Computer Science
Department

Technion-Israel Institute of Technology

Numerical Geometry

Ron Kimmel

Information Society
Technologies

Partly supported
European FP 6 NoE
grant No. 507752
MUSCLE

Students


Michal Aharon
CS Technion


Ilya Blayvas
CS Technion


Alexander Bronstein
EE-CS Technion

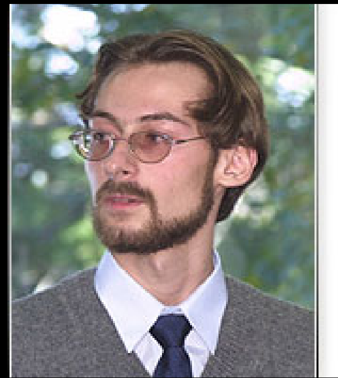
Michael Bronstein
CS Technion

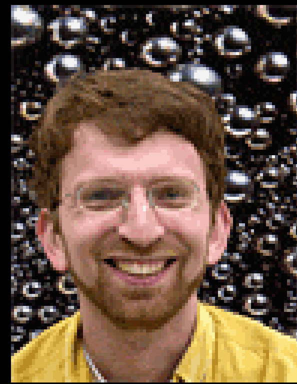
Alexander Brook
Math Technion


Asi Elad
IO-Image


Roman Goldenberg
CS Technion

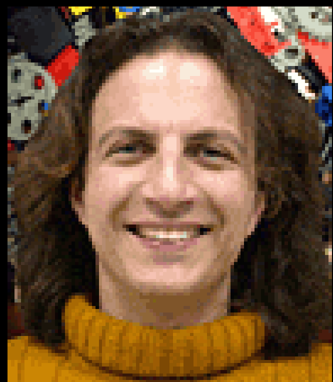
Eyal Gordon
CS Technion

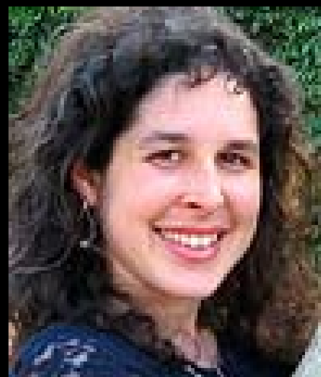
Michal Holzman Gazit
EE-CS Technion


Alon Spira
CS Technion


Gil Zigelman
3DV Systems

Colleagues

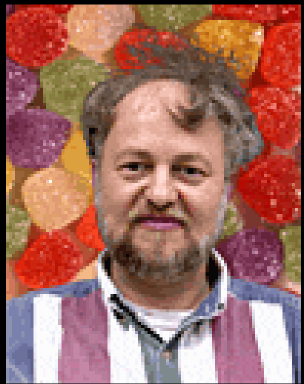

Alfred M. Bruckstein
CS Technion

Laurent Cohen
Paris Dauphine

Vicent Caselles
Math. Barcelona

Nahum Kiryati
Tel Aviv Univ.

Ehud Rivlin
CS Technion

Michael Rudszky
CS Technion

Guillermo Sapiro
EE Minnesota

James Sethian
UC Berkeley

Nir Sochen
Tel Aviv Univ.

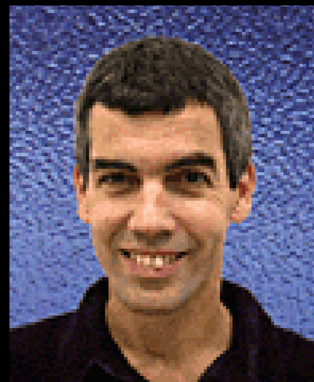

Irad Yavneh
CS Technion

On Edge Detection and Integration

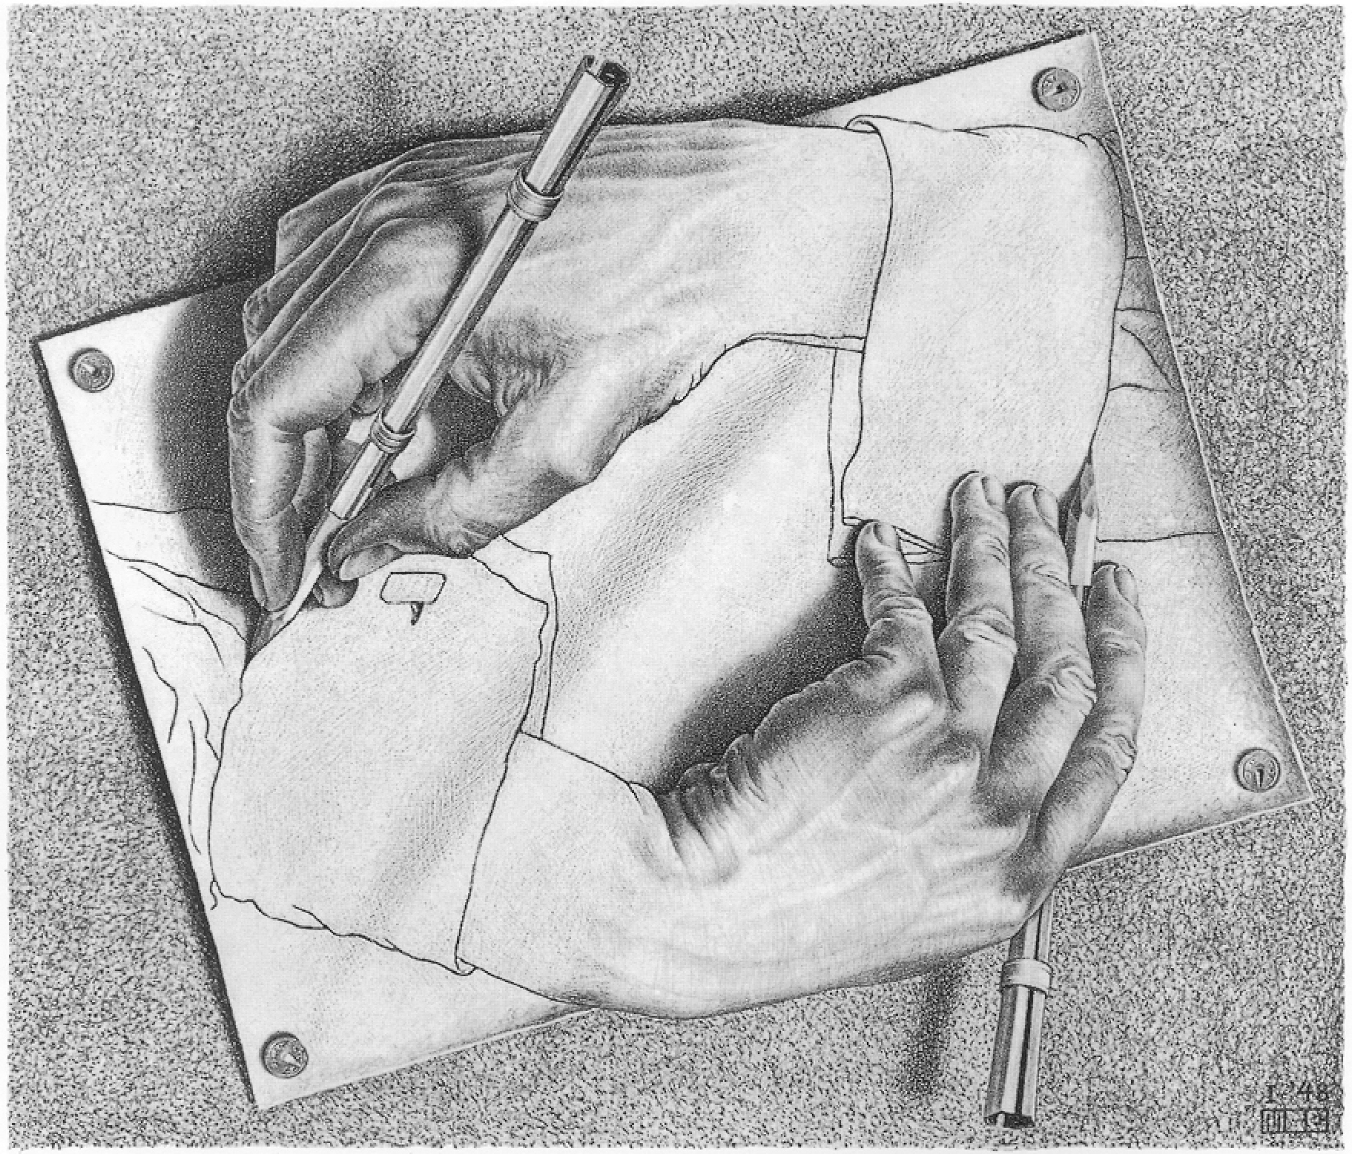

Edge Detection

- Edge Detection:
  - The process of labeling the locations in the image where the gray level's "rate of change" is high.
    - OUTPUT: "edgels" locations, direction, strength

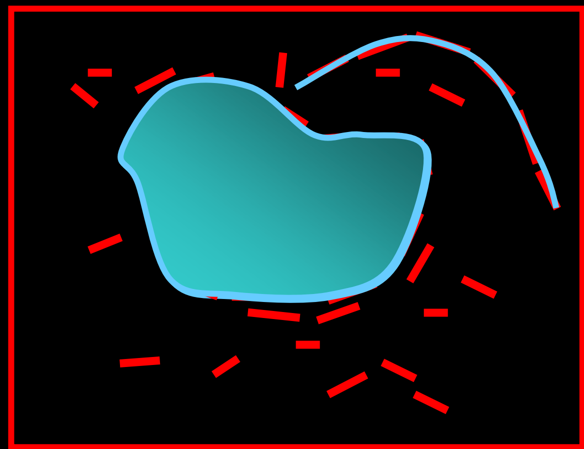

- Edge Integration:
  - The process of combining "local" and perhaps sparse and non-contiguous "edgel"-data into meaningful, long edge curves (or closed contours) for segmentation
    - OUTPUT: edges/curves consistent with the local data

The Classics

- Edge detection:
  - Sobel, Prewitt, Other gradient estimators
  - Marr Hildreth
    zero crossings of $\Delta G * I$
  - Haralick/Canny/Deriche et al.
    "optimal" directional local max of derivative
    $I_{\xi\xi} = 0$

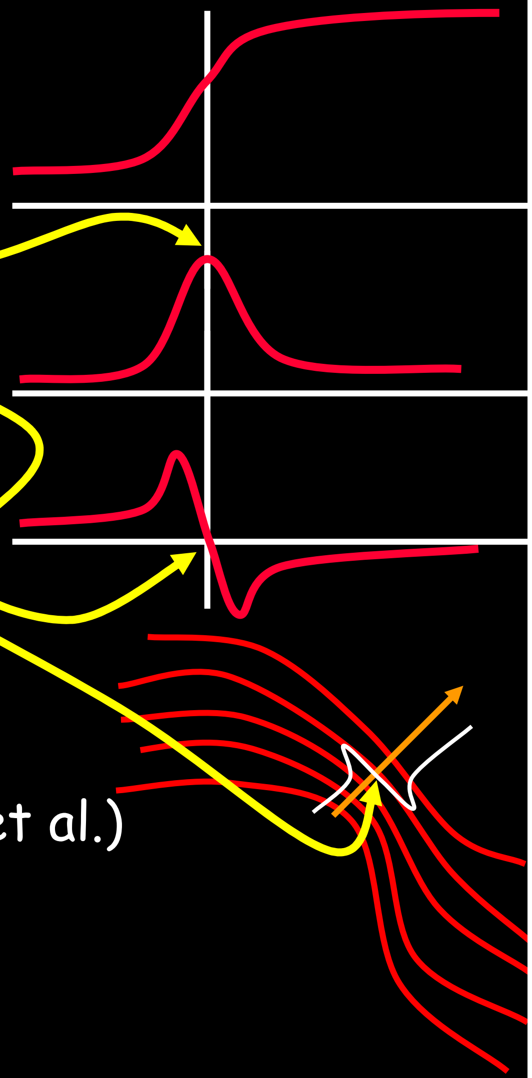

- Edge Integration:
  - tensor voting (Rom, Medioni, Williams, …)
  - dynamic programming (Shashua & Ullman)
  - generalized "grouping" processes (Lindenbaum et al.)

