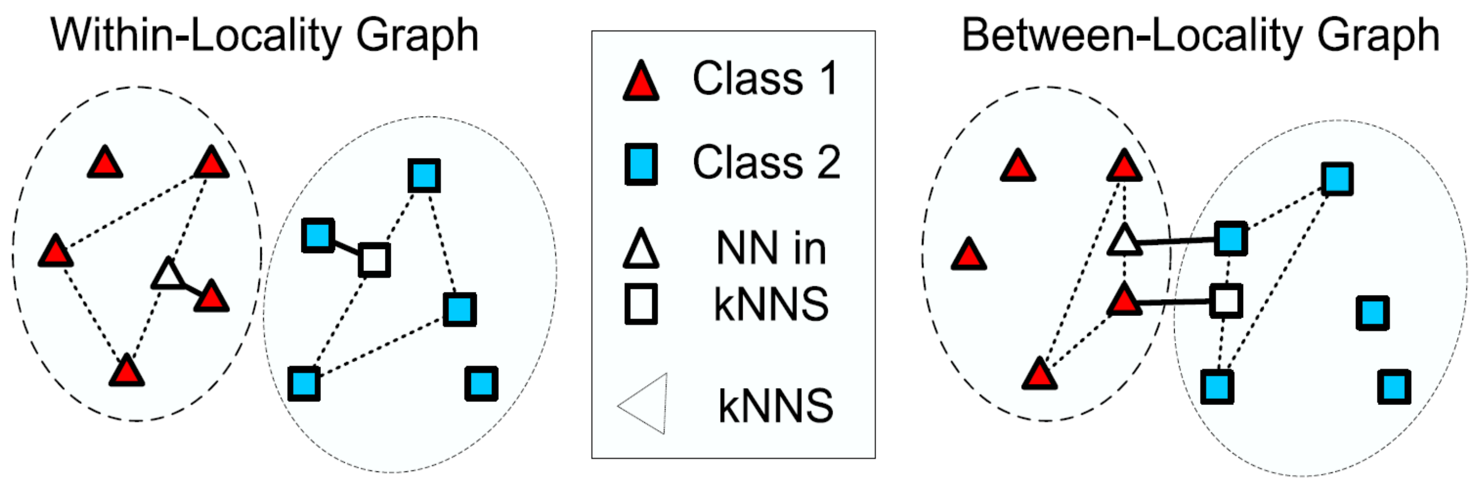

$$\mathbf{x} \Rightarrow \tilde{\mathbf{y}} = \arg\max \frac{\varepsilon_{\mathrm{b}}(\mathbf{y})}{\varepsilon_{\mathrm{w}}(\mathbf{y})} = \arg\max \frac{\mathbf{y}^T\mathbf{L}_{\mathrm{b}}\mathbf{y}}{\mathbf{y}^T\mathbf{L}_{\mathrm{w}}\mathbf{y}}.$$

$$\begin{cases} \varepsilon_{\mathrm{w}}(\mathbf{P}) = \sum_{i=1}^{n} \left\|\mathbf{P}^T\mathbf{x}_i - \sum_{j=1}^{k_{\mathrm{w}}} c_{\mathrm{w}(j)}^{(i)}\mathbf{P}^T\mathbf{x}_{\mathrm{w}(j)}^{(i)}\right\|^2 \\ \varepsilon_{\mathrm{b}}(\mathbf{P}) = \sum_{i=1}^{n} \left\|\mathbf{P}^T\mathbf{x}_i - \sum_{j=1}^{k_{\mathrm{b}}} c_{\mathrm{b}(j,l^*)}^{(i)}\mathbf{P}^T\mathbf{x}_{\mathrm{b}(j,l^*)}^{(i)}\right\|^2 \end{cases}$$

$$\mathbf{X}\mathbf{L}_{\mathrm{b}}\mathbf{X}^T\mathbf{P} = \mathbf{X}\mathbf{L}_{\mathrm{w}}\mathbf{X}^T\mathbf{P}\Lambda$$

Y. Fu, et. al., IEEE Transactions on Information Forensics and Security, 2008.

Level 3: Similarity Metric

- Single-Sample Metric
  - Euclidean Distance and Pearson Correlation Coefficient.

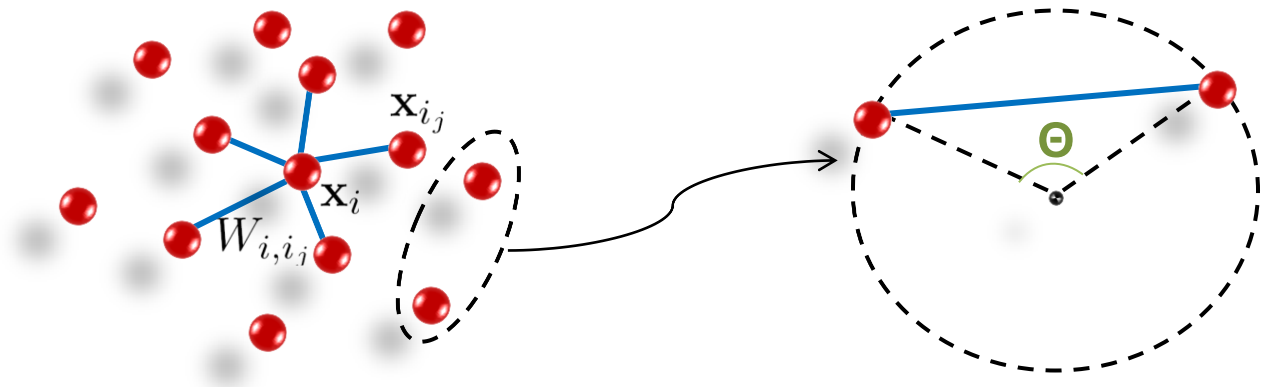

- Multi-Sample Metric
  - k-Nearest- Neighbor Simplex $\mathcal{S}(\mathbf{x}_1, \mathbf{x}_2, \cdots, \mathbf{x}_k) = \{\sum_{i=1}^{k} \ell_i \mathbf{x}_i \mid \sum_{i=1}^{k} \ell_i = 1, \ell_i \geq 0\}.$

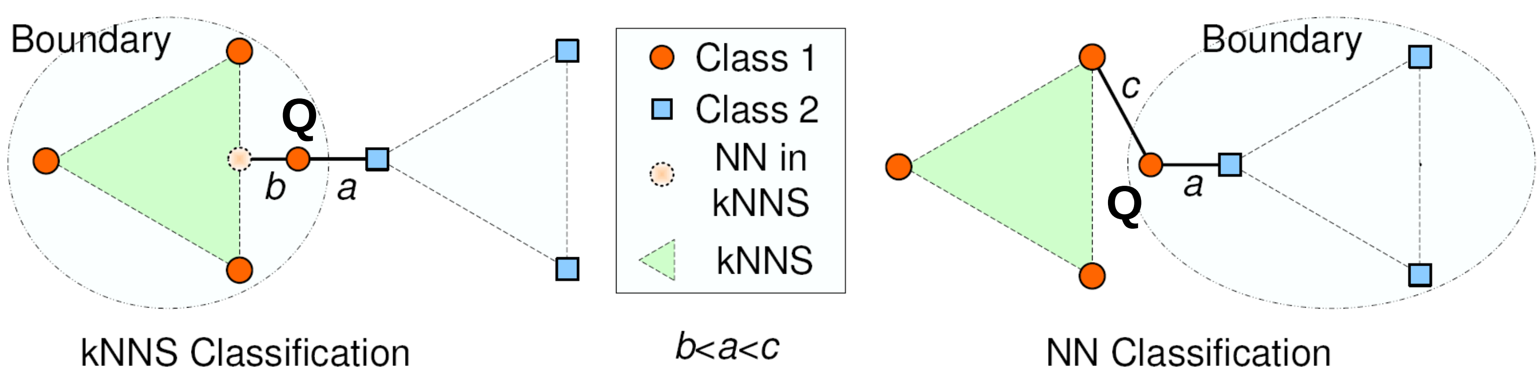

Correlation Embedding Analysis

- Objective Function

$$\arg\max_{\mathbf{P}} \quad \varepsilon(\mathbf{P}) = \sum_{i,j=1}^{n} \underbrace{(1 - \frac{\mathbf{y}_i^T \mathbf{y}_j}{\| \mathbf{y}_i \| \| \mathbf{y}_j \|})}_{\text{Correlation Distance}} \cdot \underbrace{(w_{ij}^{(\mathrm{d})} - w_{ij}^{(\mathrm{s})})}_{\text{Fisher Graph}}$$

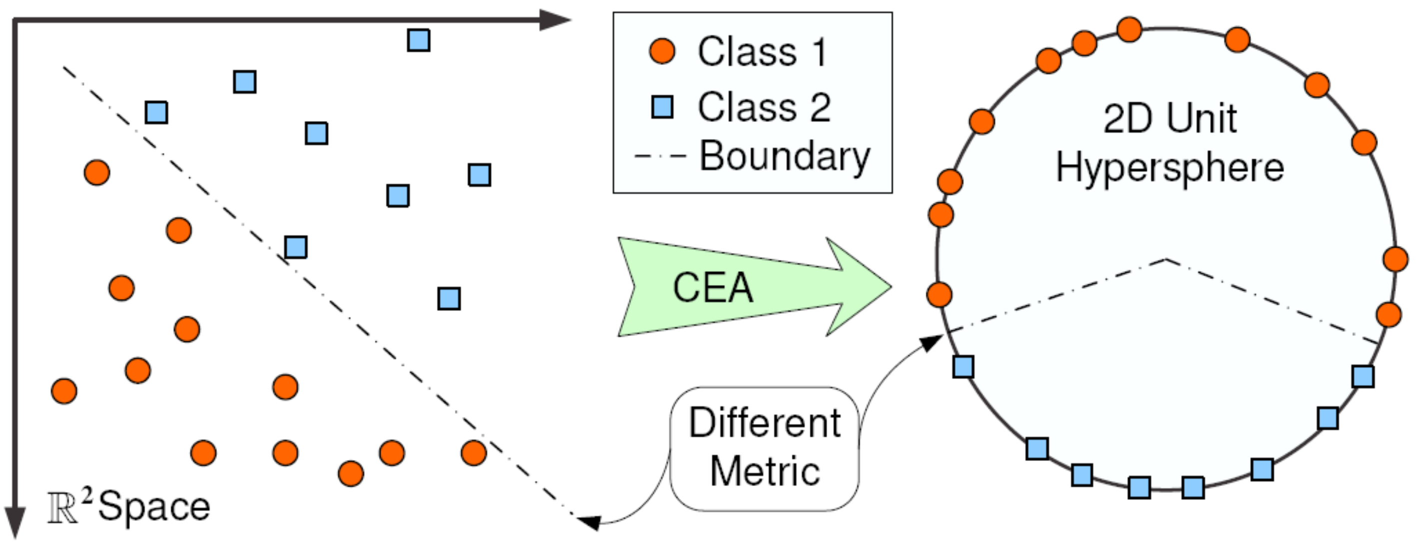

Y. Fu, et. al., IEEE Transactions on Pattern Analysis and Machine Intelligence, 2008.

Level 3: High-Order Data Structure

- *m*-th order tensors $\{\mathbf{X}_i \mid \mathbf{X}_i \in \mathbb{R}^{D_1 \times D_2 \times \cdots \times D_m}\}_{i=1}^n$
- Representation $\{\mathbf{Y}_i \mid \mathbf{Y}_i \in \mathbb{R}^{d_1 \times d_2 \times \cdots \times d_m}\}_{i=1}^n$ where $d_i \leqslant D_i$
- Define $\mathbf{Y}_i = \mathbf{X}_i \times_1 \mathbf{U}_1 \times_2 \ldots \times_m \mathbf{U}_m$ where $\mathbf{U}_j \in \mathbb{R}^{D_j \times d_j}$
- Here, tensor means multilinear representation.

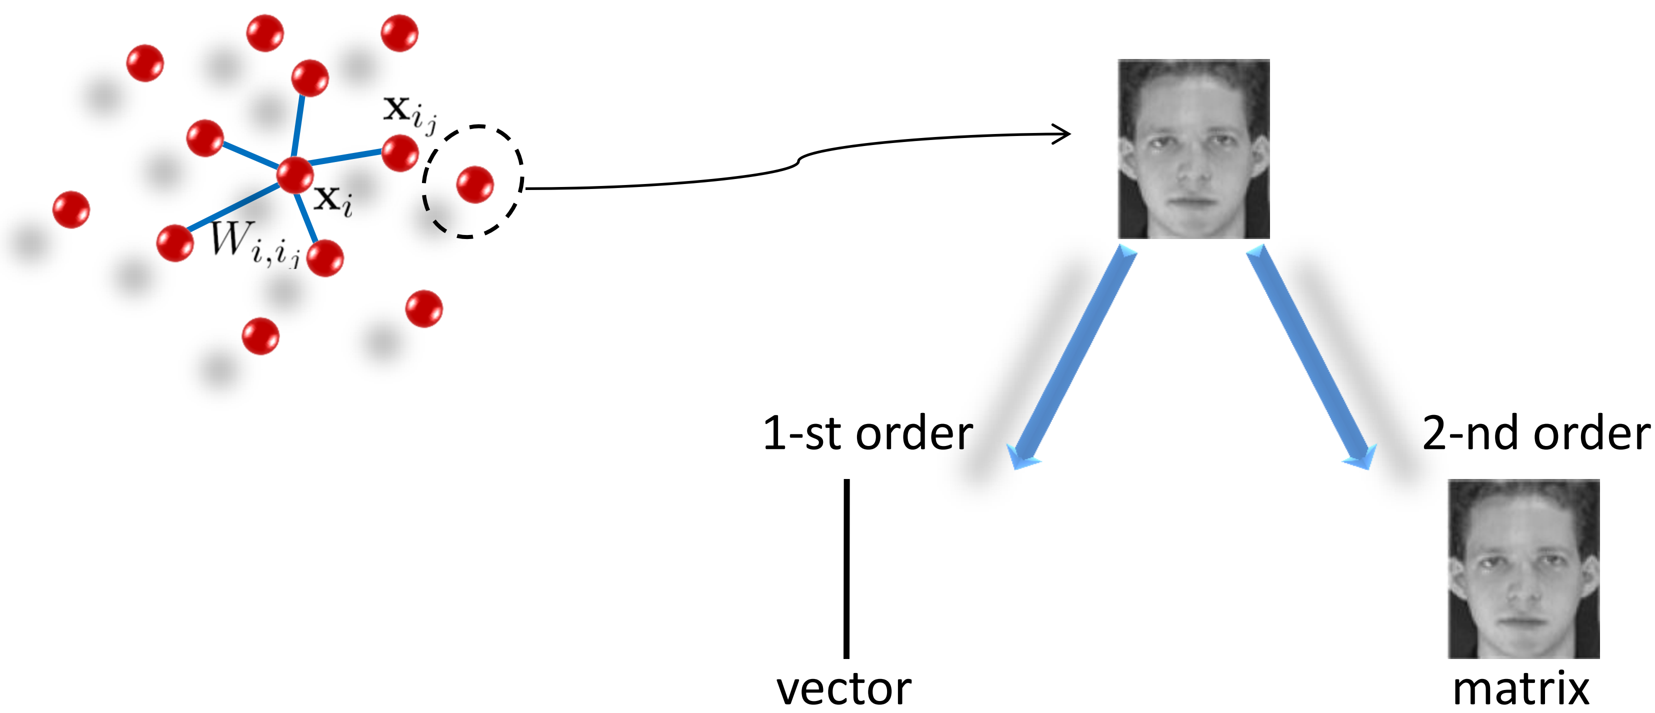

Tensor

(a)

(b)

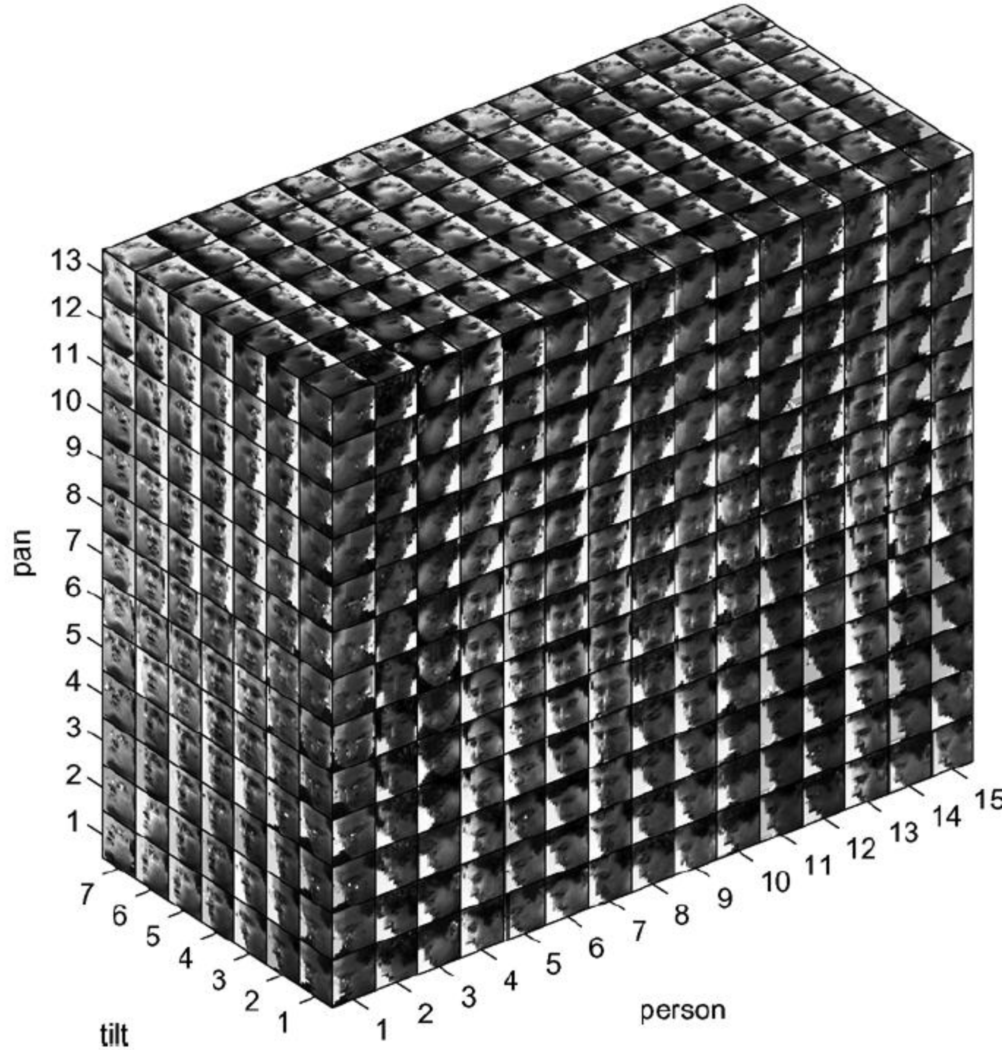

Y. Fu, et. al., IEEE Transactions on Circuits and Systems for Video Technology, 2009.

Correlation Tensor Analysis

Given two m-th order tensors, $\mathbf{X}_1, \mathbf{X}_2 \in \mathbb{R}^{D_1 \times D_2 \times \cdots \times D_m}$

Pearson Correlation Coefficient (PCC):

$$\mathrm{Corr}(X_1, X_2) = \frac{\langle \mathbf{X}_1, \mathbf{X}_2 \rangle}{\sqrt{\langle \mathbf{X}_1, \mathbf{X}_1 \rangle}\sqrt{\langle \mathbf{X}_2, \mathbf{X}_2 \rangle}},$$

$$\langle \mathbf{X}_1, \mathbf{X}_2 \rangle = \sum_{i_1=1, i_2=1, \ldots, i_m=1}^{D_1, D_2, \ldots, D_m} \mathbf{X}_{1; i_1, i_2, \ldots, i_m} \mathbf{X}_{2; i_1, i_2},$$

CTA objective function

$$\arg\max_{\mathbf{U}} \ \varepsilon(\mathbf{U}) = \sum_{i,j=1}^{n} \underbrace{\left(1 - \frac{\langle \mathbf{Y}_i, \mathbf{Y}_j \rangle}{\sqrt{\langle \mathbf{Y}_i, \mathbf{Y}_i \rangle \langle \mathbf{Y}_j, \mathbf{Y}_j \rangle}}\right)}_{\text{Correlation Distance and Multilinear Representation}} \cdot \underbrace{\left(w_{ij}^{(\mathrm{d})} - w_{ij}^{(\mathrm{s})}\right)}_{\text{Fisher Graph}}$$

$$\mathbf{Y}_i = \mathbf{X}_i \times_1 \underbrace{\mathbf{U}_1 \times_2 \ldots \times_m \mathbf{U}_m}_{\text{m different subspaces}}, \ \ \mathbf{U}_j \in \mathbb{R}^{D_j \times d_j} \ \text{for} \ j = 1, \ldots, m.$$

Y. Fu, et. al., IEEE Transactions on Image Processing, 2008.

Large Scale Manifold Learning

- Graph based methods require spectral decomposition of matrices of $n \times n$, where n denotes the number of samples.
- The storage cost and computational cost of building neighborhood maps are $O(n^2)$ and $O(n^3)$, it is almost intractable to apply these methods to large-scale scenarios.
- Neighborhood search is also a large scale aspect.

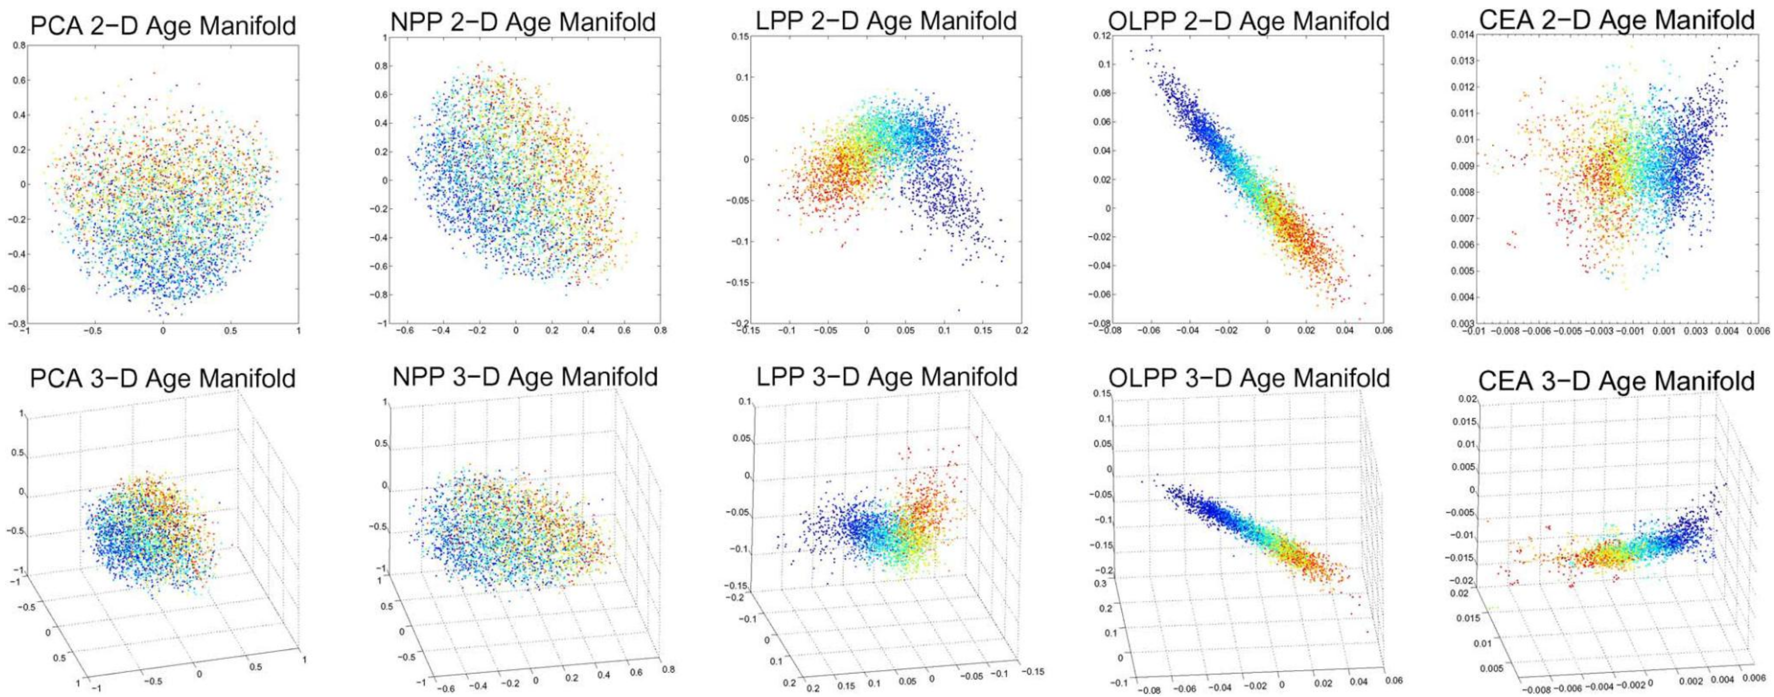

Large Scale Manifold Learning

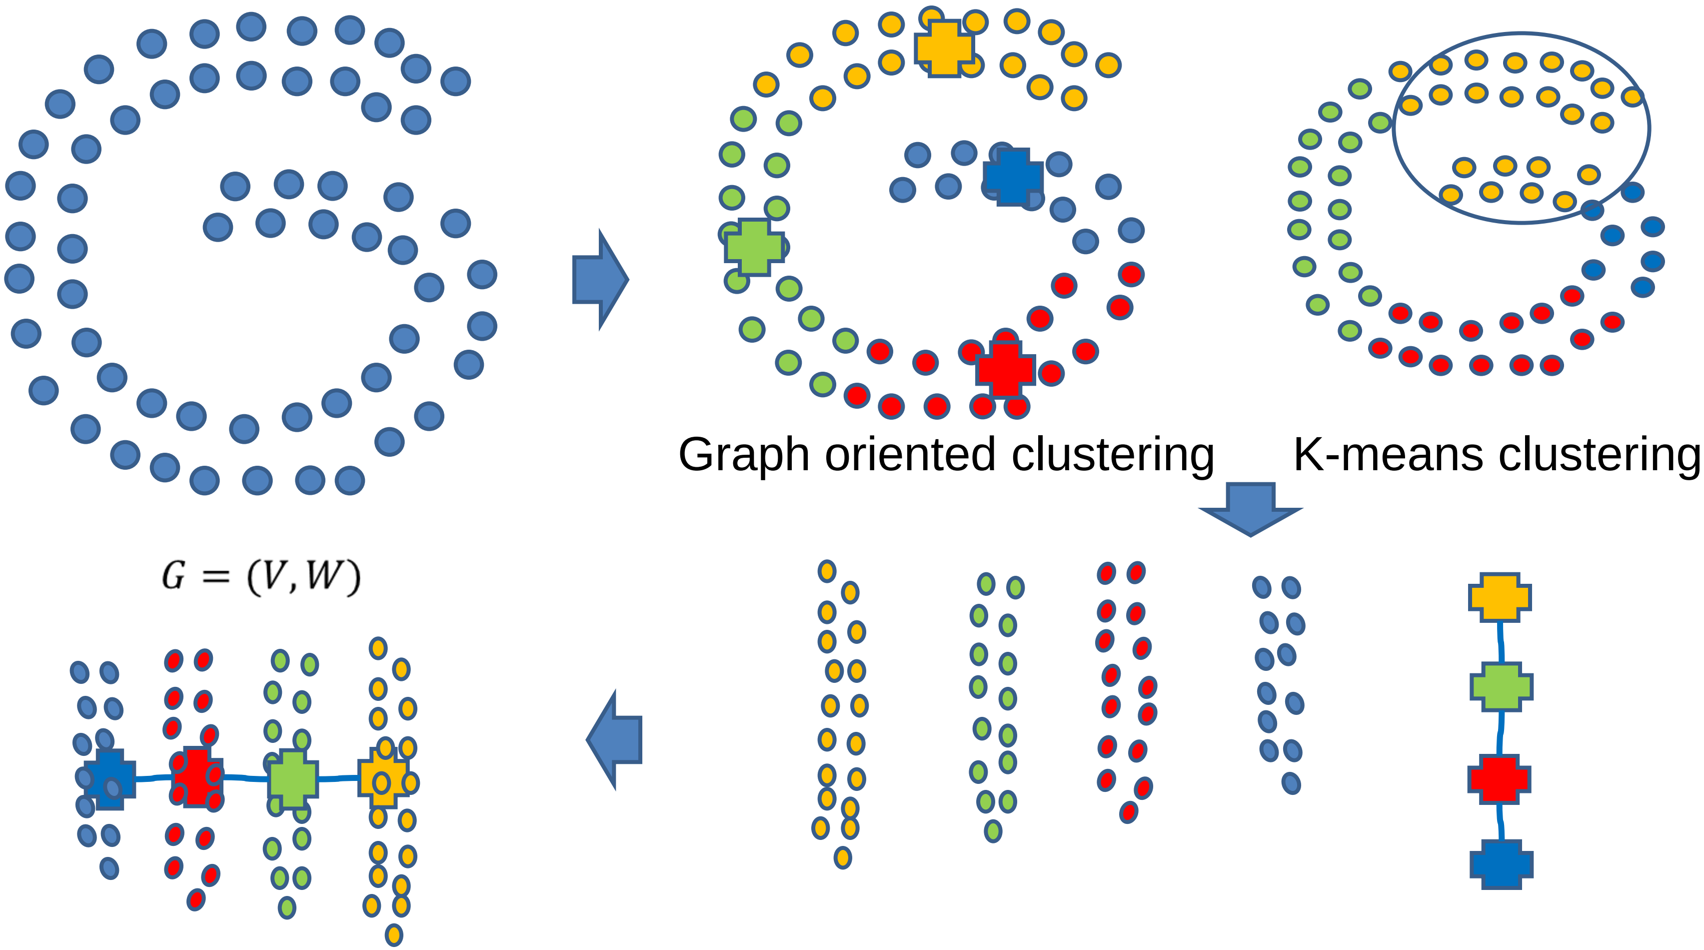

$$[\widehat{W}_{Inter} \backslash \widehat{W}_{Intra}]$$
$$= LowRankApproximation(W_{Inter} \backslash W_{Intra})$$
$$s.t. \left\| [\widehat{W}_{Inter} \backslash \widehat{W}_{Intra} \right\|_F \leq \varepsilon$$

$$G = (V, \widehat{W} = \widehat{W}_{Inter} + \widehat{W}_{Intra})$$

$$W_{Intra} = diag(W_1, \cdots, W_k)$$

$$W_{Inter} = \begin{bmatrix} 0 & W_{12} & \cdots & W_{1k} \\ W_{21} & 0 & \cdots & \vdots \\ \vdots & \vdots & \ddots & W_{1k-1} \\ W_{k1} & \cdots & W_{k-11} & 0 \end{bmatrix}$$

Previous and Current Work:

Social Media Scenario

Expression Manifold

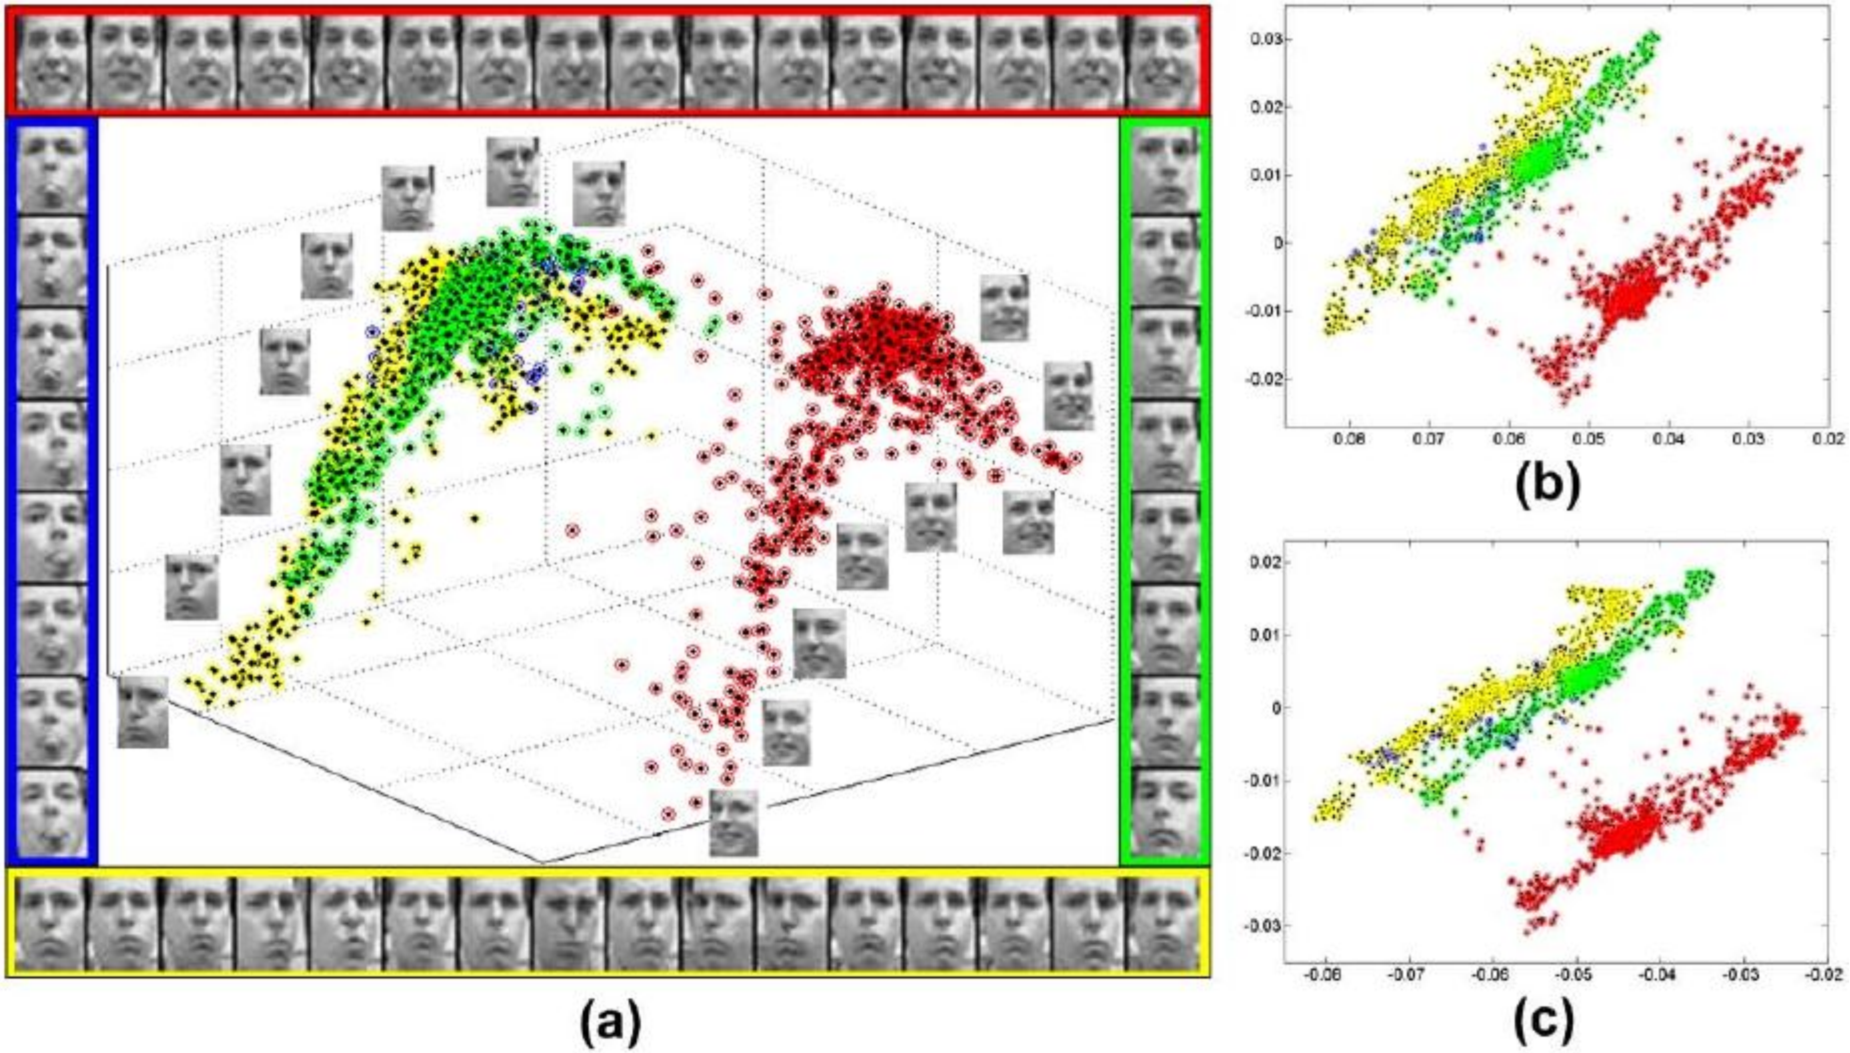

Manifold visualization of 1,965 Frey's face images by LEA using *k = 6 nearest* neighbors.

Yun Fu, et. al. "Locally Adaptive Subspace and Similarity Metric Learning for Visual Clustering and Retrieval", CVIU, Vol. 110, No. 3, pp: 390-402, 2008.

Emotion State Manifold

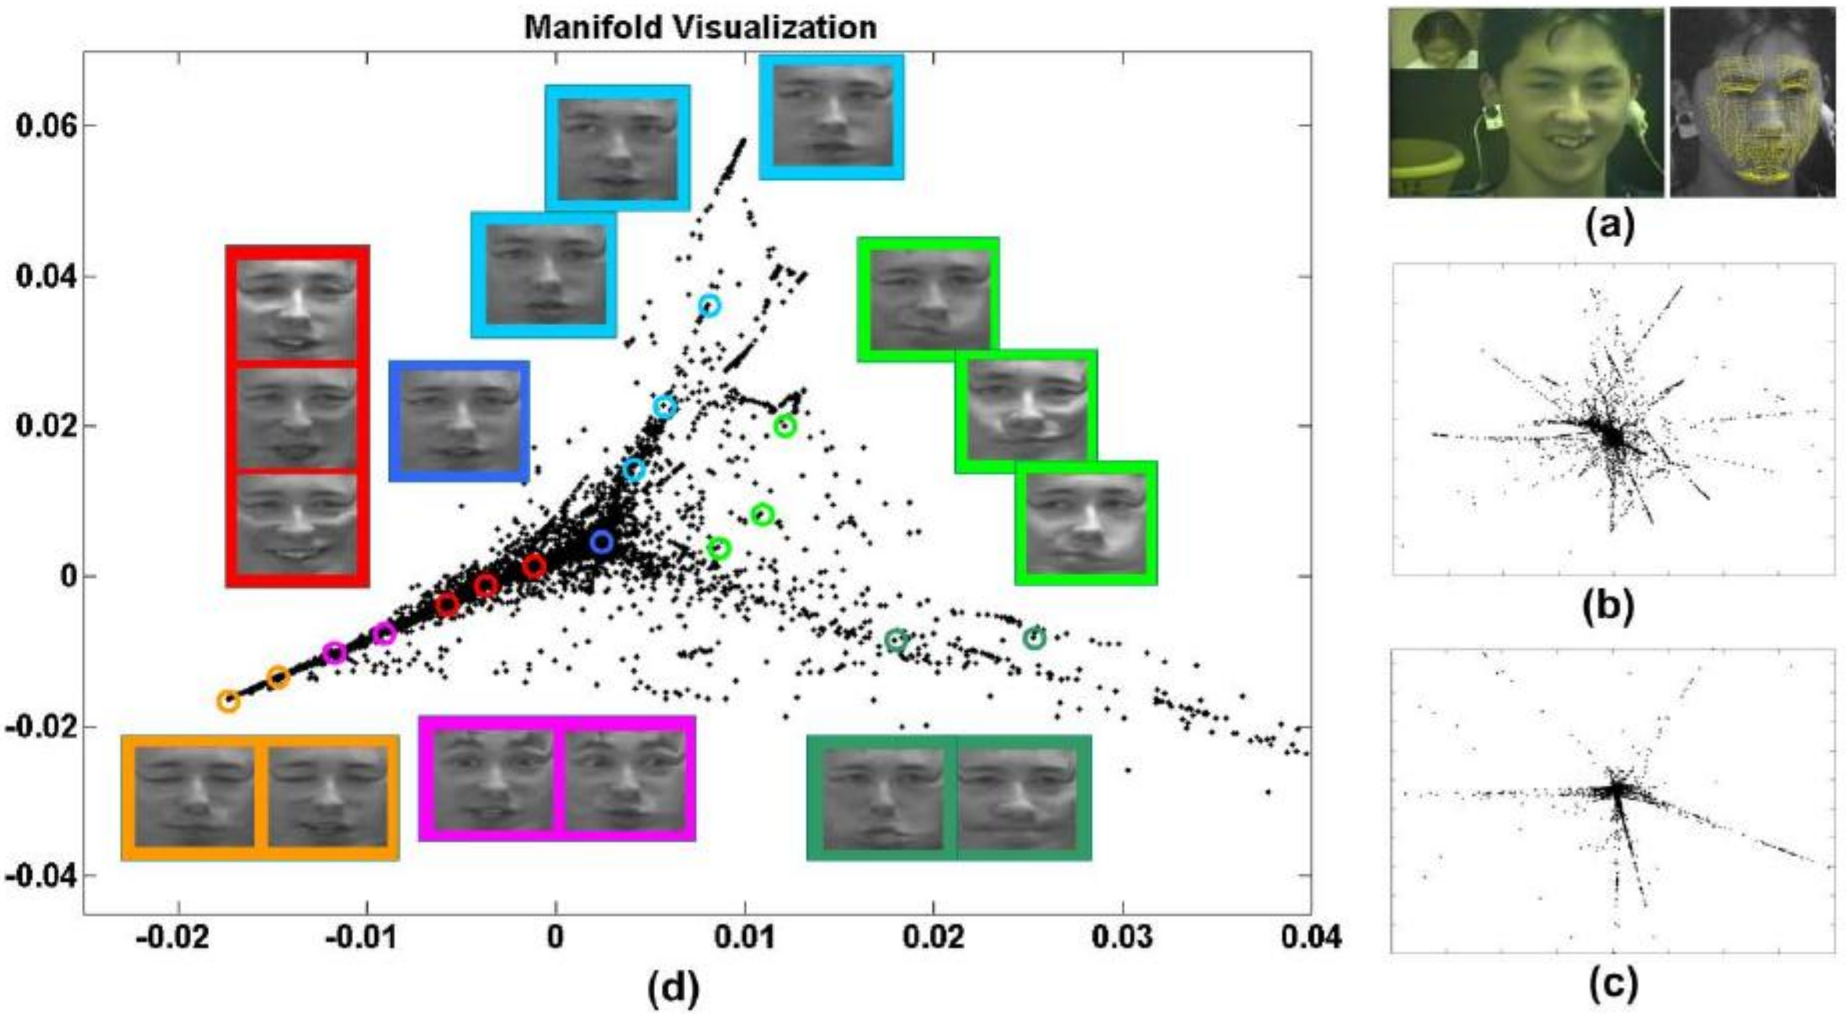

Manifold visualization for 11,627 AAI sequence images of a male subject using LLE algorithm. (a) A video frame snapshot and the 3D face tracking result. The yellow mesh visualizes the geometric motion of the face. (b) Manifold visualization with k=5 nearest neighbors. (c) k=8 nearest neighbors. (d) k=15 nearest neighbors and labeling results.

Application for Age Estimation

AS International, *How Old Are You?*, www.asmag.com Vol. 120, Page 40-41, Dec. 2008.
PhysOrg.com, *Intelligent Computers See Your Human Traits*, May 2008.
Roland Piquepaille's Technology Trends, *Computers can now guess our age*, Sep. 2008.
UIUC News Bureau, *Step right up, let the computer look at your face and tell you your age*, Sep. 2008.
ABC Science, *Age recognition software has a human eye*, Oct. 2008.
UPI.com, *Age estimation software is created*, Sep. 2008.
Eureka! Science News, *Step right up, let the computer look at your face and tell you your age*, 2008
Zdnet.com, *Computers can now guess our age*, Sep. 2008.
Webindia123.com, *Age estimation software is created*, Sep. 2008.
Newkerala.com, *Now, a computer software that can tell age just by looking at your face!*, 2008.
Hindustantimes.com, *Computer that says how old you are*, Sep. 2008.
TXonline.net, *Age estimation software is created*, Sep. 2008.
Topnews.in, *Now, computer software that can tell age just by looking at your face*, Oct. 2008.


32

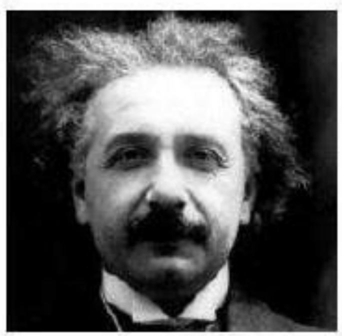
52

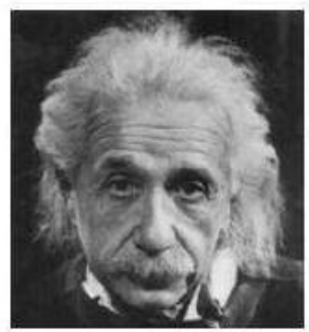
76

Age estimation on Einstein's faces. The estimated ages below each face might be a little bit older than the true ages (unknown to us) but reasonable. Our training data are all Asian faces.

This might be a good example to echo the phenomenon that Asian faces often aesthetically look younger than the Western.

Y. Fu, et. al., IEEE TPAMI, CVPR, ICCV, 2009, 2010, 2011.

Why Regression on Manifold?