The "New-Wave"

- Snakes
- Geodesic Active Contours
- (Variational) Model Driven Edge Detection

Image

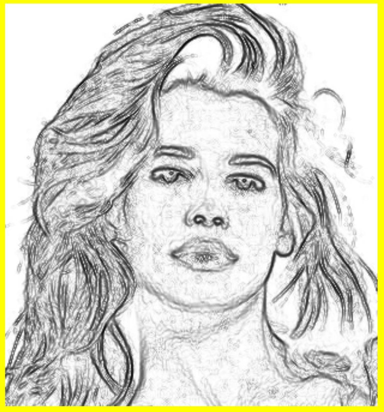

Edge Indicator Function

"nice" curves that optimize a functional of g( ), i.e.

$$\int_{curve} g(\,)ds$$

nice: "regularized", smooth, fit some prior information

Edge Curves

$$g(x,y) = \frac{1}{1+|\nabla(G_\sigma * I)|^2}$$

Segmentation

Segmentation

- Ultrasound images

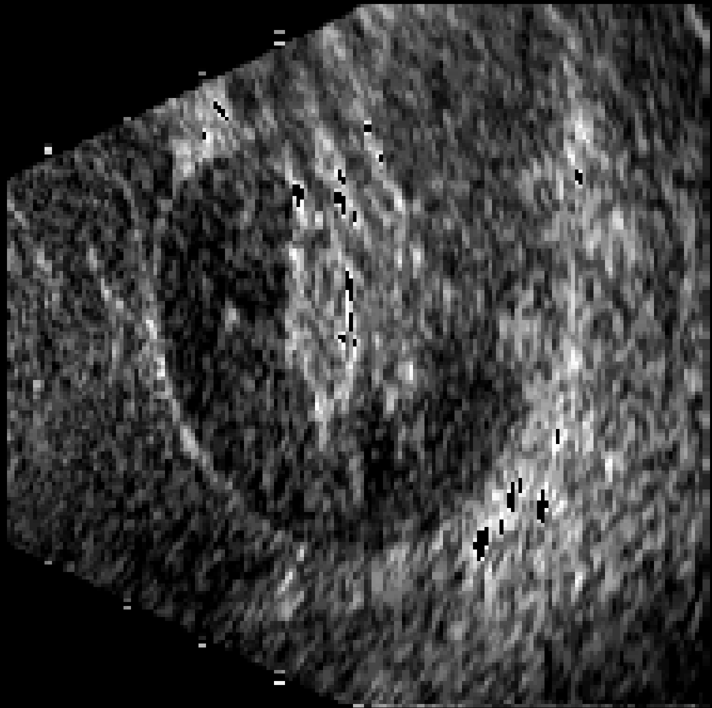

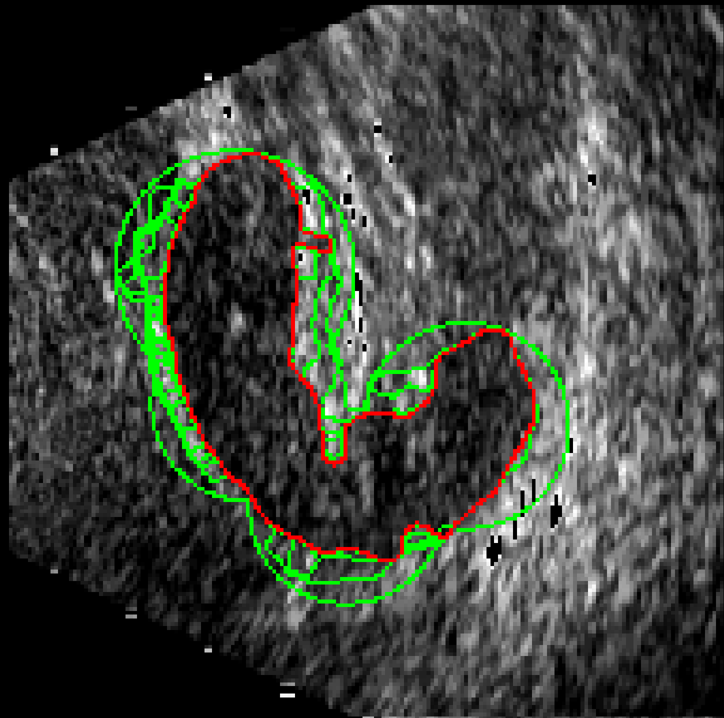

Caselles Kimmel Sapiro *1995*

Pintos

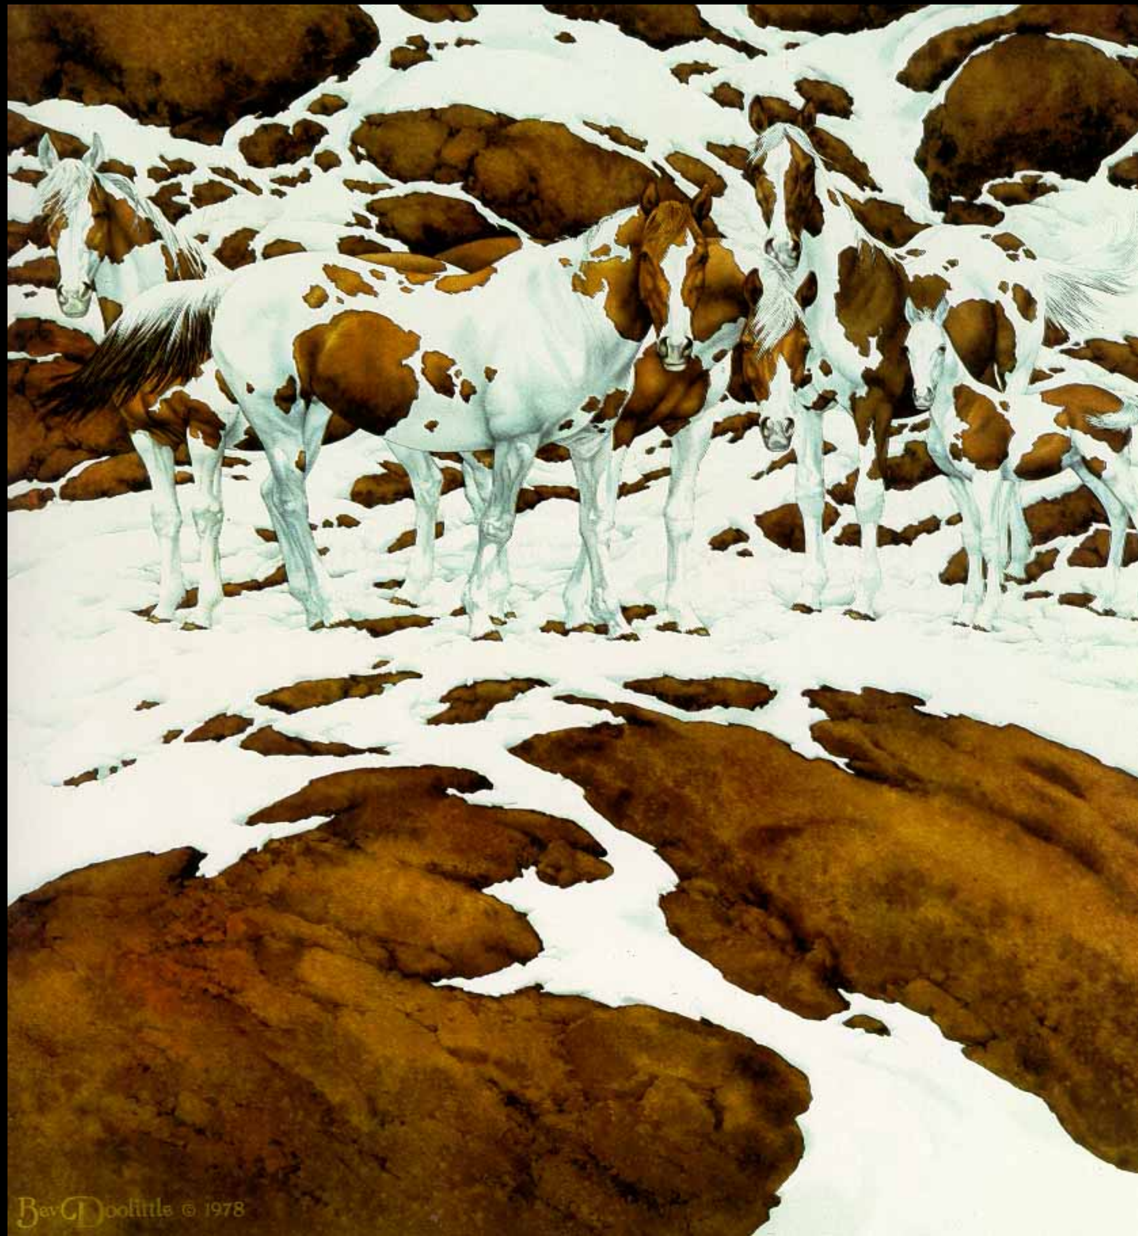

Segmentation

- With a good prior who needs the data...

Wrong Prior???

$$\int g(C)ds \quad \Rightarrow$$

$$\frac{dC}{dt} = \left(g(C)\kappa - \left\langle \nabla g(C), \vec{N} \right\rangle\right)\vec{N}$$

Experiments - Color Segmentation

Goldenberg, Kimmel, Rivlin, Rudzsky,

IEEE T-IP 2001

$$\iint g(S)da \Rightarrow$$

$$\frac{dS}{dt} = \left(g(S)H - \left\langle \nabla g(S), \vec{N} \right\rangle\right)\vec{N}$$

Tumor in 3D MRI

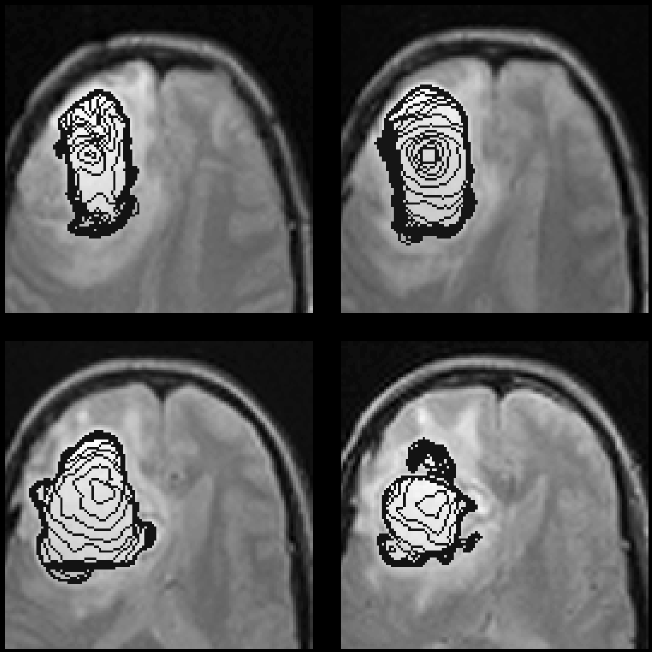

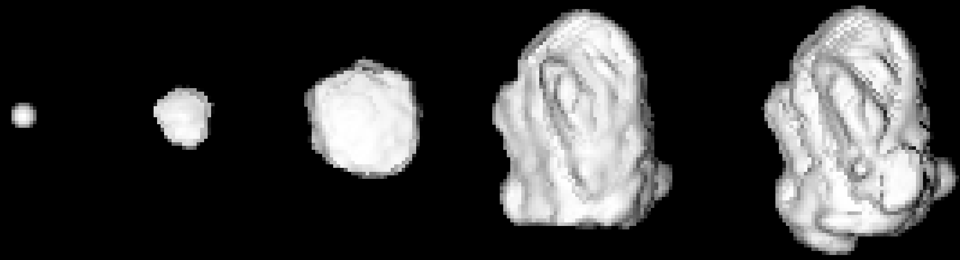

Caselles,Kimmel, Sapiro, Sbert, IEEE T-PAMI 97

$$\iiint g(M)dv \quad \Rightarrow$$

$$\frac{dM}{dt} = \left(g(M)H - \left\langle \nabla g(M), \vec{N} \right\rangle\right)\vec{N}$$