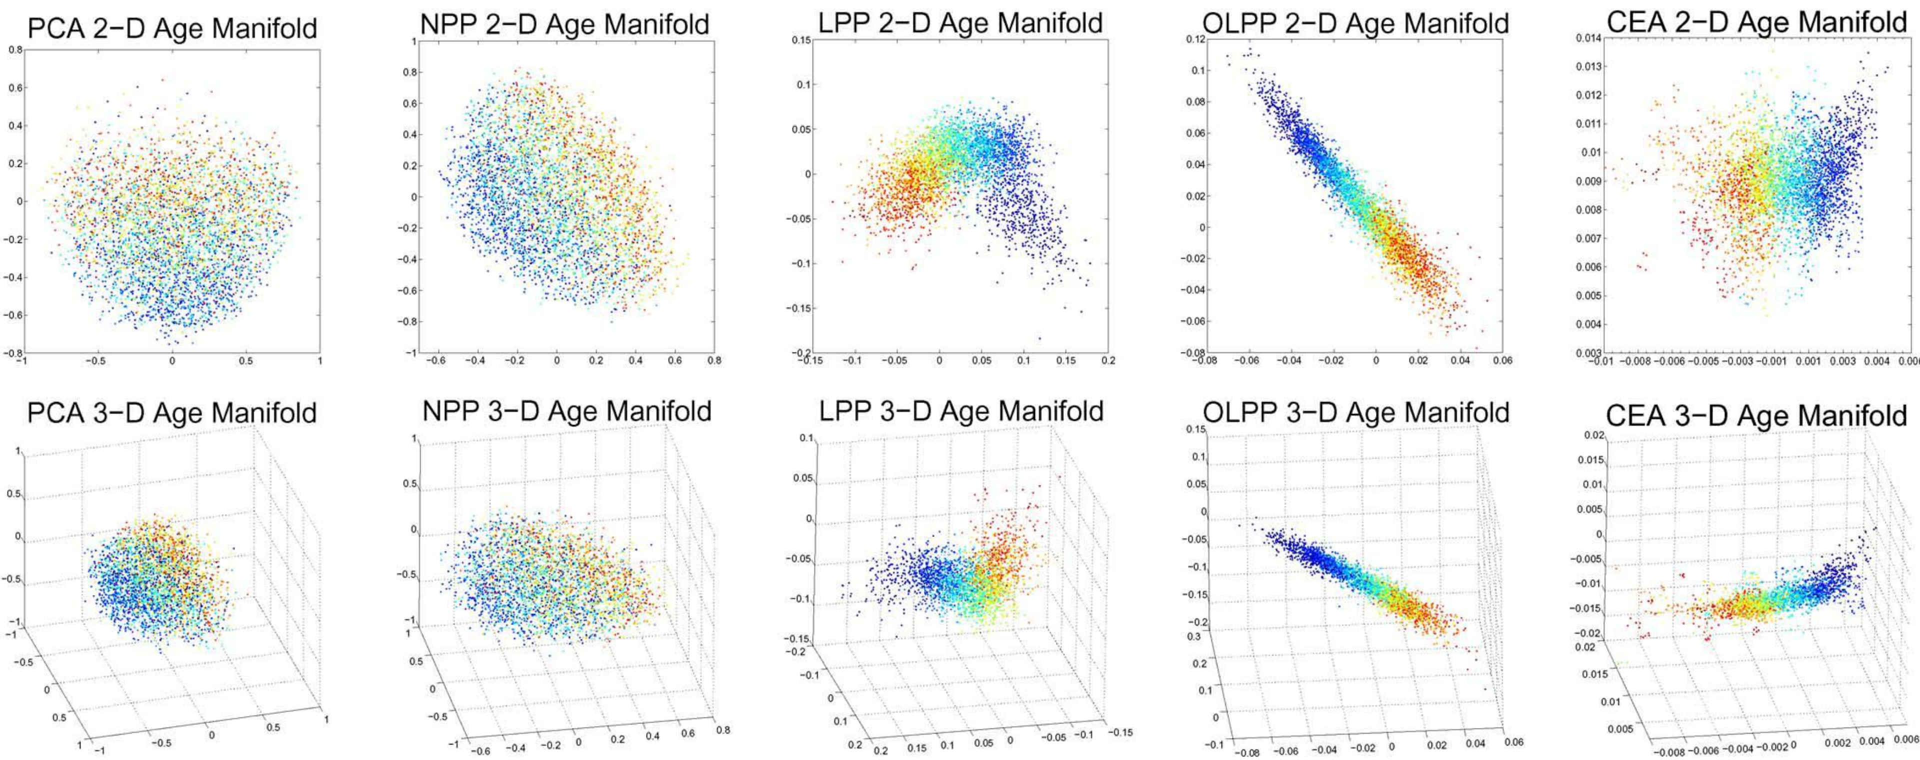

- YGA database
  - 1600 Asian subjects
  - Age range from 0 to 93 years
  - 60x60 gray-level patches
  - 8000 images in total. 4000 female and 4000 male

Y. Fu, et. al., IEEE Transactions on Multimedia, 2008.

Regression Framework

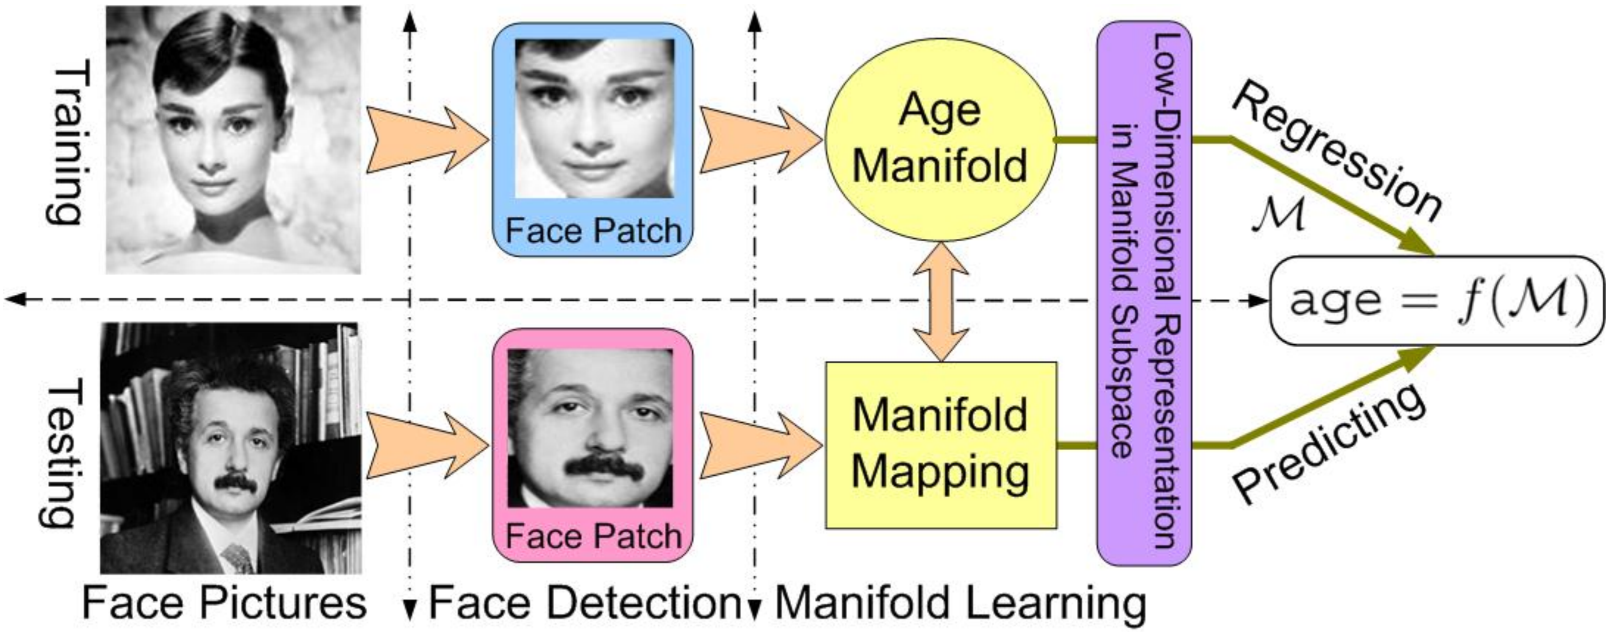

Multiple linear regression $\mathbf{L} = \widetilde{\mathbf{Y}}\mathbf{B} + \mathbf{e}, \quad \mathrm{Var}(\mathbf{e}) = \sigma^2 \mathbf{I},$

Model fitting $\hat{\mathbf{L}} = \widetilde{\mathbf{Y}}\hat{\mathbf{B}}.$ Ordinary Least Squares $\hat{\mathbf{B}} = (\widetilde{\mathbf{Y}}'\widetilde{\mathbf{Y}})^{-1}\widetilde{\mathbf{Y}}'\mathbf{L} = \mathbf{H}\mathbf{L}$

Residuals $\hat{\mathbf{e}} = \mathbf{L} - \hat{\mathbf{L}}.$ $\mathrm{E}(\hat{\mathbf{e}}) = 0$ $\mathrm{Var}(\hat{\mathbf{e}}) = \sigma^2(\mathbf{I} - \mathbf{H}).$

Quadratic function $\hat{l}_i = \hat{b}_0 + \hat{\mathbf{b}}_1^T \mathbf{y}_i + \hat{\mathbf{b}}_2^T \mathbf{y}_i^2,$

$$\hat{\mathbf{L}} = [\hat{l}_1 \cdots \hat{l}_n]^T, \quad \hat{\mathbf{B}} = [\hat{b}_0 \;\; \hat{b}_1^{(1)} \cdots \hat{b}_1^{(d)} \;\; \hat{b}_2^{(1)} \cdots \hat{b}_2^{(d)}]^T$$

$$\widetilde{\mathbf{Y}} = [\mathbf{1}_{n\times 1} \;\; [\mathbf{y}_1 \cdots \mathbf{y}_n]^T \;\; [\mathbf{y}_1^2 \cdots \mathbf{y}_n^2]^T].$$

Y. Fu, et. al., IEEE Transactions on Multimedia, 2008.

CEA for Age Estimation

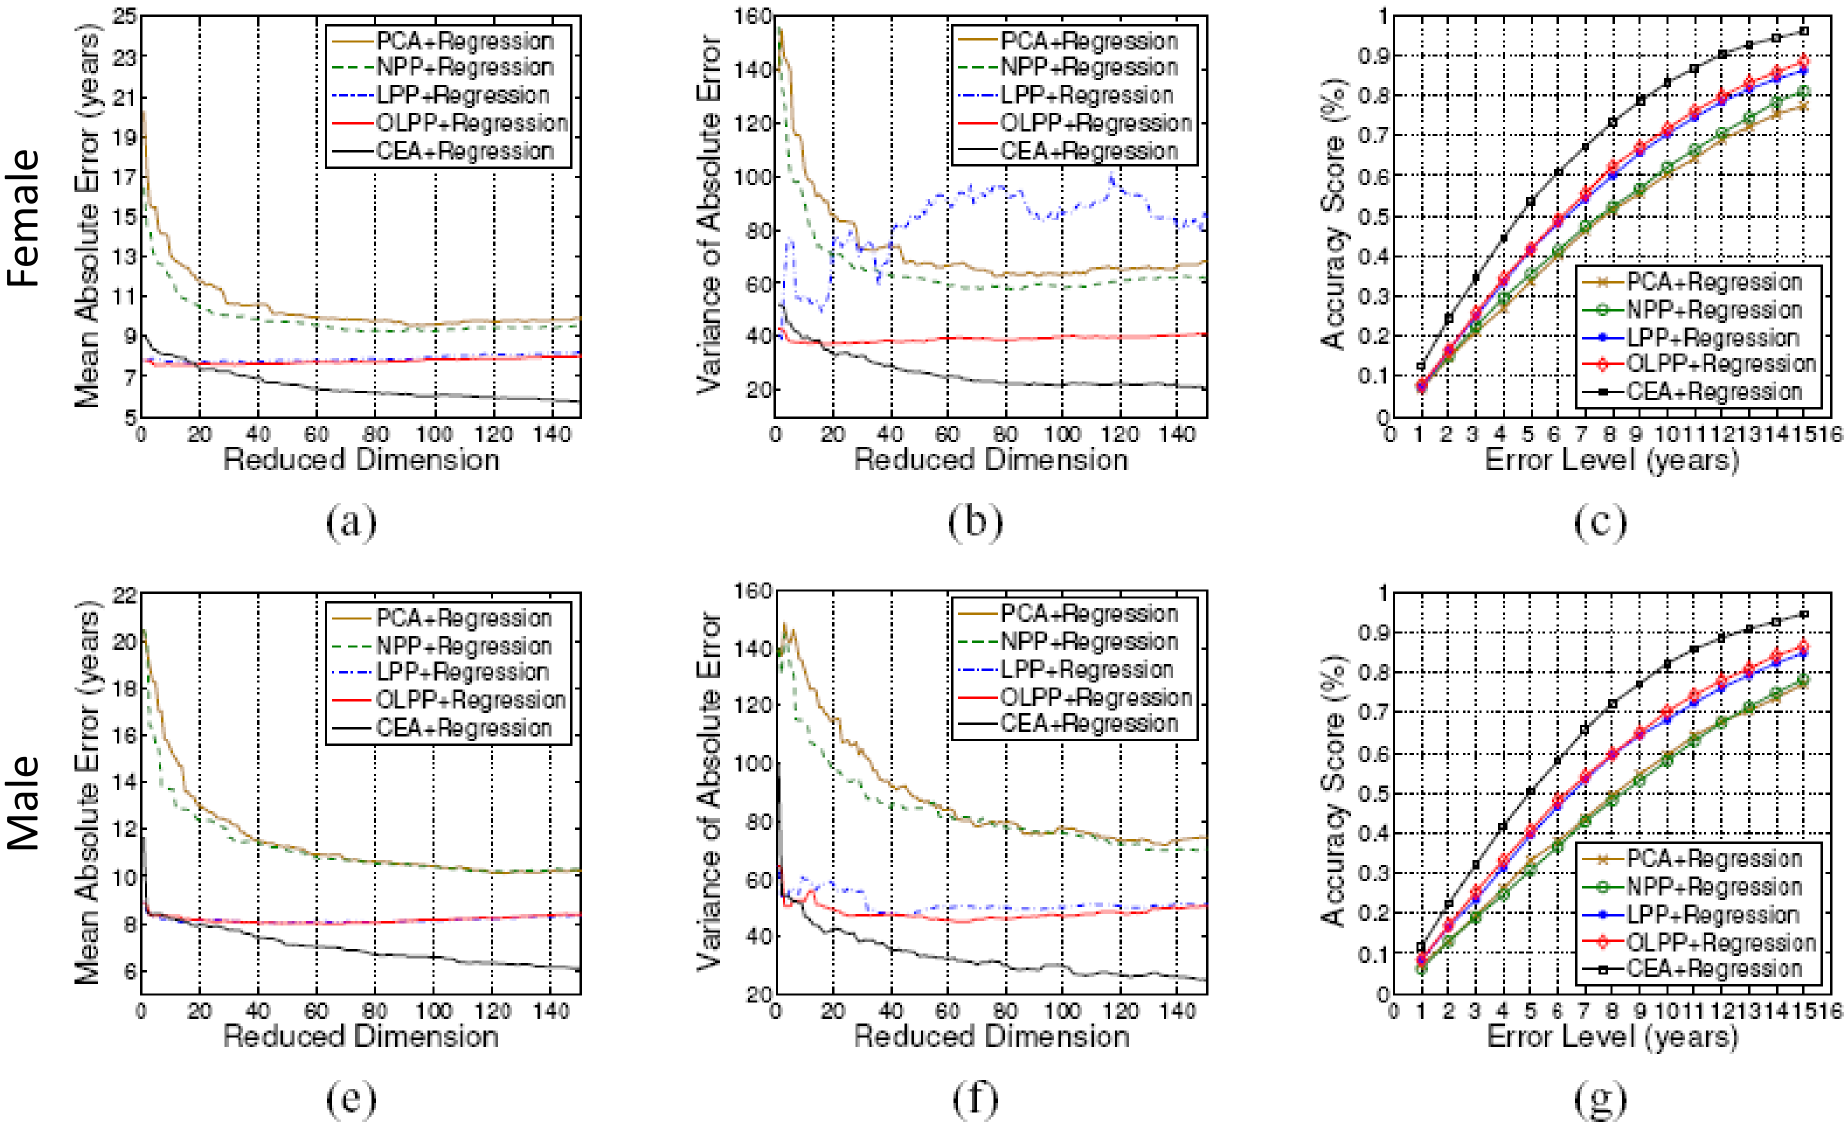

Y. Fu, et. al., IEEE Transactions on Multimedia, 2008.

Automatic Age Estimation

Representation for Age and Gender Group Classification	Representations for Age Estimation in Each Group		Age Estimation Error (MAE)
BIF+LSDA	Young Female	BIF+MFA	**2.66**
	Adult Female	BIF+LSDA	
	Senior Female	BIF+MFA	
	Young Male	BIF+MFA	
	Adult Male	BIF+OLPP	
	Senior Male	BIF+LSDA	

Method	MAE	Error Reduction Rate
RPK [33]	4.66	0
F1	2.95	36.7%
F2	2.84	39.1%
F3	2.87	38.4%
F2.F	**2.66**	**42.9%**

MAEs (in years) comparison with the result in [33] that uses manual separation of gender.

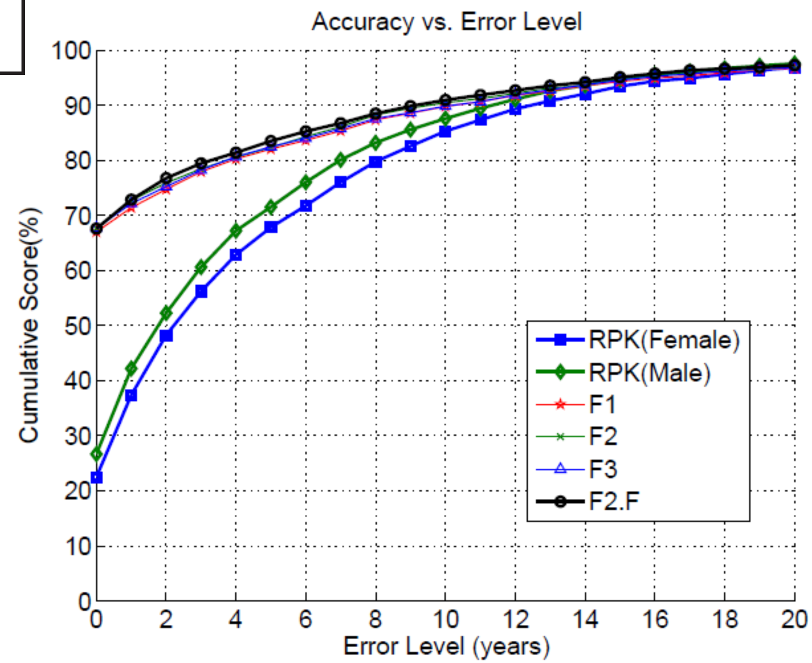

Y. Fu, et. al., IEEE CVPR, ICCV, 2009.

Gender Recognition from Body

Method	Frontal View	Back View	Mixed View
PBGR [7]	76.0±1.2%	74.6±3.4%	75.0±2.9%
BIO+PCA	79.1±3.0%	82.8±5.1%	**79.2±1.4%**
BIO+OLPP	78.3±3.5%	82.8±4.6%	77.1±1.9%
BIO+LSDA	**79.5±2.6%**	**84.0±3.9%**	78.2±1.1%
BIO+MFA	79.1±2.2%	81.7±5.2%	75.2±2.1%

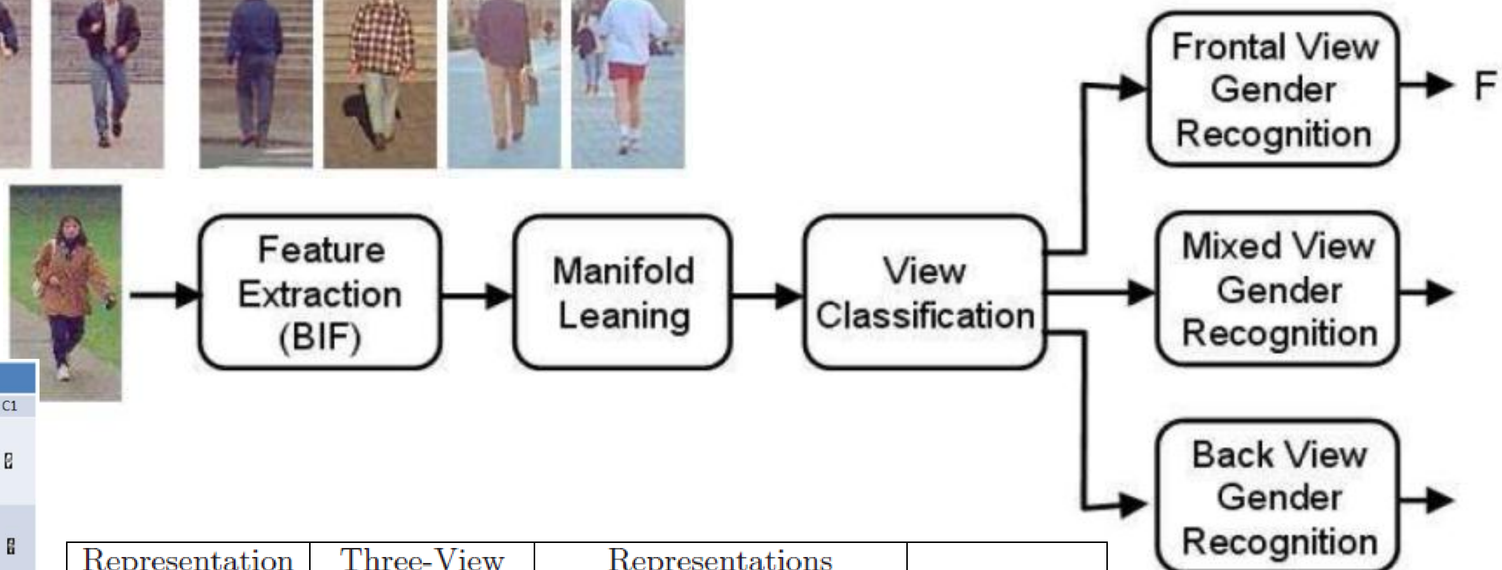

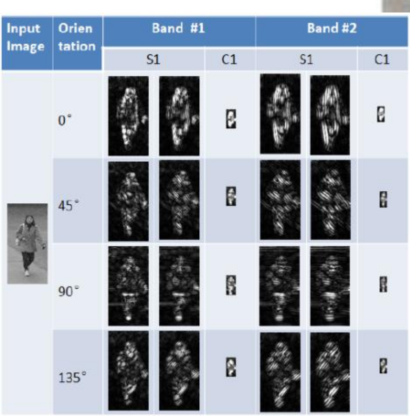

Bio-Inspired Feature

Representation for Human Body View Classification	Three-View Classification Accuracy (Frontal, Back, and Mixed)	Representations for Gender Recognition in Each View		Gender Recognition Accuracy
BIF+PCA	**99.3±0.5%**	Frontal View	BIF+LSDA	**80.6±1.2%**
		Mixed View	BIF+PCA	
		Back View	BIF+LSDA	

Y. Fu, et. al., ACCV, 2009.

Kinship Recognition

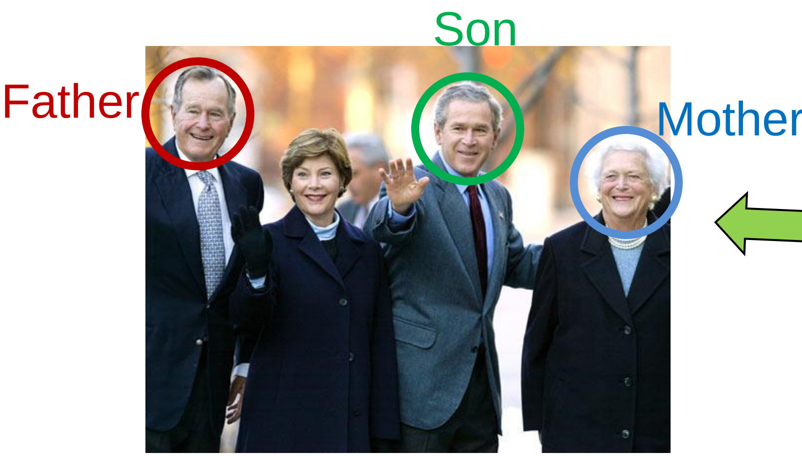

Family Album

KinFace Database

Son

Young Father

Father

- Hypothesis: most of children look like their parents at young ages
- Utilizing transfer learning method to bridge the gap

UB KinFace Database

(a) Old parents

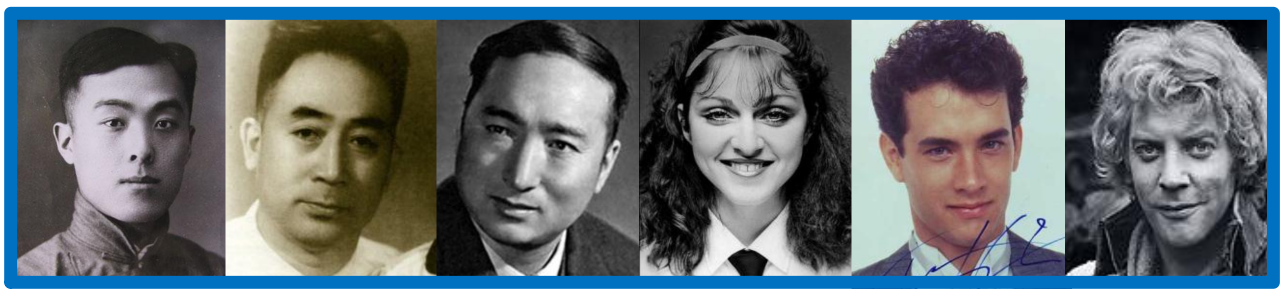

(b) Young parents

(c) Childern

KinFace is the first database which contains 600 images of both children and parents at different ages.

All the images are real-world images of public figures downloaded from the Internet.