Segmentation in 4D

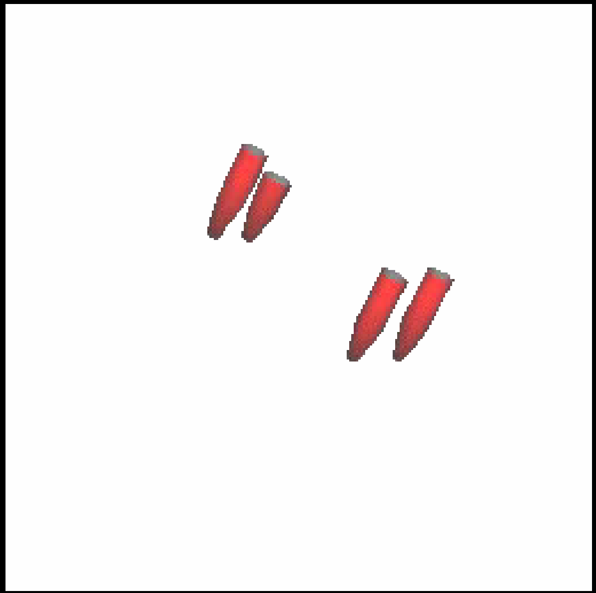

Malladi, Kimmel, Adalsteinsson,
Caselles, Sapiro, Sethian
SIAM Biomedical workshop 96

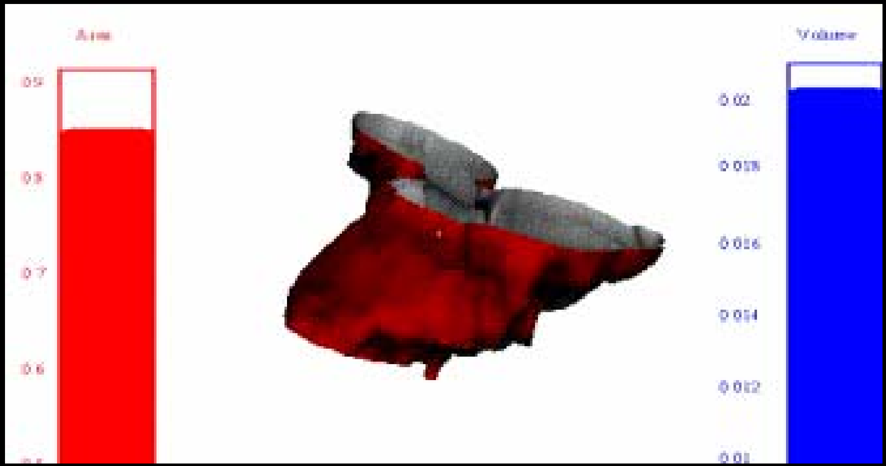

$$\int g(C)ds \quad \Rightarrow$$

$$\frac{dC}{dt} = \left(g(C)\kappa - \left\langle \nabla g(C), \vec{N} \right\rangle\right)\vec{N}$$

Tracking in Color Movies

Goldenberg, Kimmel, Rivlin, Rudzsky,

IEEE T-IP 2001

$$\int g(C)ds \quad \Rightarrow$$

$$\frac{dC}{dt} = \left(g(C)\kappa - \left\langle \nabla g(C), \vec{N} \right\rangle\right)\vec{N}$$

Tracking in Color Movies

Goldenberg, Kimmel, Rivlin, Rudzsky,
IEEE T-IP 2001

Edge Gradient Estimators

$$I(x, y) \longrightarrow \nabla I$$

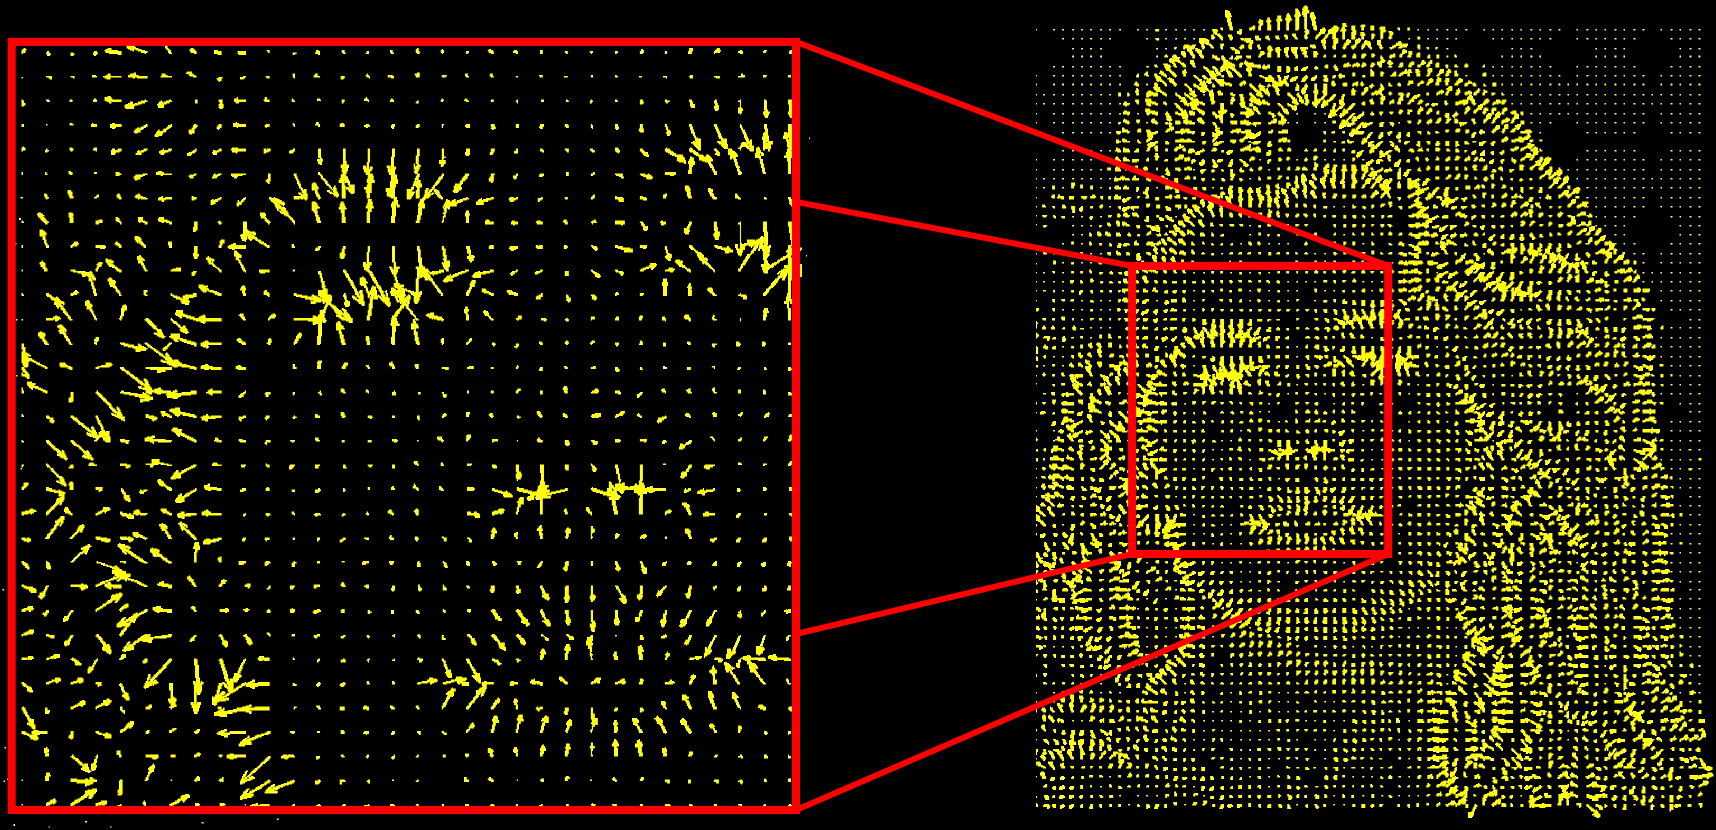

Edge Gradient Estimators

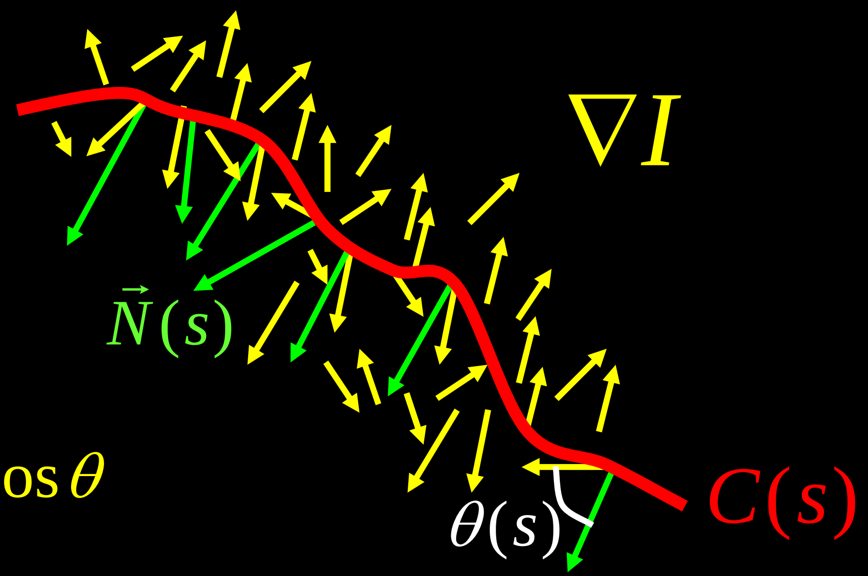

$$\langle \vec{N}, \nabla I \rangle = |\nabla I| \cos\theta$$

- We want a curve with large $|\nabla I|$ points and small θ's so:
- Consider the functional $E(C) = \int_C \langle \vec{N}, \nabla I \rangle ds$

The Classic Connection

Suppose $\rho(\alpha) = \alpha$ and we consider a closed contour for $C(s)$.
We have

$$E(C) = \oint_{C(s)} \left\langle \nabla I, \vec{N} \right\rangle ds$$

and by Green's Theorem we have

$$= \iint_{\text{Area within } C(s)} \operatorname{div}\left(\nabla I\right) dxdy$$

$$= \iint_{\text{Area within } C(s)} \Delta I(x, y) dxdy$$

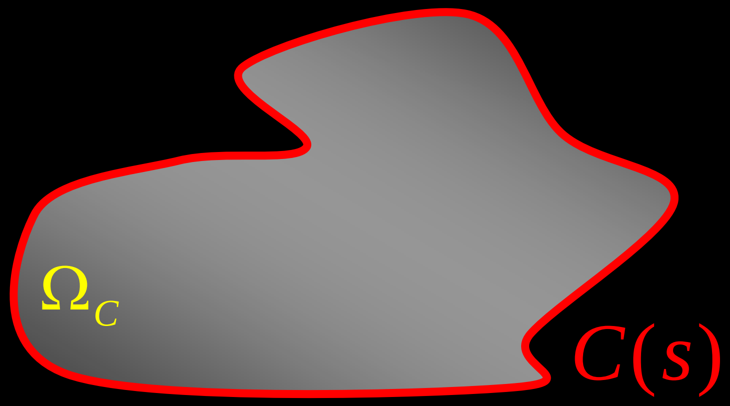

Haralick/Canny-like Edge Detector

- Haralick suggested $I_{\xi\xi} = 0$ as edge detector

$$\Delta I = I_{xx} + I_{yy} = I_{\xi\xi} + I_{\eta\eta}$$

$$I_{\xi\xi} = \Delta I - I_{\eta\eta}$$

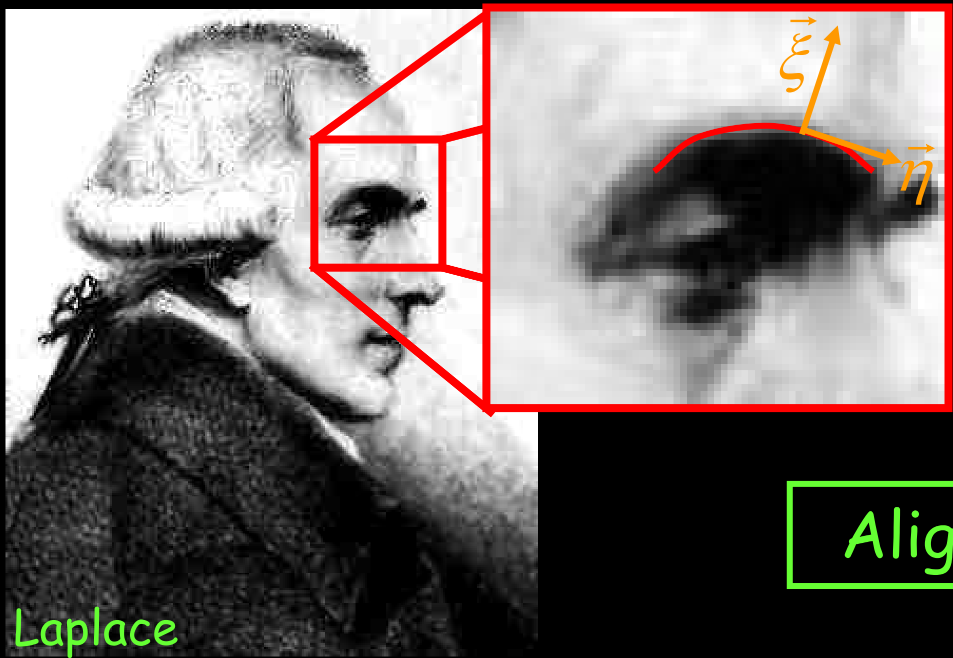

Alignment (ΔI)

Topological Homogeneity ($I_{\eta\eta}$)

Haralick/Canny Edge Detector $I_{\xi\xi} = 0$

- Haralick $I_{\xi\xi} = \Delta I - I_{\eta\eta}$

$$I_{\eta\eta} = \mathrm{div}\left(\frac{\nabla I}{|\nabla I|}\right)|\nabla I| = \kappa_I |\nabla I|$$

co-area

Kronrod

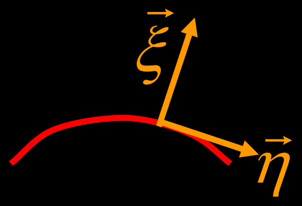

$$\iint_{\Omega_C} I_{\eta\eta}\, dxdy = \iint_{\Omega_C} \kappa_I |\nabla I|\, dxdy = \iint_{\Omega_C} \kappa_I\, ds dI = 2\pi \int dI$$

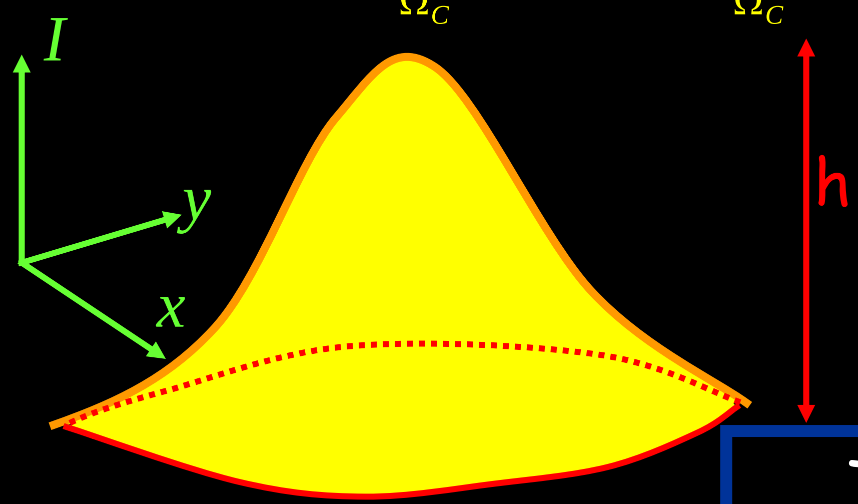

$$\iint_{\Omega_C} I_{\eta\eta}\, dxdy = 2\pi h$$

Thus, $I_{\xi\xi} = 0$ indicates optimal alignment + topological homogeneity

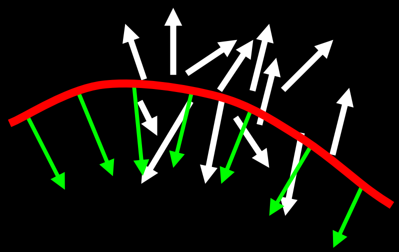

Closed contours $L_\rho + \varepsilon L_g$

EL eq. $$\left(\text{sign}(\langle \vec{N}, \vec{V} \rangle)\,\text{div}(\vec{V}) + \varepsilon\left(\kappa g - \langle \vec{N}, \nabla g \rangle\right)\right)\vec{N} = 0$$

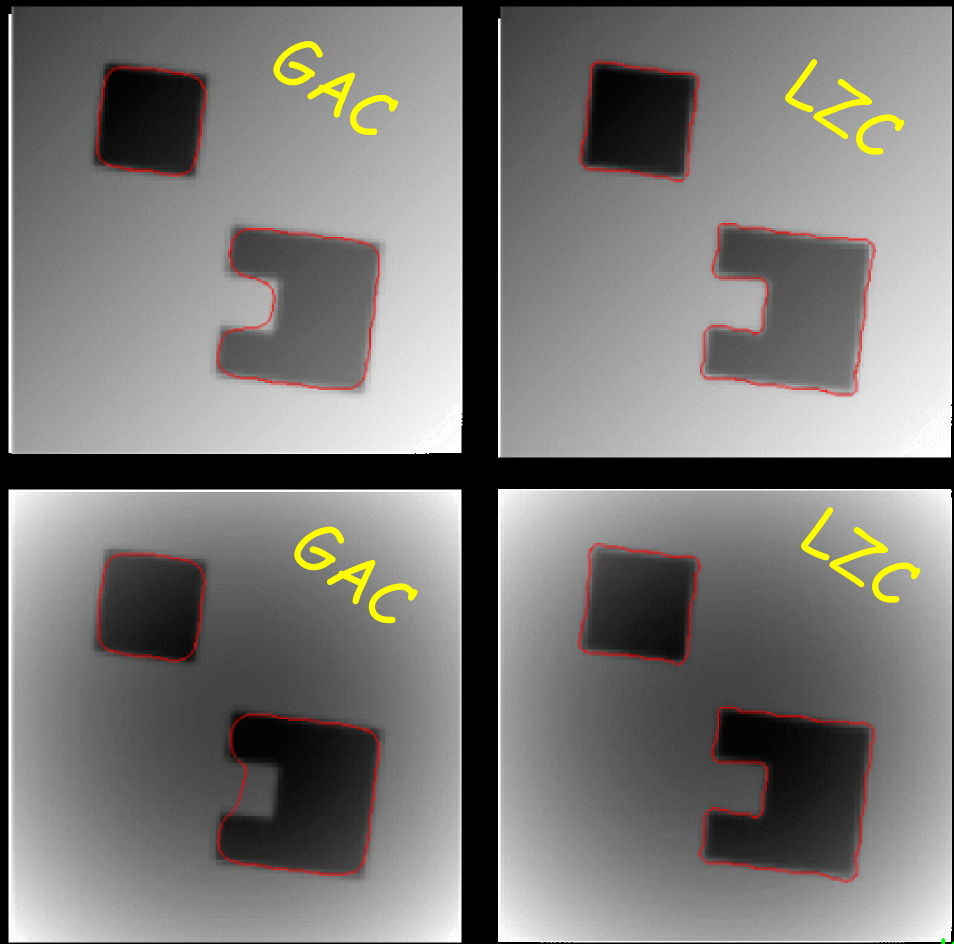

Kimmel-Bruckstein IJCV 2003

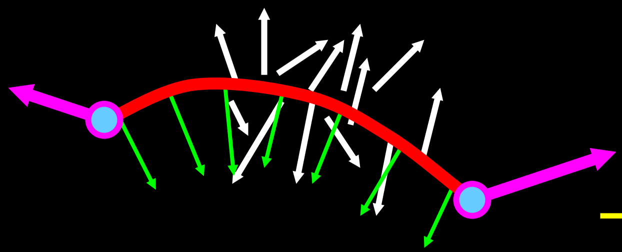

Open contours $L_\rho - L$

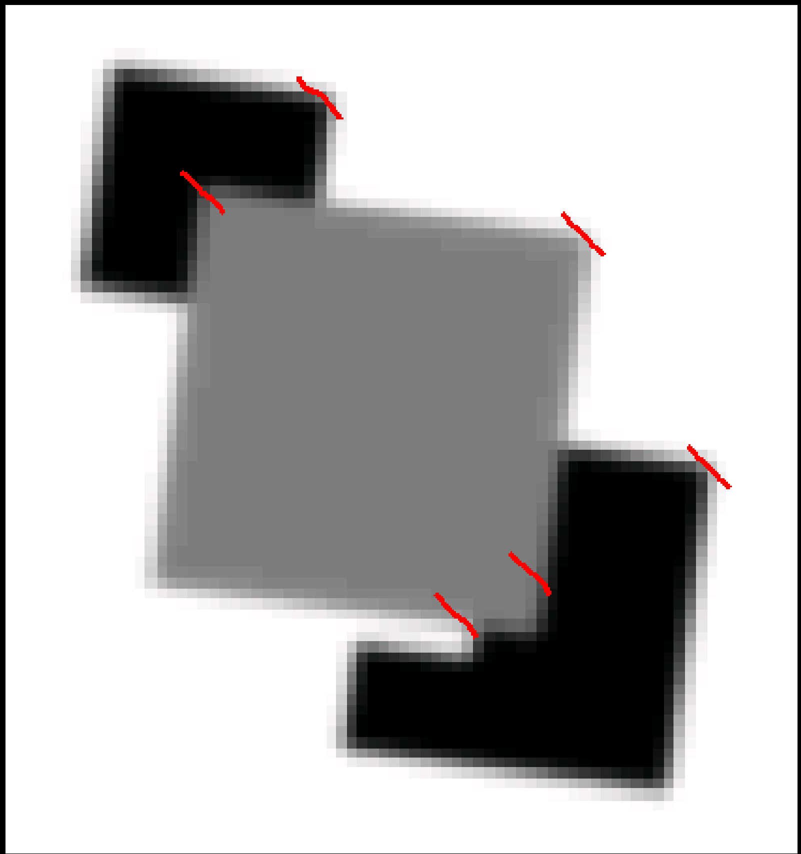

Kimmel-Bruckstein IJCV 2003

Geometric Measures

Weighted arc-length $\int_C g(C(s))ds \Rightarrow (\kappa g - \langle \nabla g, \vec{N} \rangle)\vec{N} = 0$

Weighted area $\iint_\Omega g(x,y)da \Rightarrow g\vec{N} = 0$

Alignment $\int_C \langle \vec{N}, \vec{V} \rangle ds \Rightarrow \mathrm{div}(\vec{V})\vec{N} = 0$

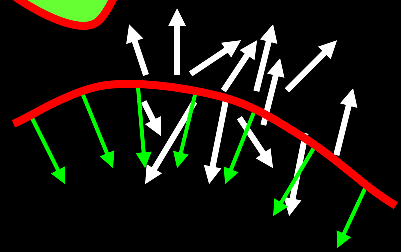

Robust-alignment $\int_C |\langle \vec{N}, \vec{V} \rangle| ds \Rightarrow \mathrm{sign}(\langle \vec{N}, \vec{V} \rangle)\mathrm{div}(\vec{V})\vec{N} = 0$

e.g. $\int_C \langle \vec{N}, \nabla I \rangle ds \Rightarrow \Delta I \vec{N} = 0$

Variational meaning for Marr-Hildreth edge detector

Kimmel-Bruckstein IJCV 2003

Geometric Measures

Minimal variance $$\iint_{\Omega_C} (I - c_1)^2 da + \iint_{\Omega \backslash \Omega_C} (I - c_2)^2 da$$

Simplified Mumford-Shah, Zhu-Yuille,
Chan-Vese, Max-Lloyd, regularized VQ, Threshold,… $\Rightarrow$ $$(c_1 - c_2)(I - \tfrac{c_1 + c_2}{2})\vec{N} = 0$$

$$\begin{cases} c_1 = \dfrac{\int_{\Omega_C} I da}{\int_{\Omega_C} da} \\ c_2 = \dfrac{\int_{\Omega \backslash \Omega_C} I da}{\int_{\Omega \backslash \Omega_C} da} \\ C_t = (c_2 - c_1)(I - \tfrac{c_1 + c_2}{2})\vec{N} = 0 \end{cases}$$