UB KinFace Database

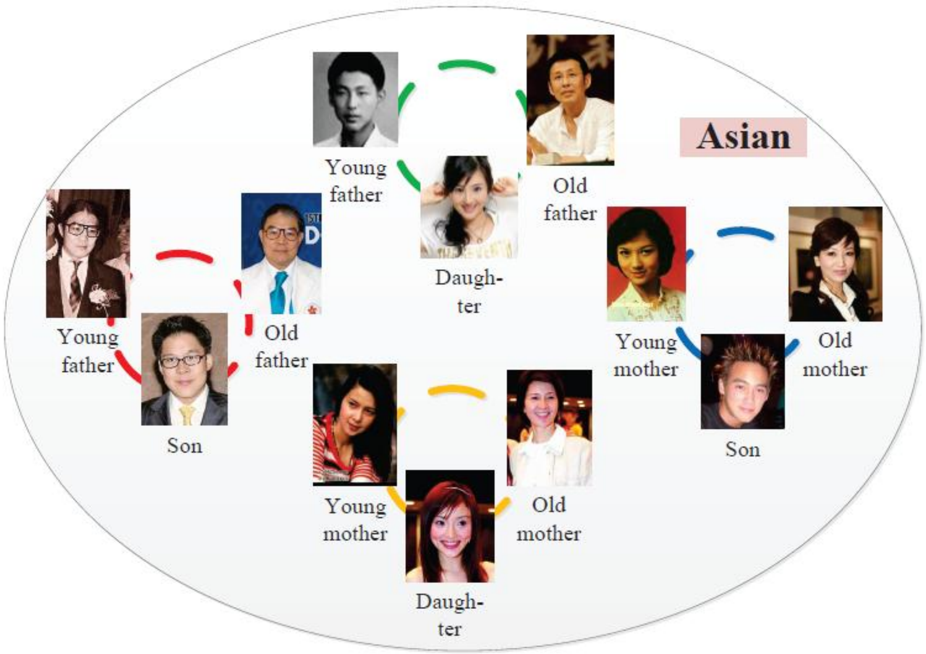

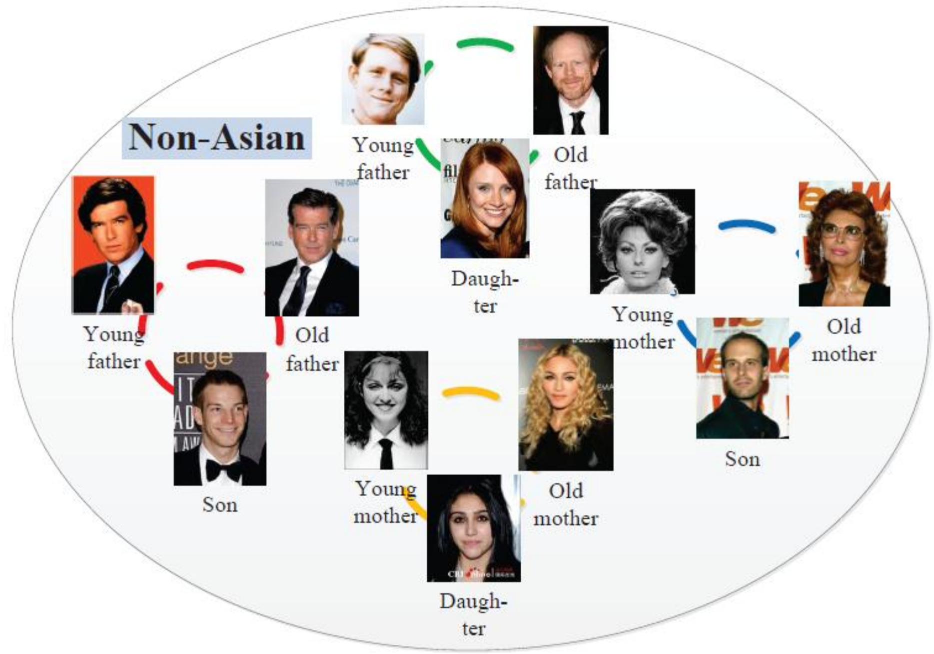

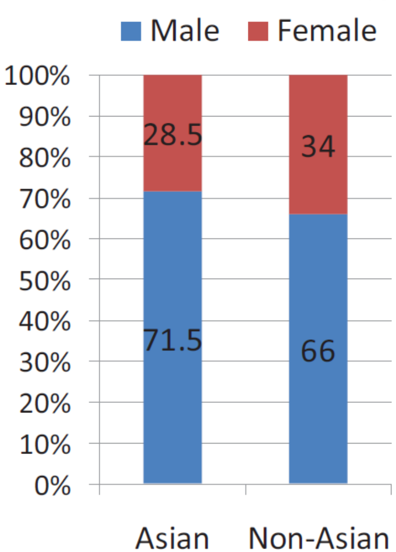

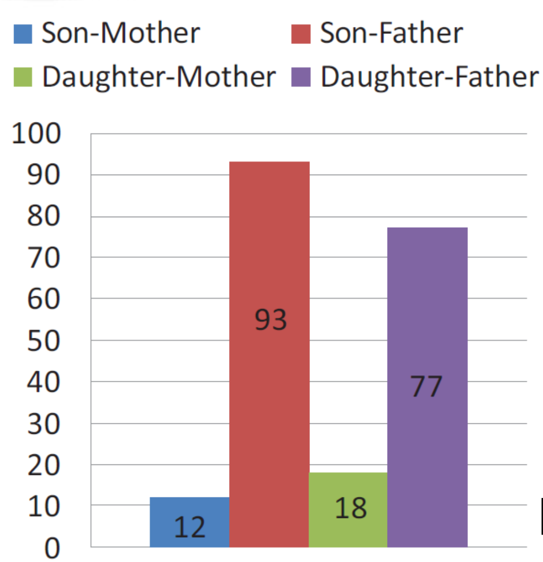

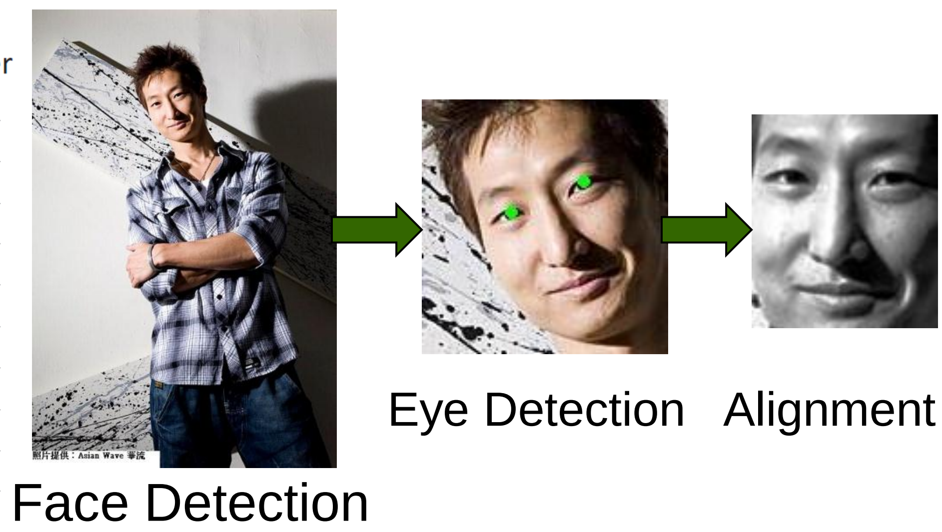

Transfer Subspace Learning

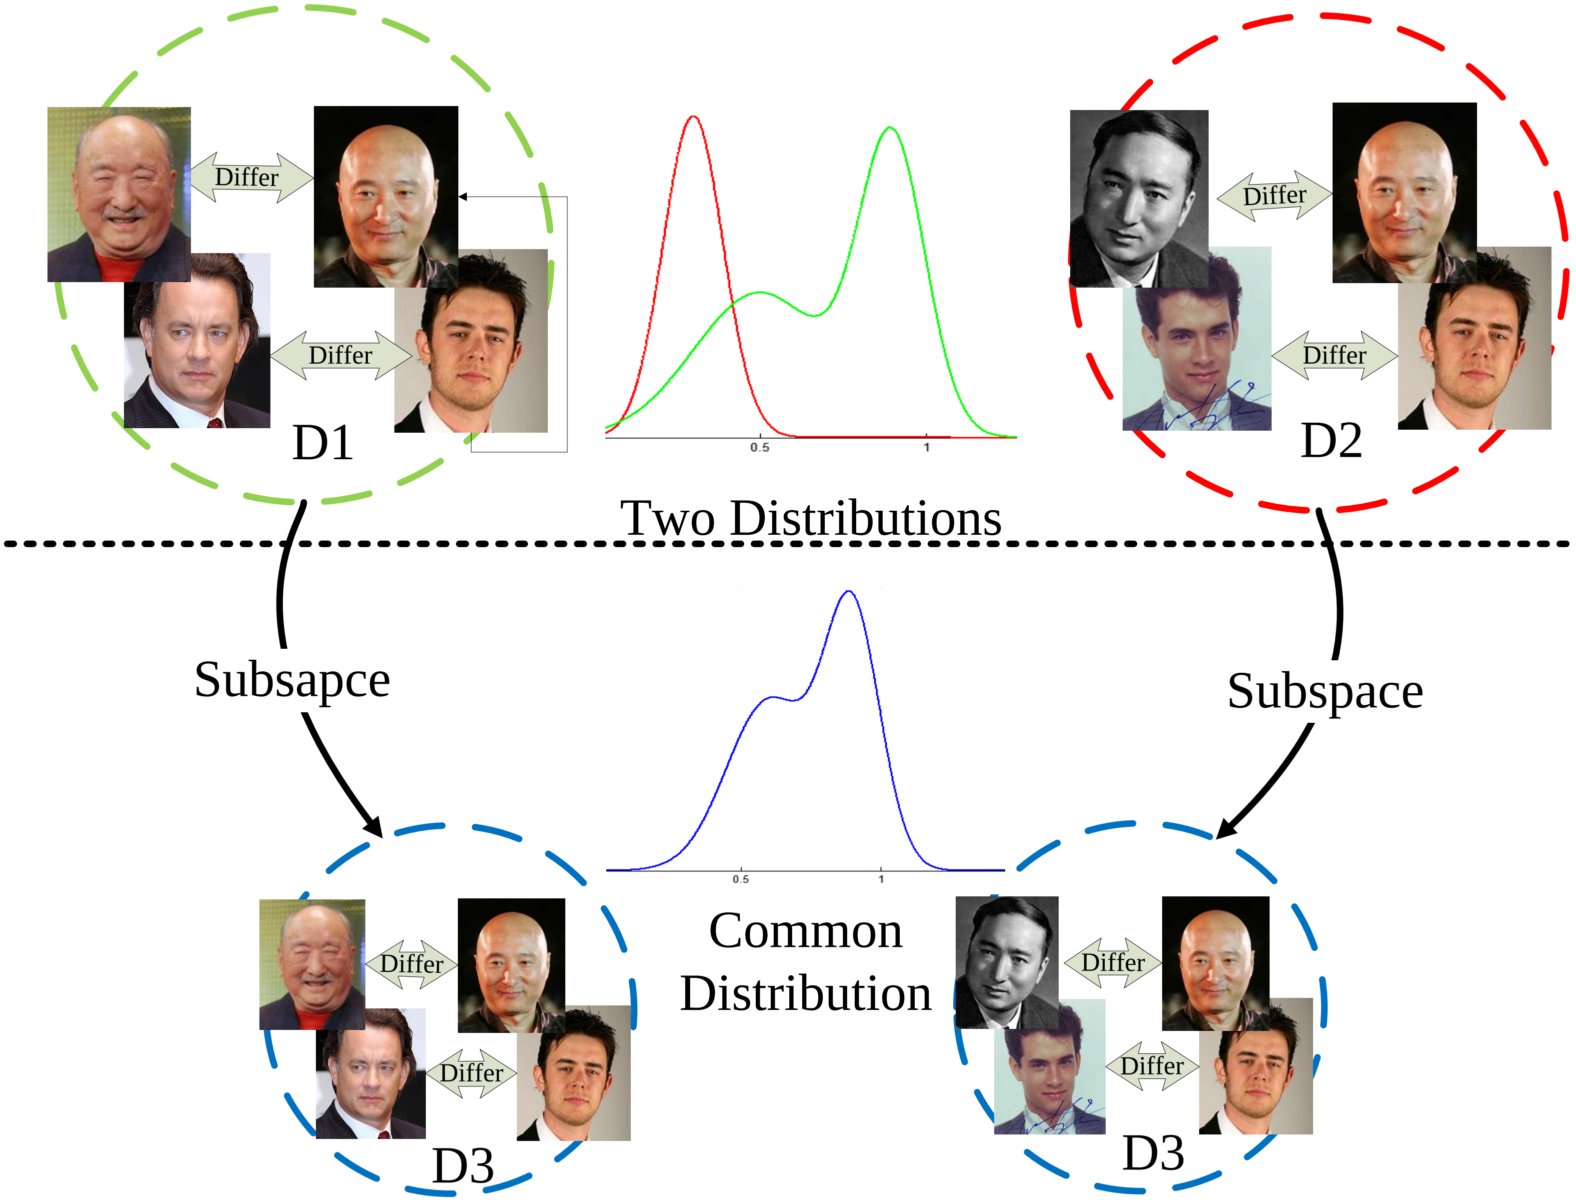

Experiments

1) Kinship Classification

"Young parents are more similar to their children based on the local Gabor features."

Hypothesis is validated!

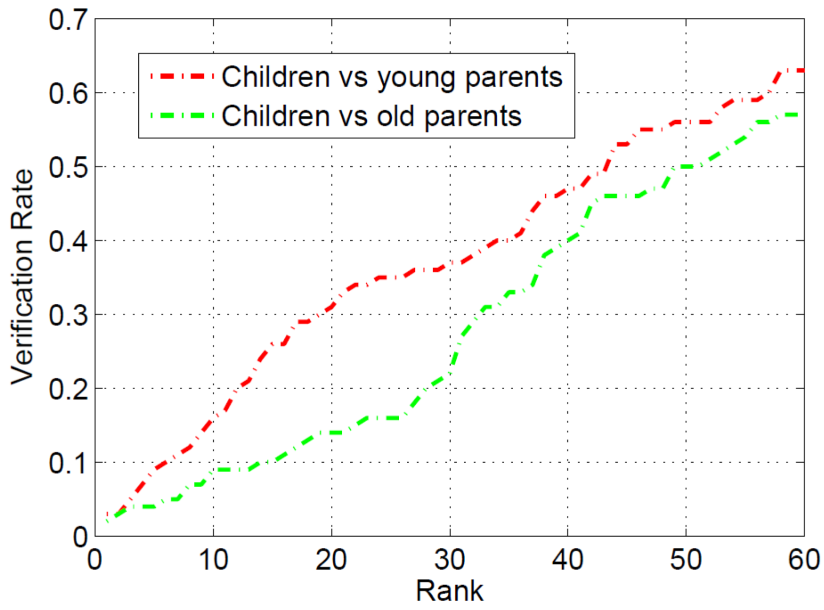

2) Pairwise Kinship Verification

Method	Feature	5-fold	Leave-one-out
Pairwise	Structure	49%	47.25%
TSL	Structure	47.25%	47.25%
Pairwise	Local Gabor	50.25%	41.25%
TSL	Local Gabor	**55.25%**	**63.75%**

Experiments

3) Verification via TSL

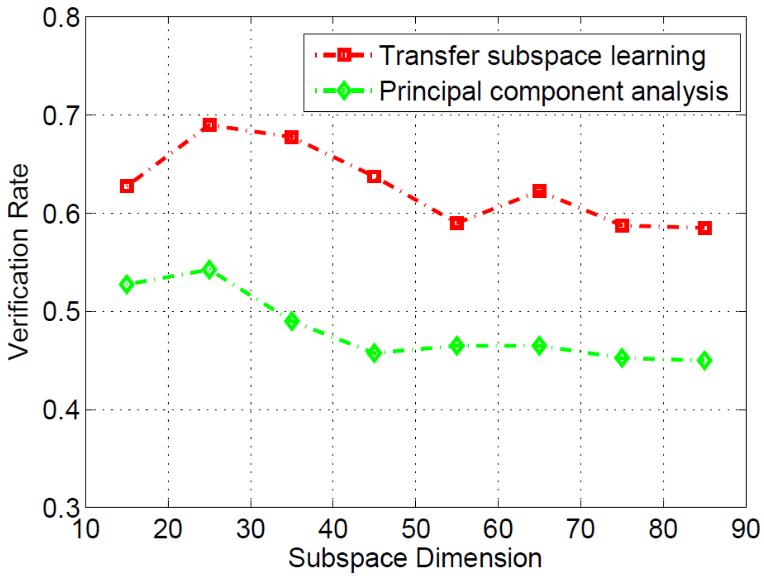

Transfer Subspace Learning

Principal Component Analysis

4) Comparison with Human Performance

Test 1 (learning without young parents' face images):
Result: 52.50%
Test 2 (learning with young parents' face images):
Result: 56.87%

Previous and Current Work:

HCI Scenario

Lipreading by Manifold Learning

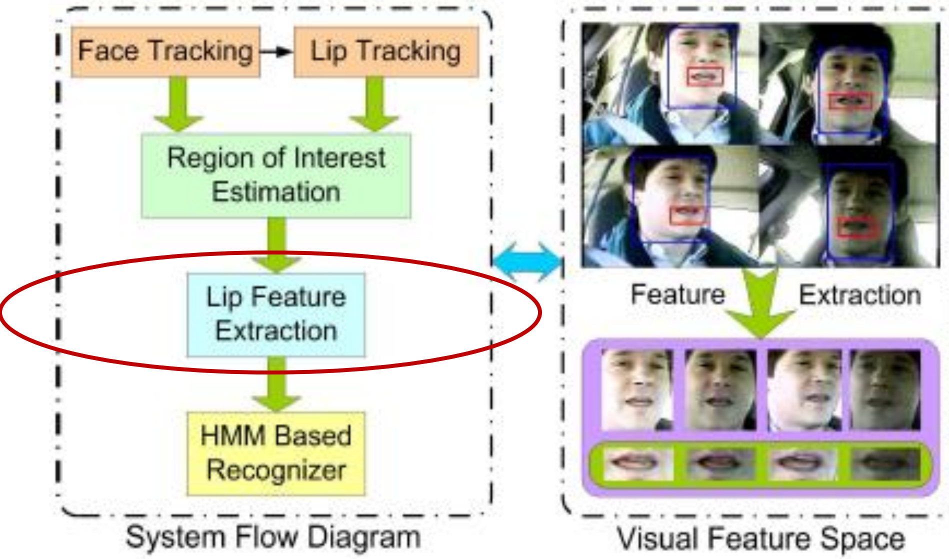

Figure 6. Lipreading system structure.

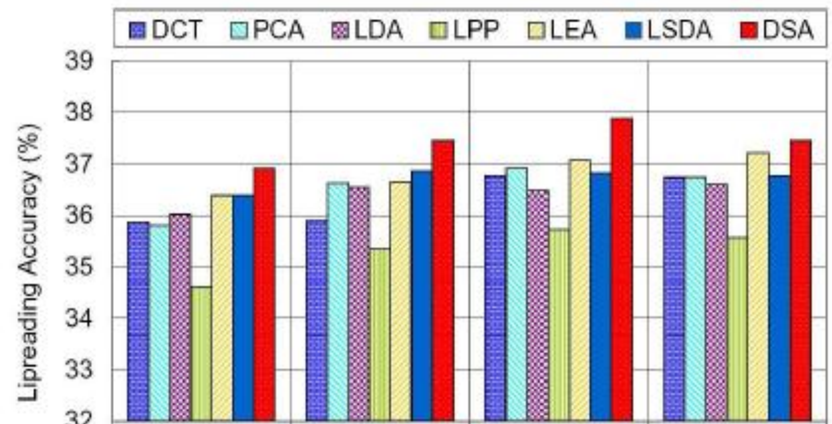

	2	4	8	32
DCT	35.87 (49)	35.90 (37)	36.76 (49)	36.73 (44)
PCA	35.80 (45)	36.63 (42)	36.91 (47)	36.73 (41)
LDA	36.02 (48)	36.55 (48)	36.48 (49)	36.60 (49)
LPP	34.61 (49)	35.34 (41)	35.72 (48)	35.57 (38)
LEA	36.38 (42)	36.65 (40)	37.08 (44)	37.21 (33)
LSDA	36.38 (48)	36.86 (49)	36.81 (39)	36.76 (47)
DSA	36.91 (41)	37.41 (47)	37.87 (41)	37.46 (50)

Number of Gaussian Mixtures

Figure 7. Comparison of the highest lipreading accuracy and dimensionality reduction of all the 7 methods under 4 different numbers of Gaussian mixtures per HMM state.

- Training: 21 talkers
- Test: 13 *different* talkers
- Dim: 584x4DCT→100PCA

The is the highest lipreading accuracy on this database ever reported.