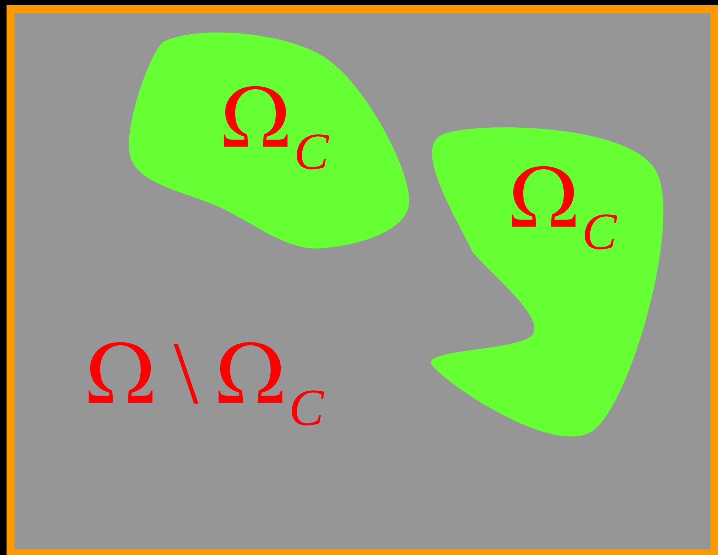

Geometric Measures

Robust minimal deviation $\iint_{\Omega_C} |I - c_1| da + \iint_{\Omega \setminus \Omega_C} |I - c_2| da$

$$\Rightarrow \left(|I - c_1| - |I - c_2|\right)\vec{N} = 0$$

$$\begin{cases} c_1 = \underset{\Omega_C}{\text{median}}\left(I(x, y)\right) \\ c_2 = \underset{\Omega \setminus \Omega_C}{\text{median}}\left(I(x, y)\right) \\ C_t = \left(|I - c_2| - |I - c_1|\right)\vec{N} \end{cases}$$

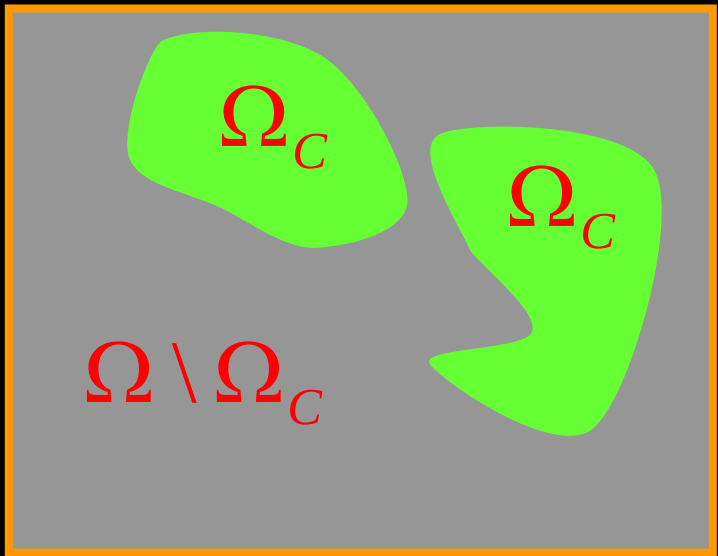

Tracking

Goldenberg, Kimmel, Rivlin, Rudzsky,
ECCV 2002

Tracking

Goldenberg, Kimmel, Rivlin, Rudzsky,
ECCV 2002

Tracking

Goldenberg, Kimmel, Rivlin, Rudzsky,
ECCV 2002

Information extraction

Goldenberg, Kimmel, Rivlin, Rudzsky,
ECCV 2002

Classification (dogs & cats)

walk run gallop cat...

Goldenberg, Kimmel, Rivlin, Rudzsky,
ECCV 2002

Classification (people)

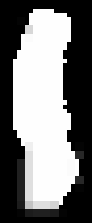

walk

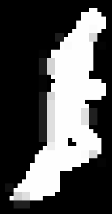

run

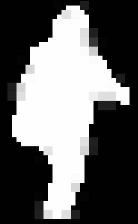

run45

Goldenberg, Kimmel, Rivlin, Rudzsky,
ECCV 2002

Thin Structures

$$I_{\xi\xi} = \Delta I - I_{\eta\eta} - I_{\gamma\gamma}$$
$$= \Delta I - H\left|\nabla I\right|$$