Y. Fu, et. al., IEEE Transactions on Information Forensics and Security, 2008.

Hand Parts Recognition from Depth Image

Reconstruct skeleton model from Kinect

depth image　　classification　　reconstruction

Difficulties

- More flexible than body
- More sensitive to environment change
- Small volume and difficult to detect
- No obviously texture to distinguish

Training data collection

Glove Aided Labeling

Labeling process

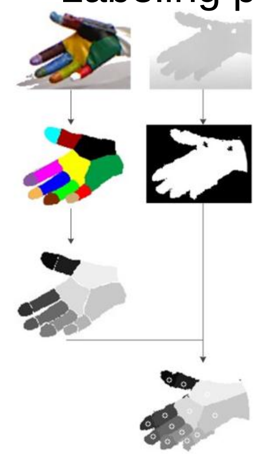

Manually labeling

Denoising Clustering

Labeled data

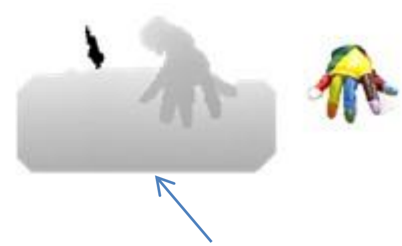

Background can be changed to achieve more training data

Classification and reconstruction

- Using Microsoft's framework as a baseline [Jamie etc., CVPR 11]
- Combine simple features

In processing…

- add more manually labeled&synthetic data
- design robust features

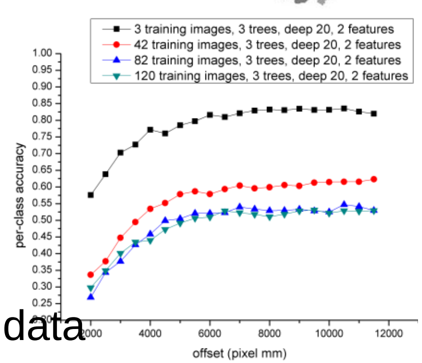

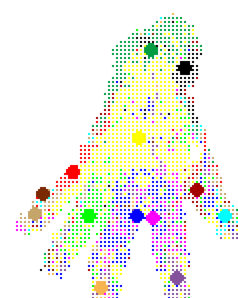

Training and test with combined features

3D Humanoid Avatar

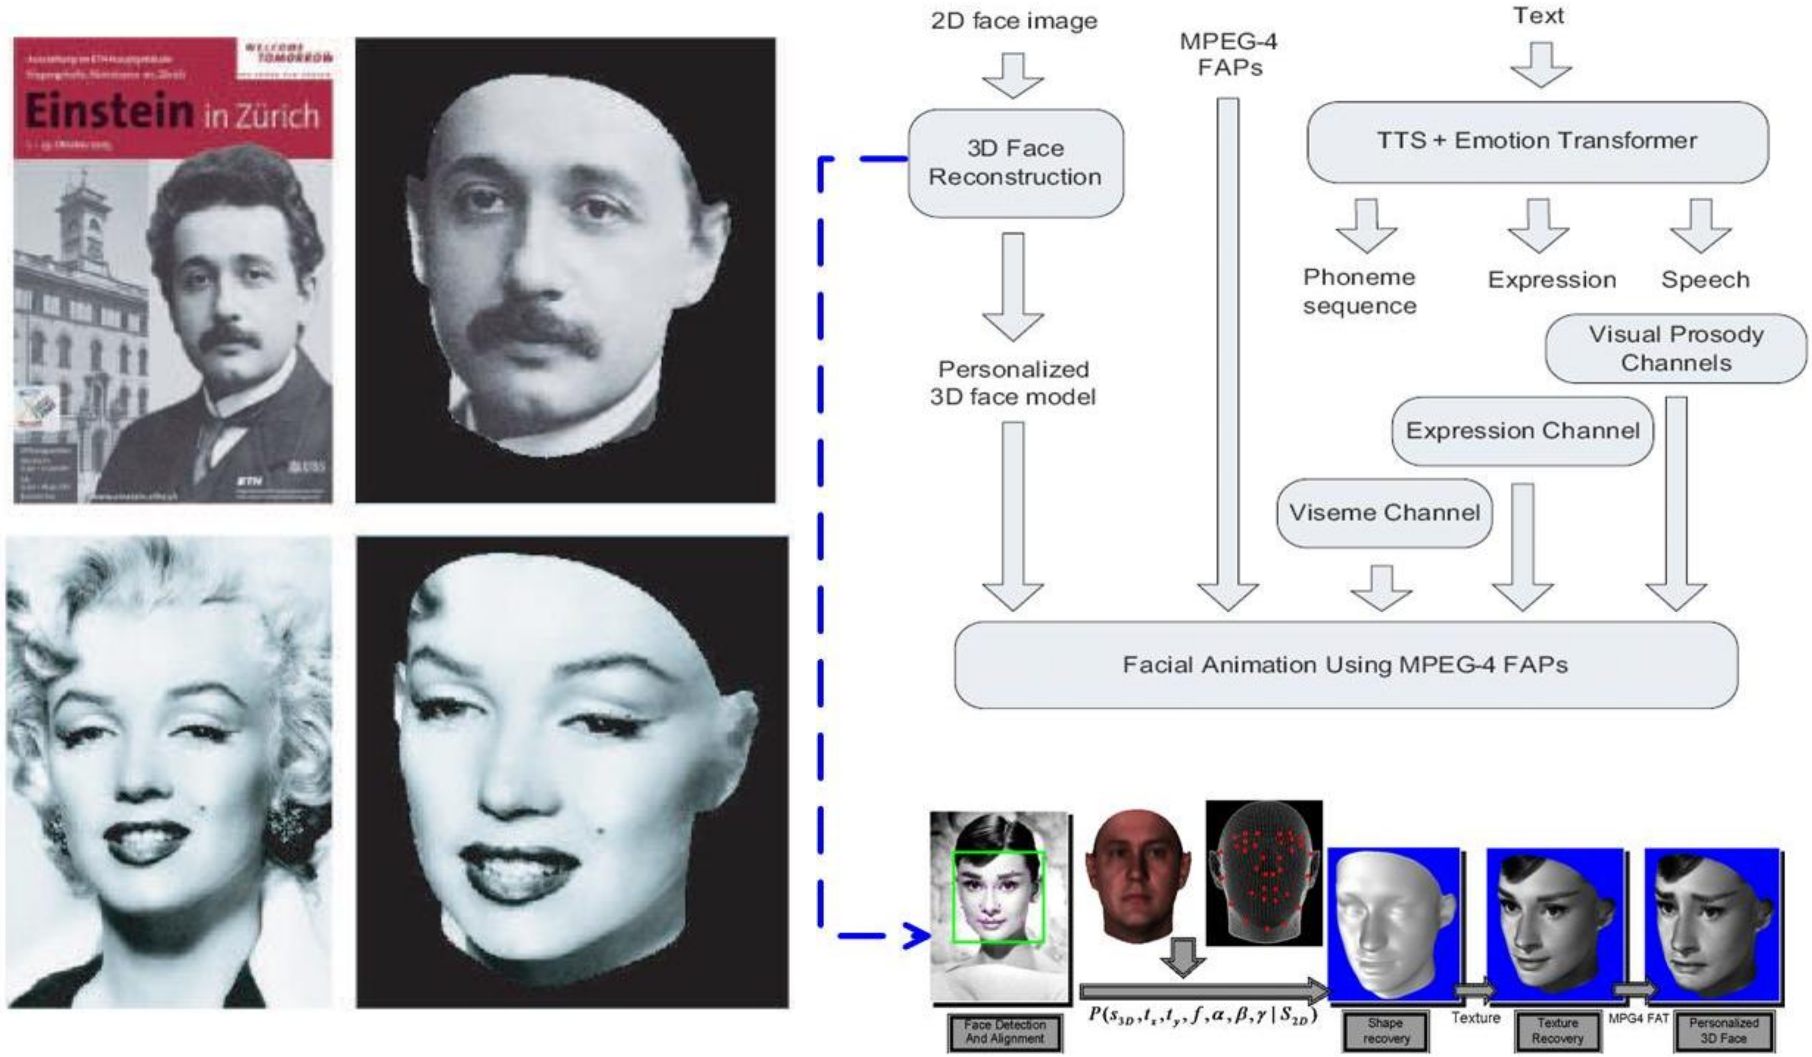

Y. Fu, et. al., IEEE Transactions on Multimedia (T-MM) , 2008.

Multimodal Human-Avatar Interaction

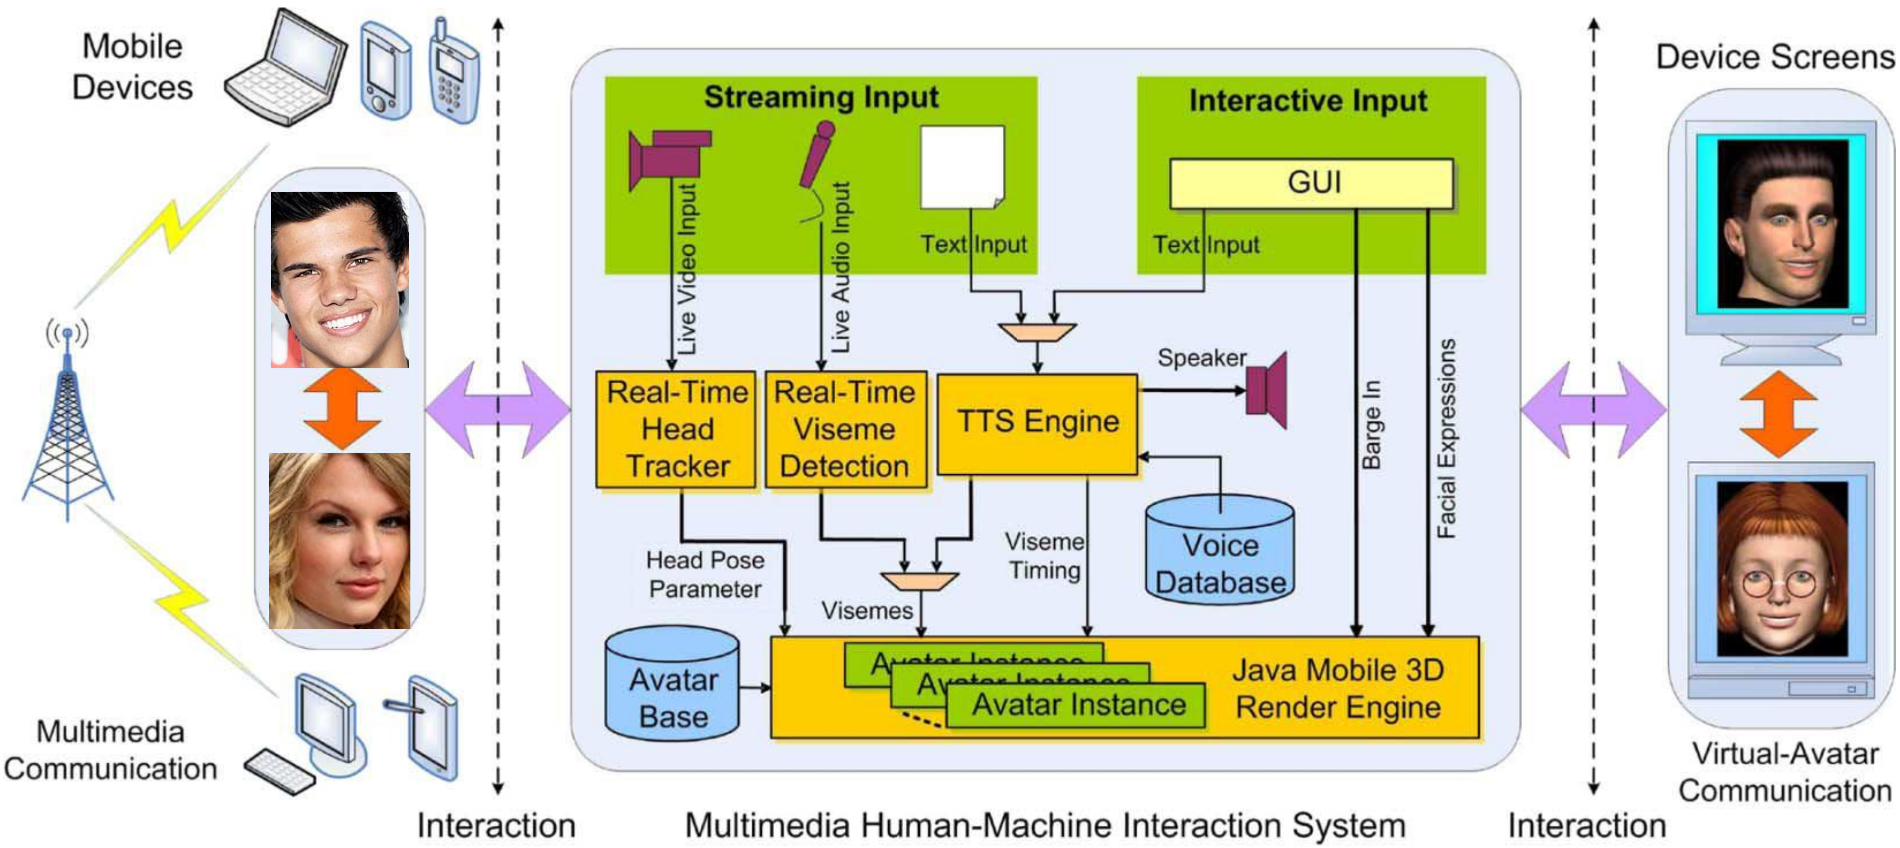

Y. Fu, et. al., IEEE Transactions on Circuits and Systems for Video Technology (TCSVT), 2008.

hMouse

Virtual Mouse Event	Description
Hold and Release	Use/activate/control and disuse/stop/quit
Move Cursor	Vertical and horizontal navigation, coarse localization
Point	Up and down fine tune, precise localization
Click Button	Left and right virtual mouse button down and up