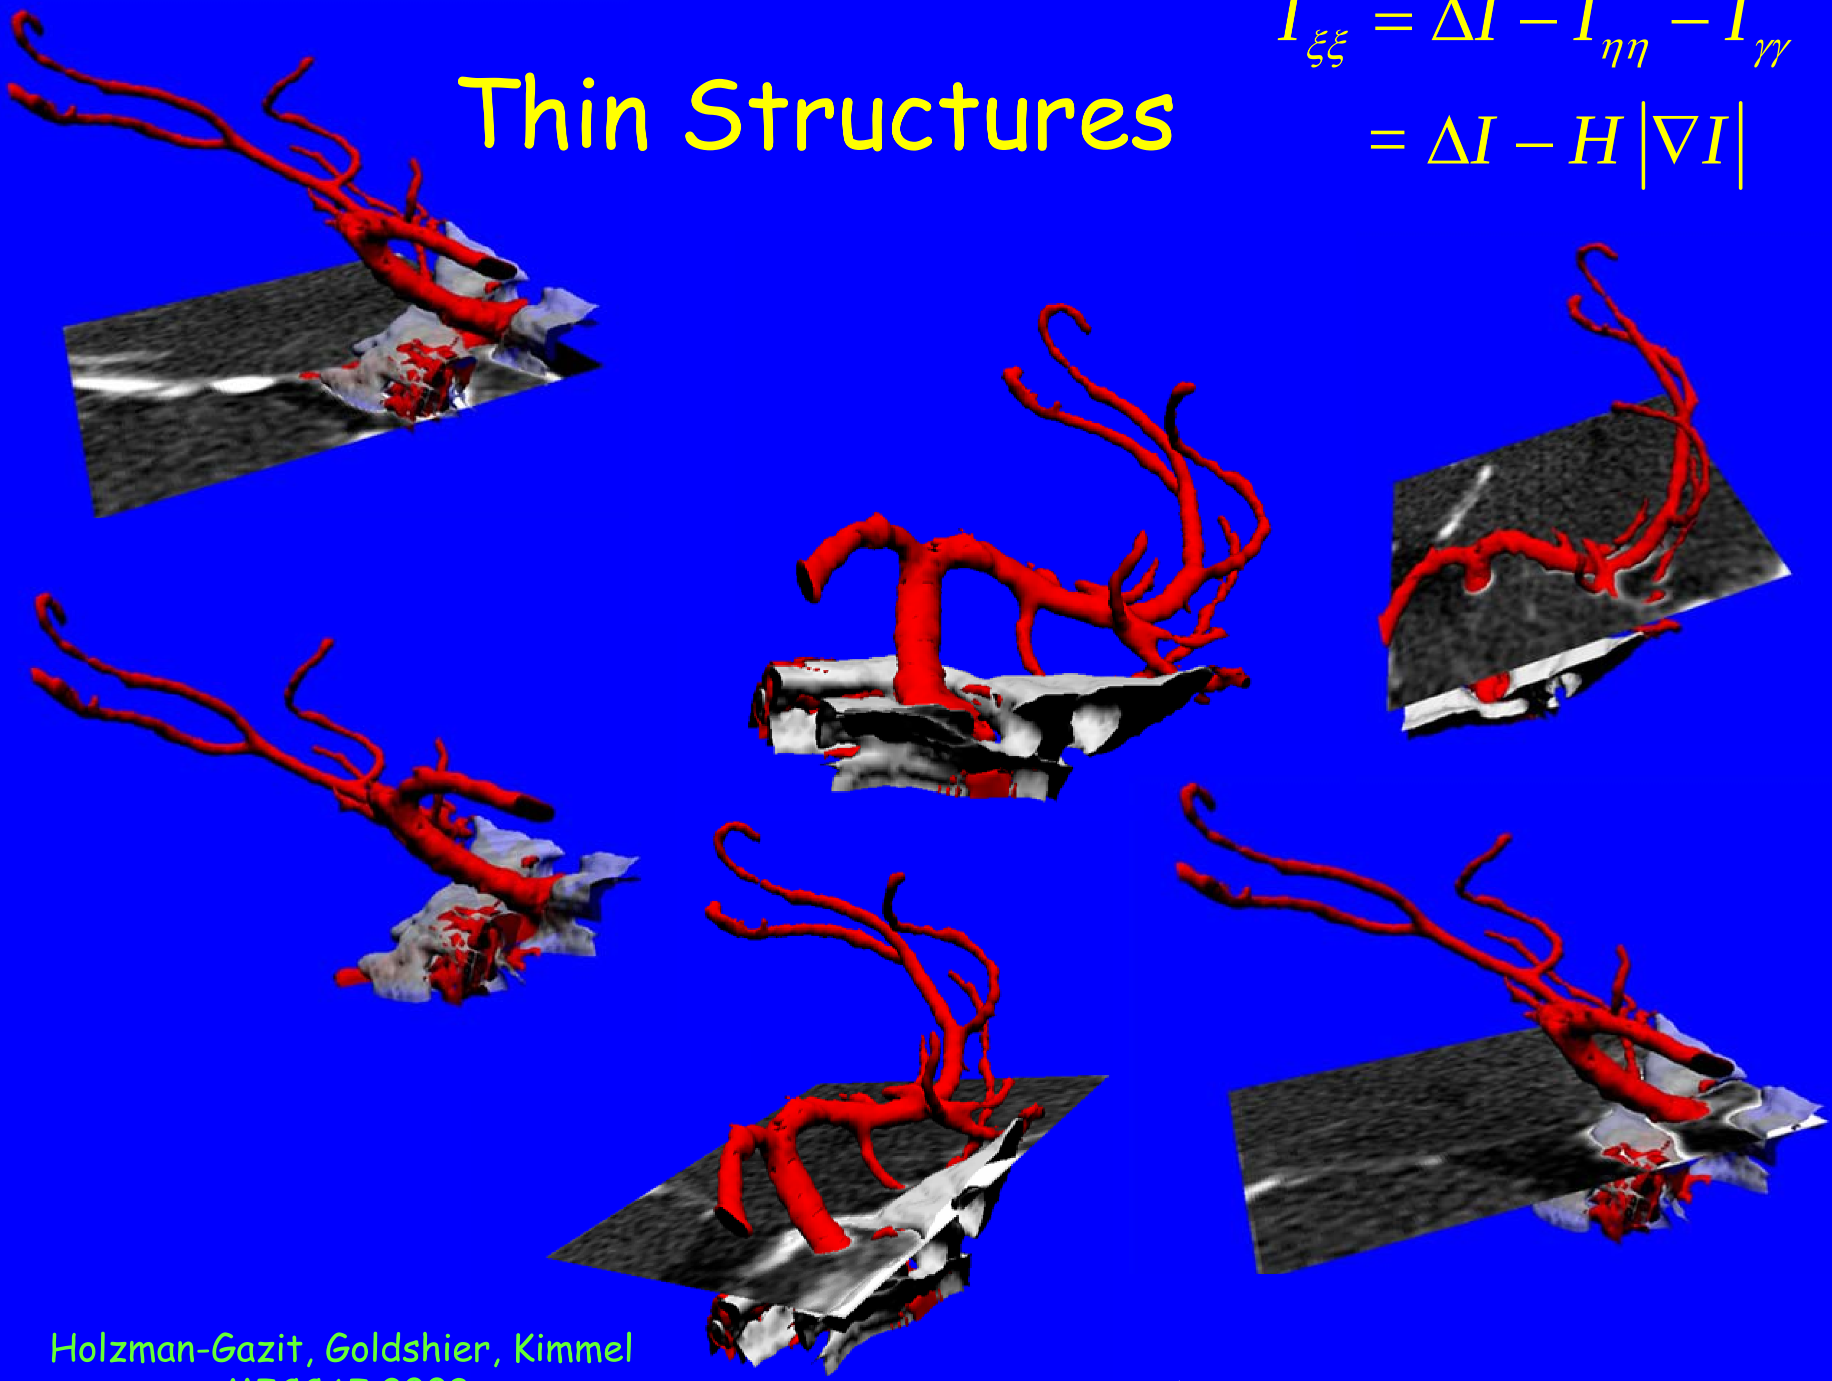

Holzman-Gazit, Goldshier, Kimmel
MICCAI 2003

Distance maps and minimal geodesics

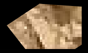

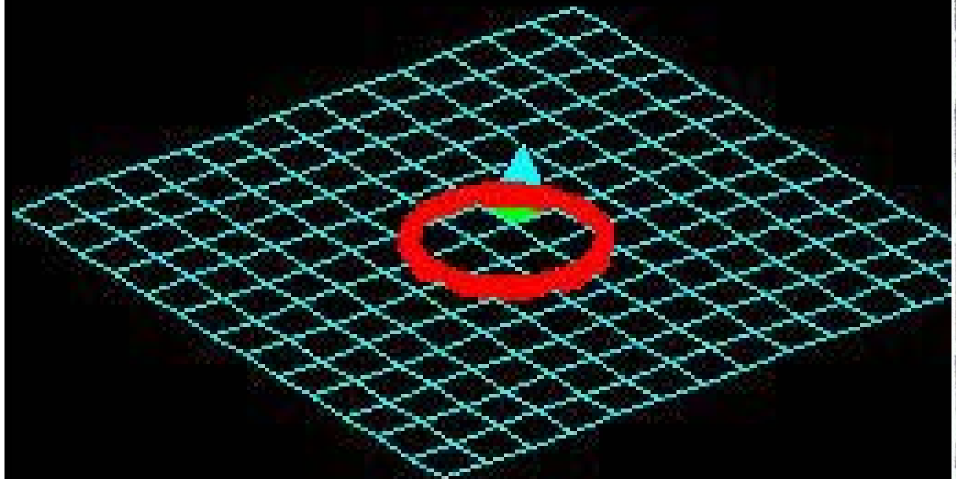

www.math.berkeley.edu/~sethian

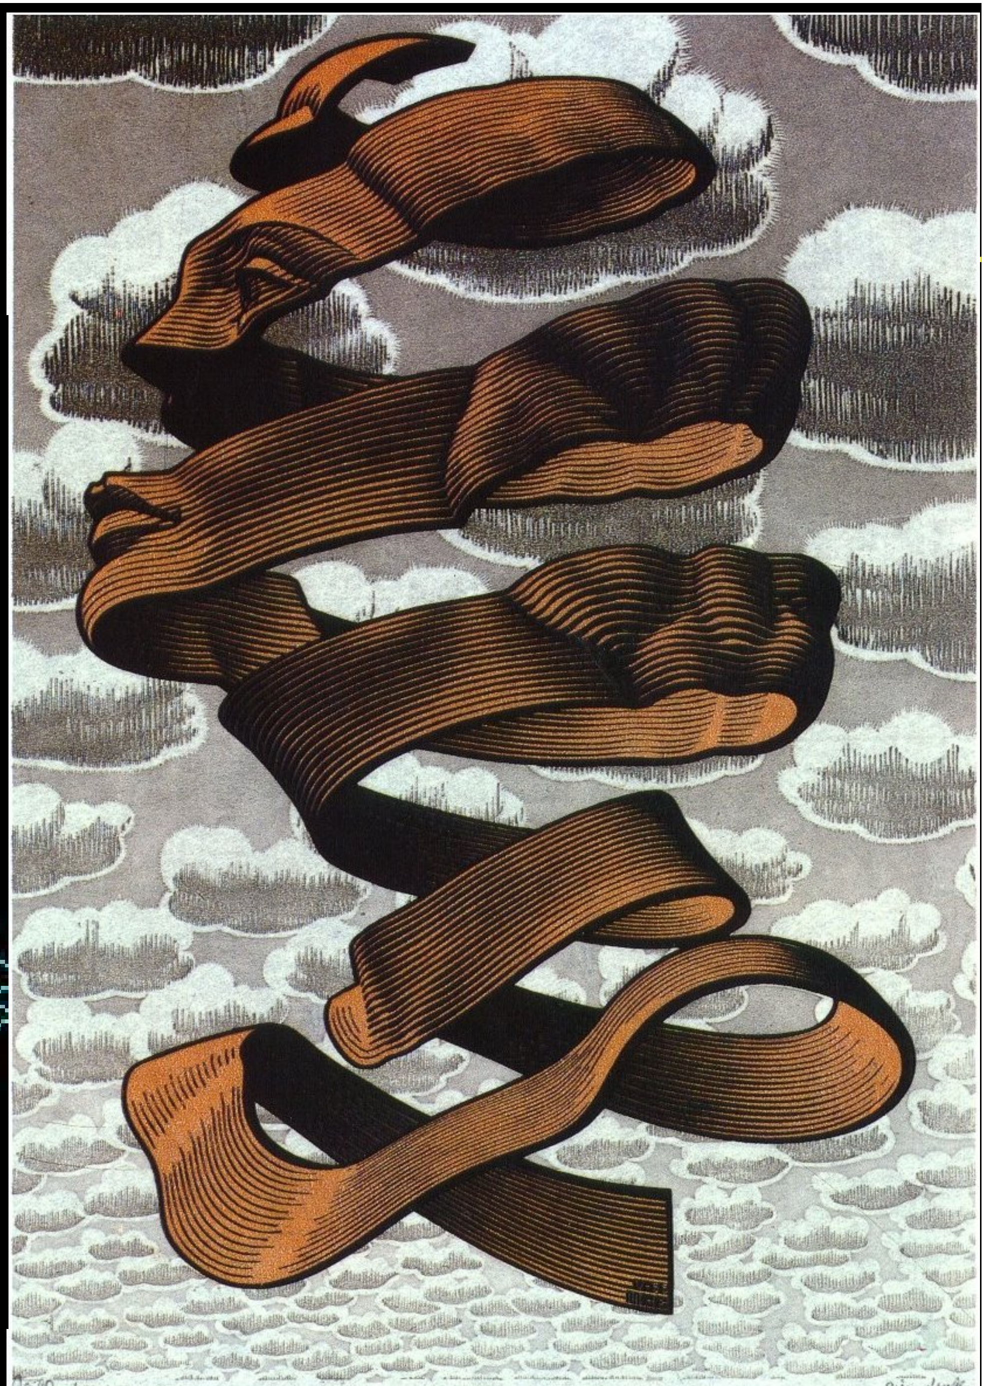

Escher, Rind 1955

Measuring geodesic distances: Fast Marching on Triangulated Domains

- Find distance *T(x)*, given *T(x0)=0*.

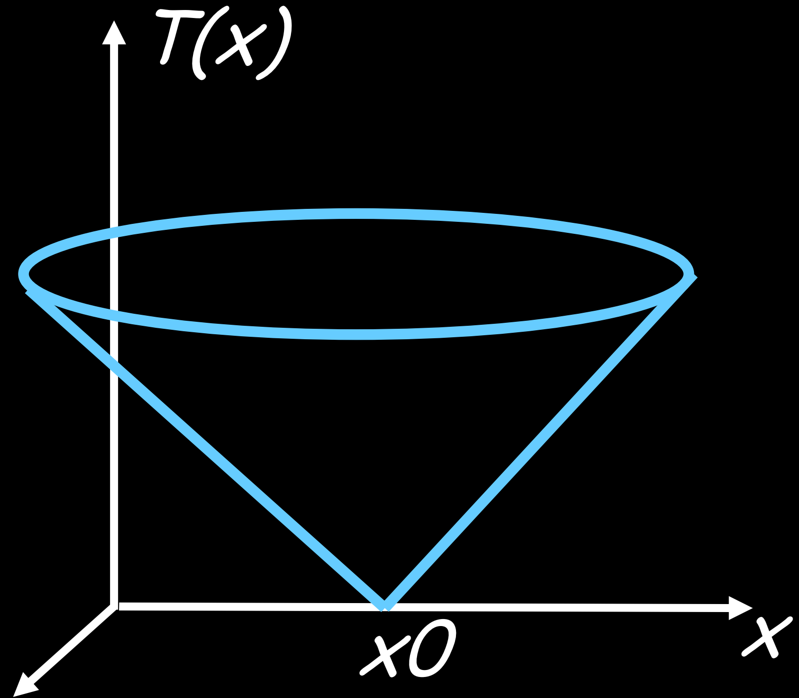

- Solution: *T(x)=|x-x0|*.
- $\left|\frac{d}{dx}T(x)\right| = 1$ except *x0*.
- Or in 2D $\left|\nabla T(x,y)\right| = 1$ where

$$\left|\nabla T(x,y)\right|^2 \equiv \left(\tfrac{d}{dx}T\right)^2 + \left(\tfrac{d}{dy}T\right)^2$$

Consistent Approximation

$$\left|\frac{d}{dx}T(x)\right| = 1$$

- $\left|\frac{d}{dx}T(x)\right| \approx \left|\max\{(T_i - T_{i-1})/h, (T_i - T_{i+1})/h, 0\}\right|$
- Updated *i* has always $T_i > \min\{T_{i-1}, T_{i+1}\}$
- `upwind' from where the `wind blows'

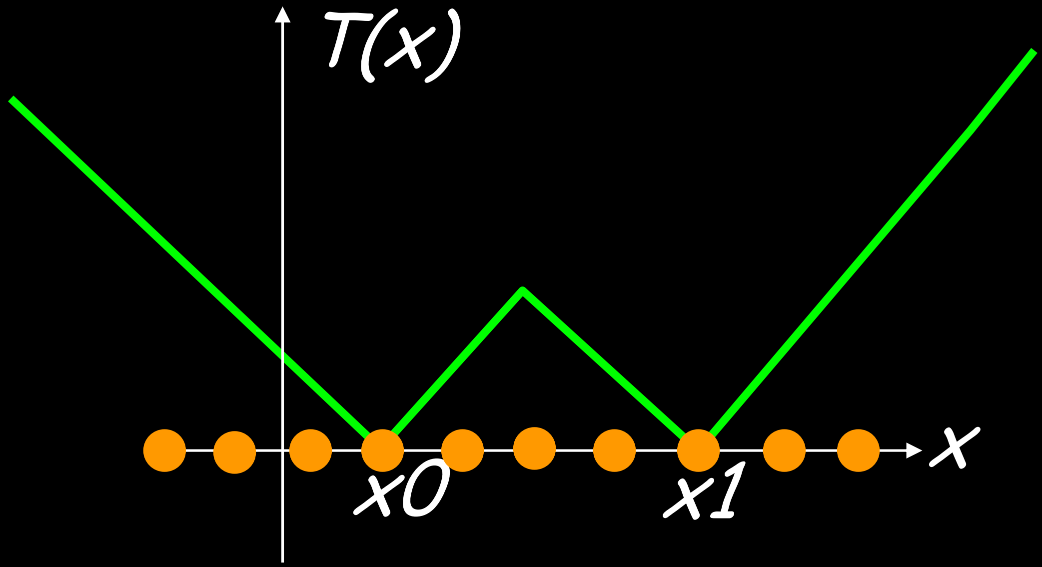

Update Procedure

Set $T_i = \infty$, and *T(x0)=T(x1)=0*.