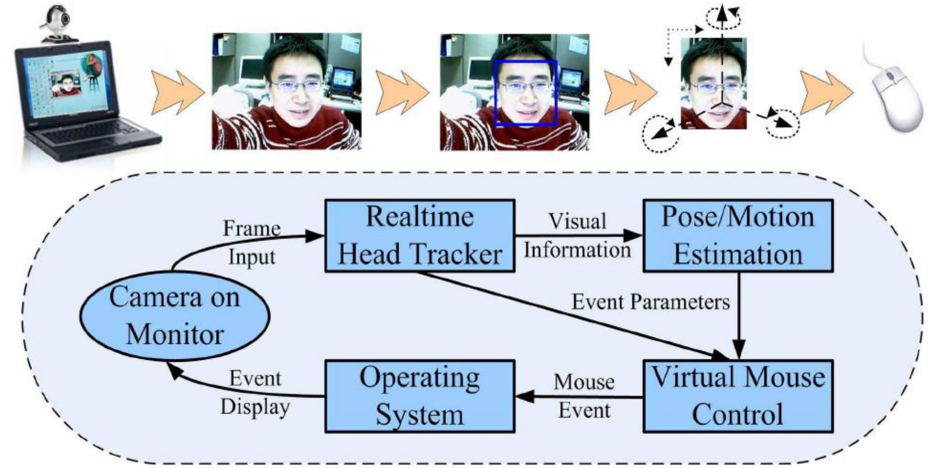

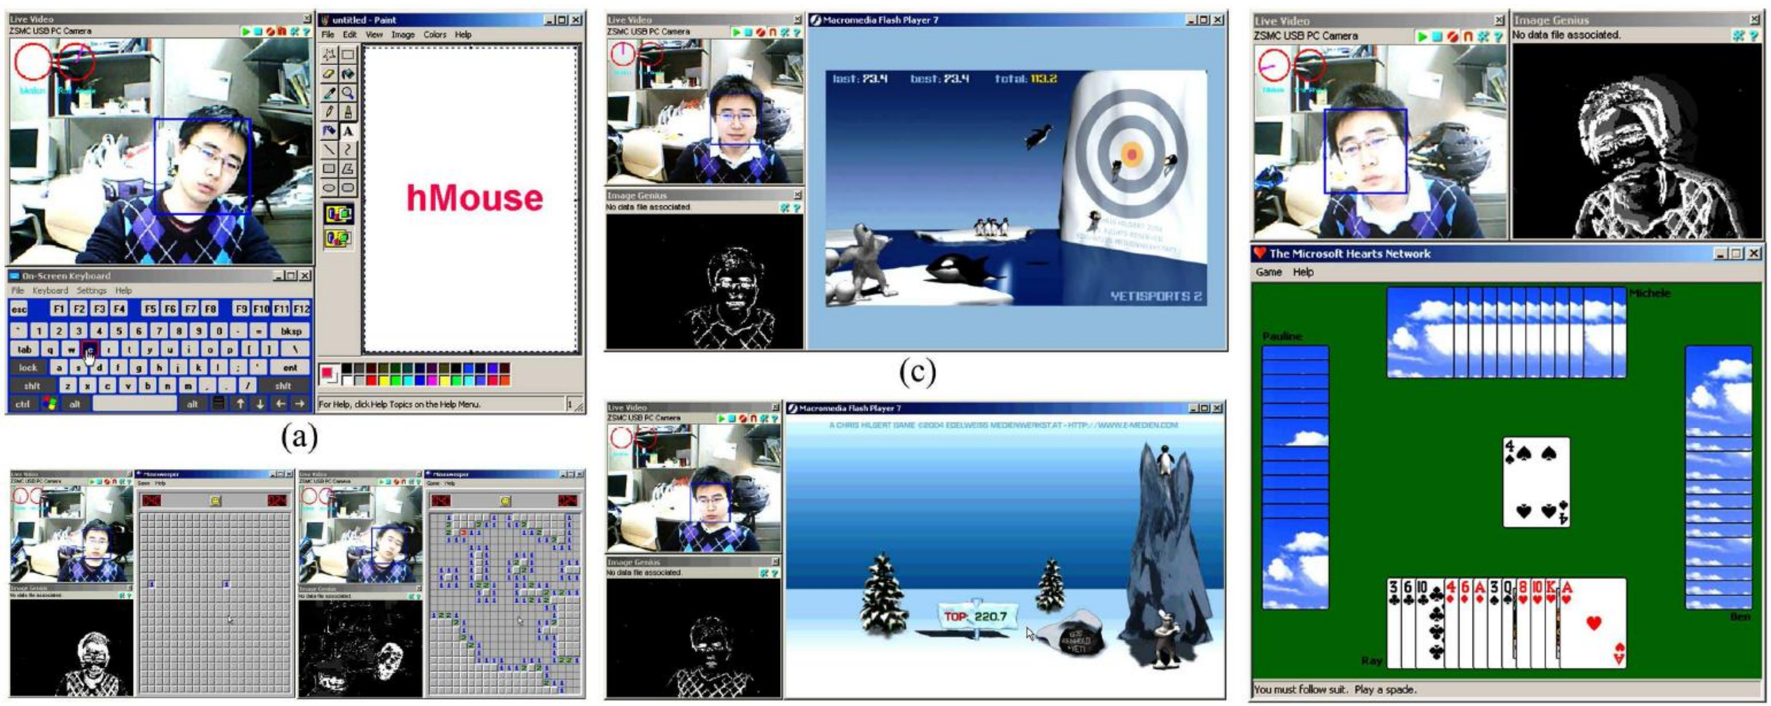

(a)

(c)

AFOSR Trust and Influence Award -- Social Media Ecosystems

Air Force Support, 11/2011 – 8/2012

(PI)Yun Raymond Fu, CSE

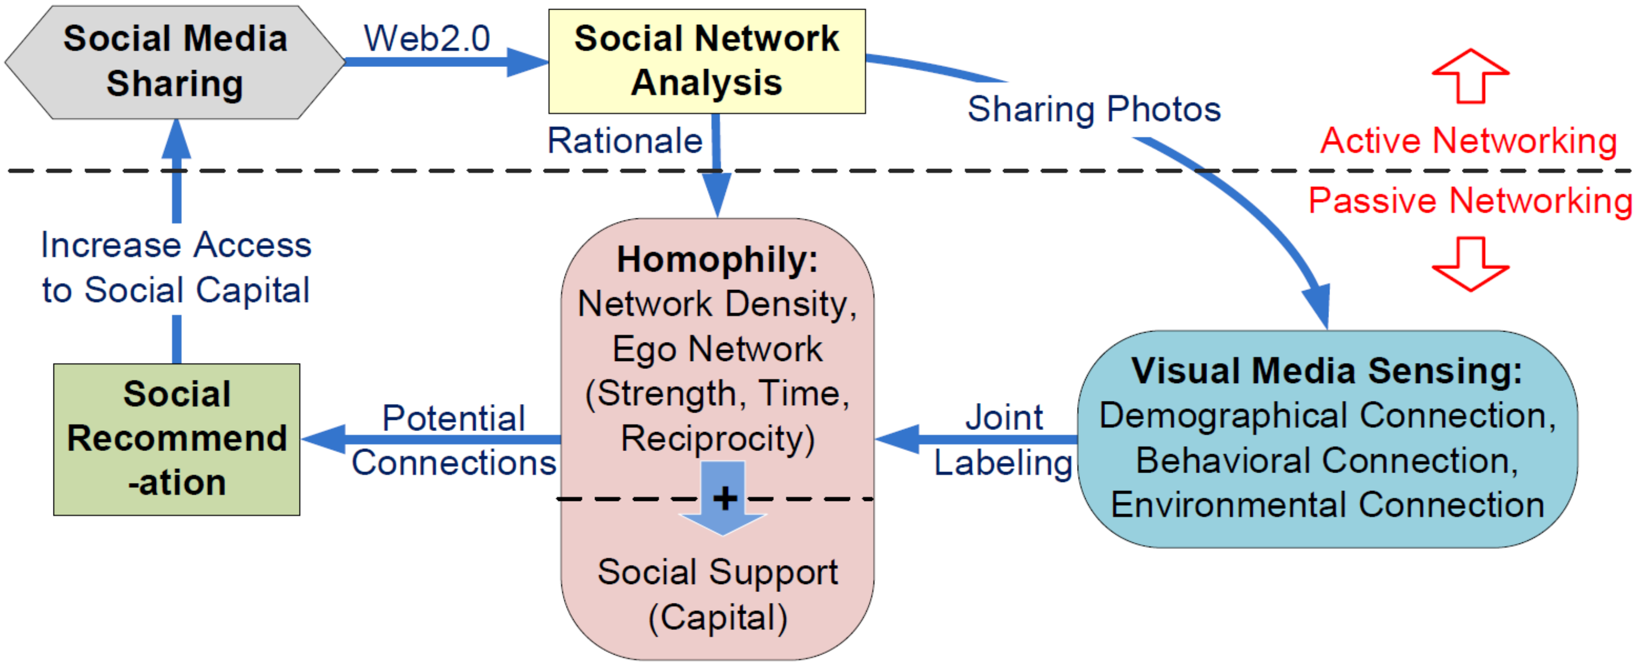

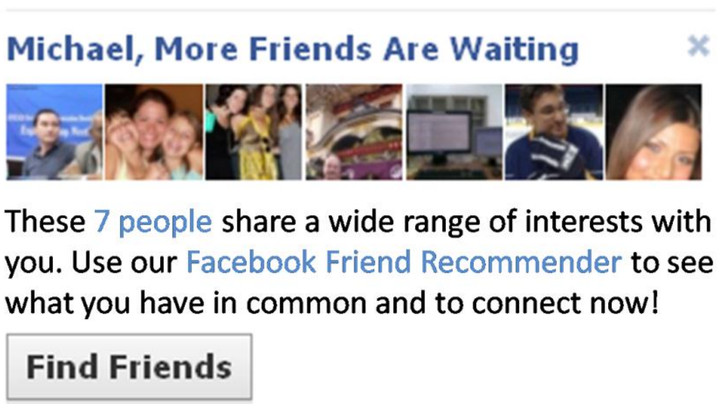

	Demographic	Behavior	Context
Left Image	Asian, Female	Musician	Violin
Right Image	Asian, Female	Musician	Violin
Correlation	1	1	1

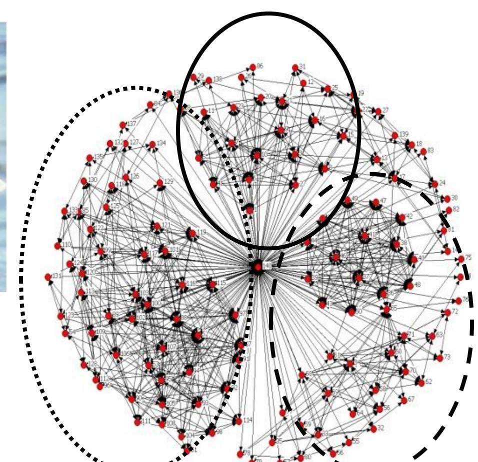

Google Faculty Research Award

Social Media Computing

1/2011 – now

(**PI**)**Yun Raymond Fu**, CSE

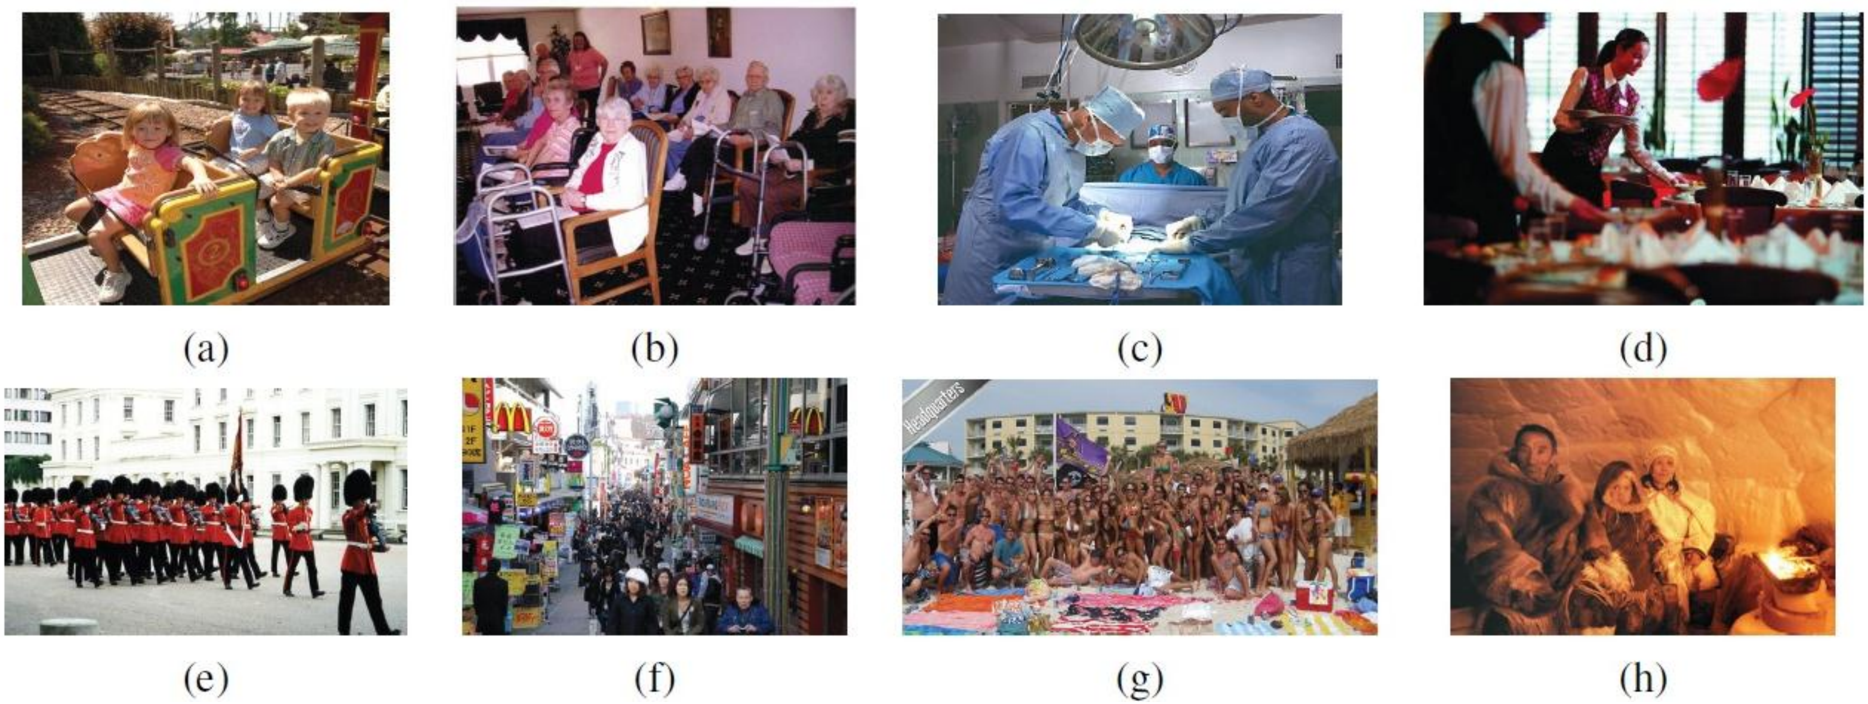

Demographic sensing. (a) Amusement park: kids (age). (b) Nursing home: old people (age). (c) Hospital: doctor (occupation). (d) Restaurant: waitress (occupation). (e) Buckingham palace: British soldier (culture and occupation). (f) Shopping area: female crowd (gender). (g) Beach: swimwear and crowd (social settings). (h) Arctic: Eskimo (ethnic group).

Others

[Augmented Reality]

[Object Recognition]

[Ethnicity, age, gender Estimation]

[Auto-Annotator]

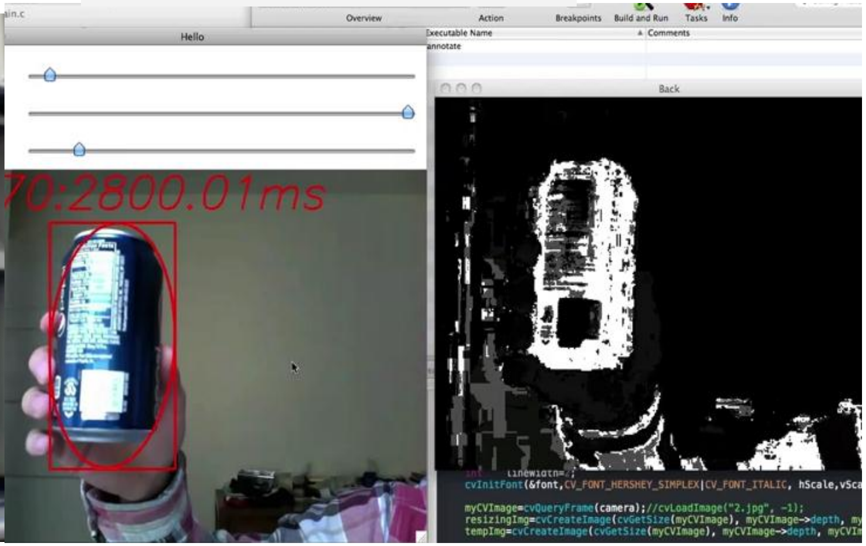

Demos

- Humanoid Avatar
- EAVA: A 3D Emotive Audio-Visual Avatar
- hMouse: Head Mouse
- Face Mask: Head Tracking and Pose Estimation
- M-Face and FaceTransfer
- Realtime Shrug Detector
- Object detection
  http://www.youtube.com/user/NotTubeIm?feature=mhee#p/u/1/9AHxGYdCL9w
- Label Ralation
  http://www.youtube.com/watch?v=LDJihDhN-qc&feature=youtu.be

http://indigo.ece.neu.edu/~yunfu/research.htm

Thank you!