REPEAT UNTIL convergence,

- FOR each *i*
  - $m = \min \{T_{i-1}, T_{i+1}\}$
  - $T_i = \min \{T_i, m + h\}$

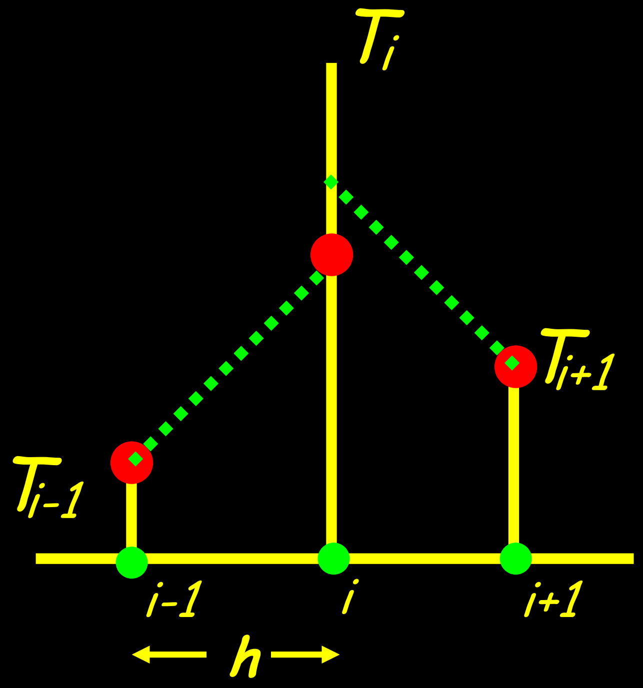

Update Order

What is the optimal order of updates?

Solution I: Scan the line successively left to right. *N* scans, i.e. $O(N^2)$

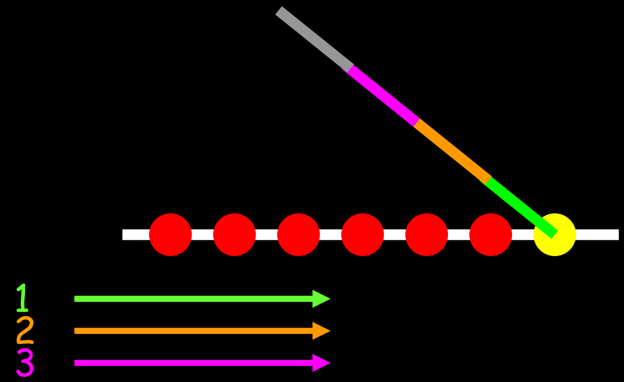

Solution II: Left to right followed by right to left. Two scans are sufficient. (Danielson`s distance map 1980)

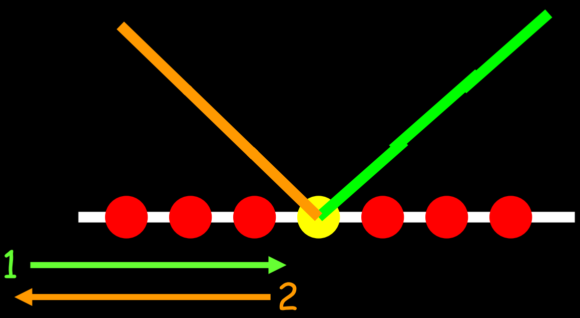

Solution III: Start from *x0*, update its neighboring points, accept updated values, and update their neighbors, etc.

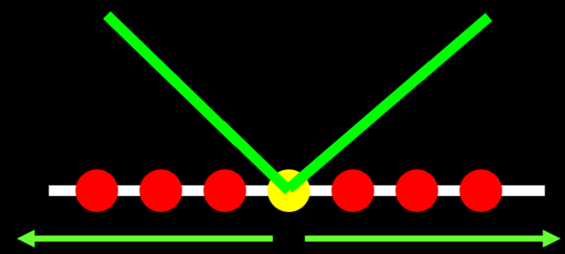

2D Approximation

Initialization:

- $\forall\{i, j\}: T_{ij} =$ given initial value or ∞

Update:

- $T_1 = \min\{T_{i-1,j}, T_{i+1,j}\}; \quad T_2 = \min\{T_{i,j-1}, T_{i,j+1}\};$

  IF $(|T_1 - T_2| < 1)$ THEN

$$T = \frac{T_1 + T_2 + \sqrt{2 - (T_1 - T_2)^2}}{2};$$

  ELSE $T = \min\{T_1, T_2\} + 1;$

  $T_{ij} = \min\{T_{ij}, T\};$

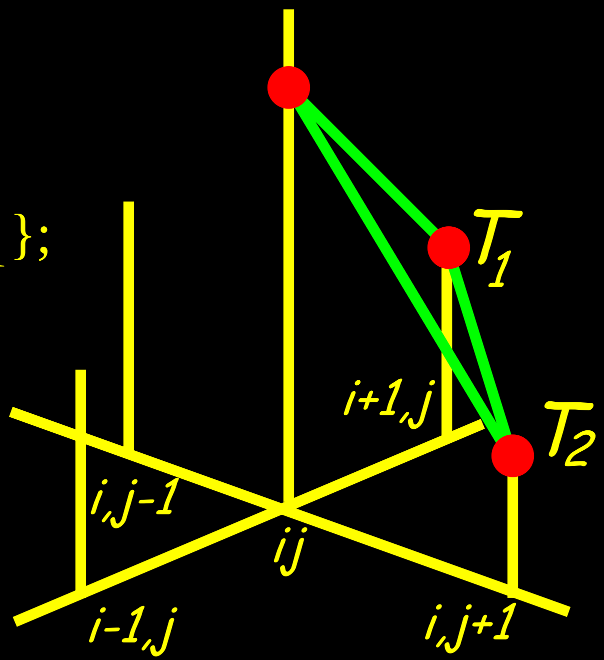

Fitting a tilted plane with gradient 1, and two values anchored at the relevant neighboring grid points.

Upwind approximation in 2D

$$\left(\max\{D^{-x}_{ij}T, -D^{+x}_{ij}T, 0\}\right)^2 + \left(\max\{D^{-y}_{ij}T, -D^{+y}_{ij}T, 0\}\right)^2 = 1,$$

where $$D^{-x}_{ij}T = \frac{(T_{ij} - T_{i-1,j})}{h}$$

$$\Leftrightarrow \quad \left(\max\{T_{ij} - \min\{T_{i-1,j}, T_{i+1,j}\}, 0\}\right)^2 + \left(\max\{T_{ij} - \min\{T_{i,j-1}, T_{i,j+1}\}, 0\}\right)^2 = h^2,$$

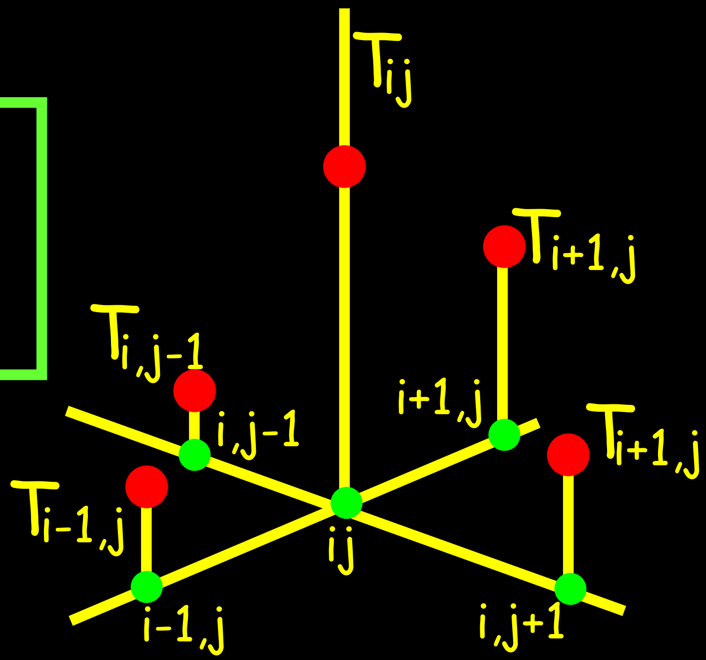

Computational Complexity

- *T* is systematically constructed from smaller to larger *T* values.
- Update of a heap element is *O(log N)*.
- Thus, upper bound of the total is *O(N log N)*.

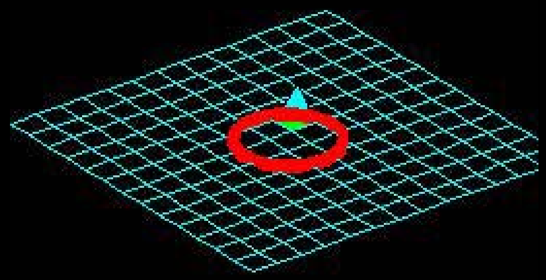

www.math.berkeley.edu/~sethian

Edge Integration

Cohen-Kimmel, IJCV, 1997.

Solve the 2D Eikonal equation

$$(D_x T)^2 + (D_y T)^2 = F_{ij}^2$$

given *T(p)=0*

Minimal geodesic w.r.t.

$$ds^2 = F^2(I)\left(dx^2 + dy^2\right)$$

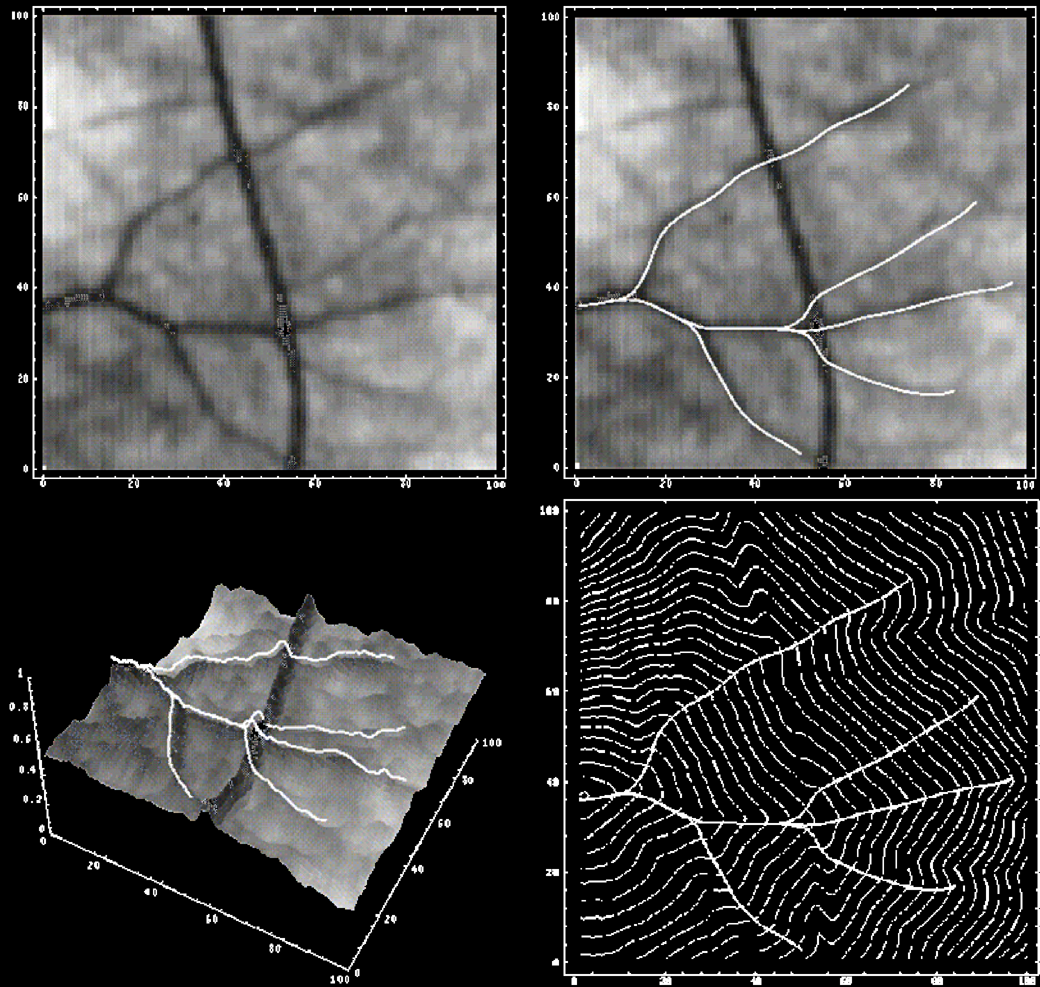

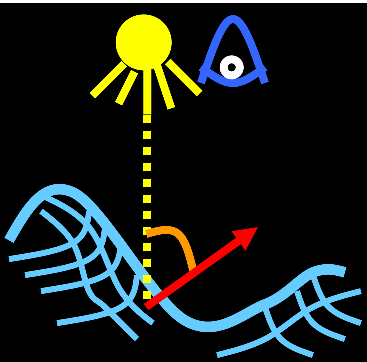

Shape from Shading

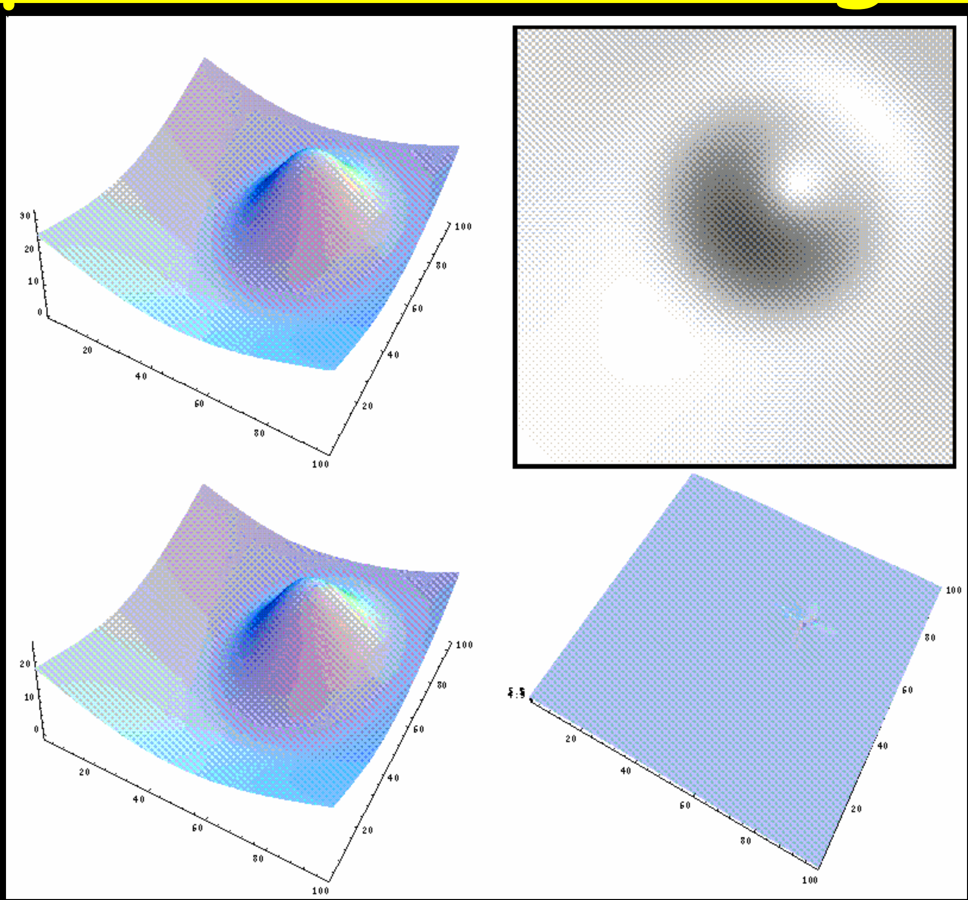

Rouy-Tourin SIAM-NU 1992,
Kimmel-Bruckstein CVIU 1994,
Kimmel-Sethian JMIV 2001.

Solve the 2D Eikonal equation

$$(D_x T)^2 + (D_y T)^2 = F_{ij}^2$$

where $F^2 = (1 - I^2)/I^2$

Minimal geodesic w.r.t.

$$ds^2 = F^2(I)\left(dx^2 + dy^2\right)$$

Path Planning 3 DOF

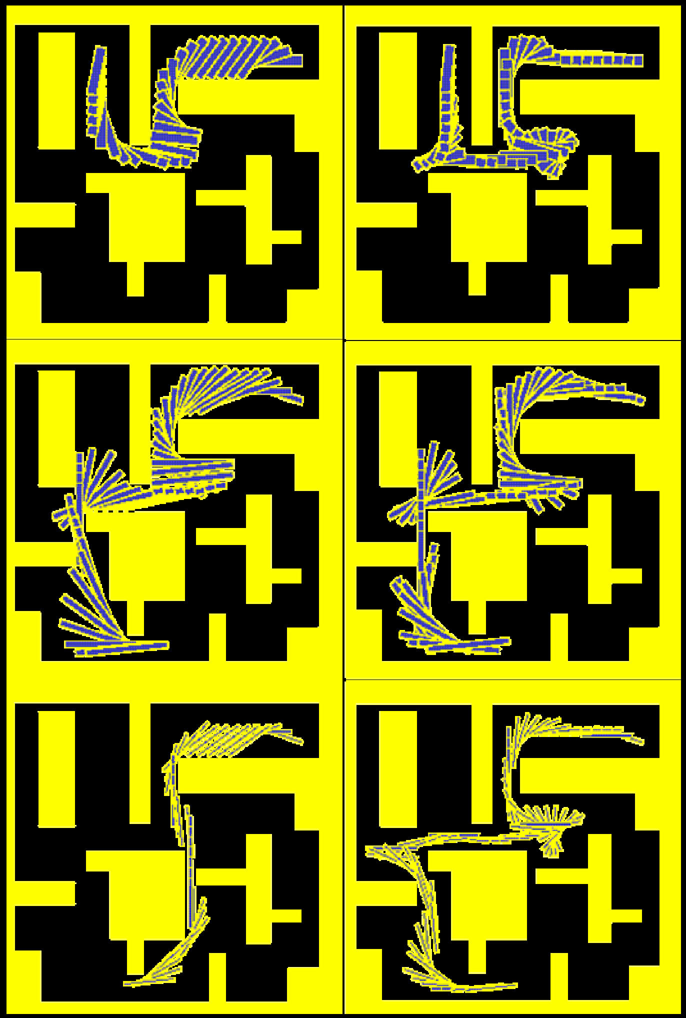

$$(D_x T)^2 + (D_y T)^2 + (D_\varphi T)^2 = F_{ijk}^2$$

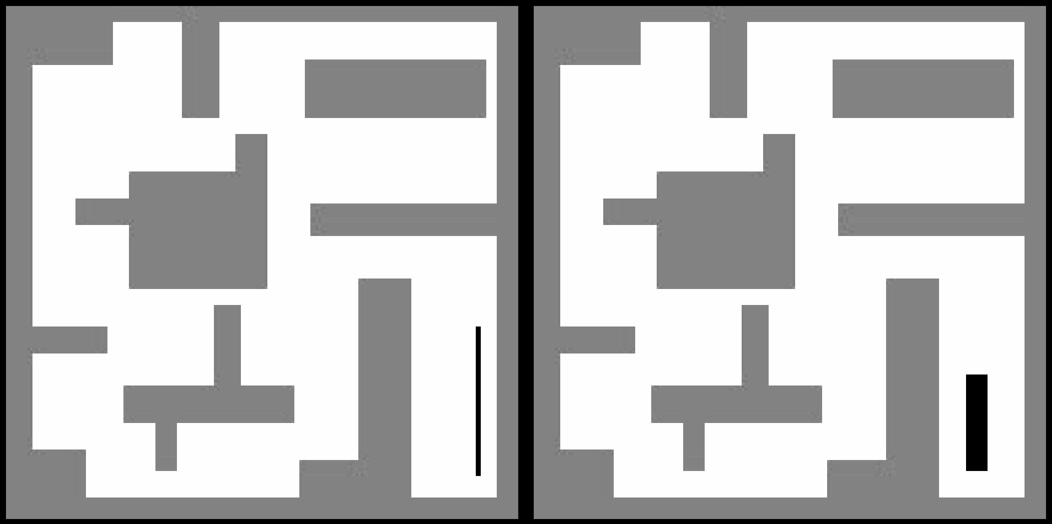

$$ds^2 = F^2(x, y, \varphi)\left(dx^2 + dy^2 + d\varphi^2\right)$$

Path Planning 4 DOF

Solve the Eikonal Eq. in 4D

$$\left(D_{\varphi_1} T\right)^2 + \left(D_{\varphi_2} T\right)^2 + \left(D_{\varphi_3} T\right)^2 + \left(D_{\varphi_4} T\right)^2 = F_{ijkl}^2$$

Minimal geodesic w.r.t.

$$ds^2 = F^2(\varphi_1, \varphi_2, \varphi_3, \varphi_4)\left(d\varphi_1^2 + d\varphi_2^2 + d\varphi_3^2 + d\varphi_4^2\right)$$

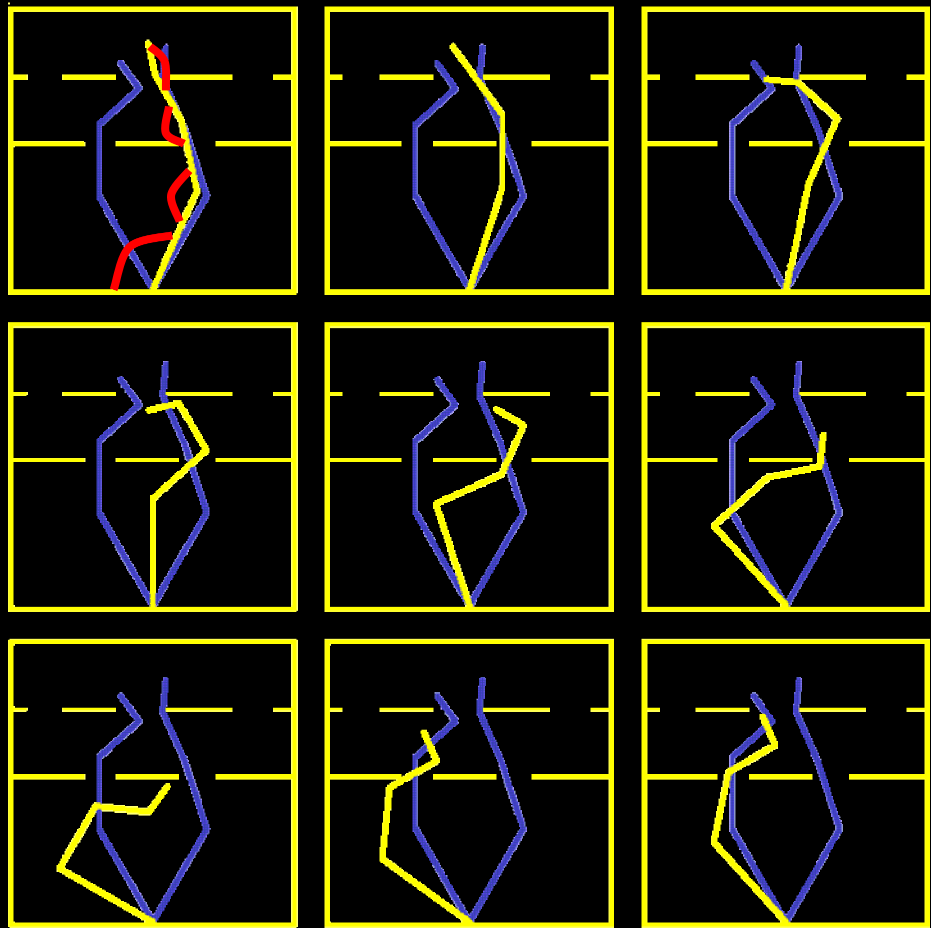

Update Acute Angle

Given *ABC*, update *C*.
Consistency and monotonicity:
Update only `from within the triangle' *h* in *ABC*
Find *t=EC* that satisfies the gradient approximation

(t-u)/h= 1.

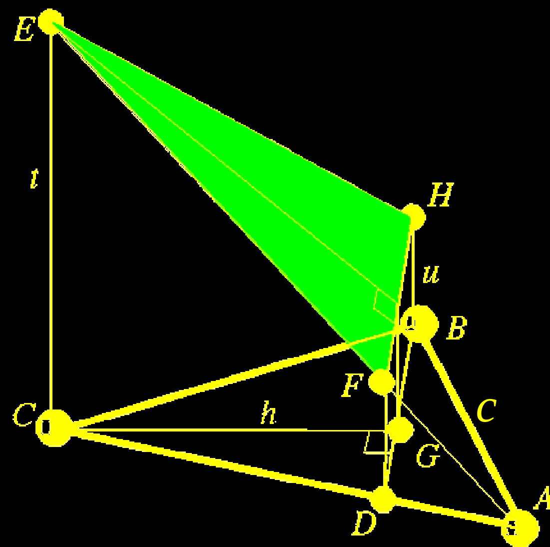

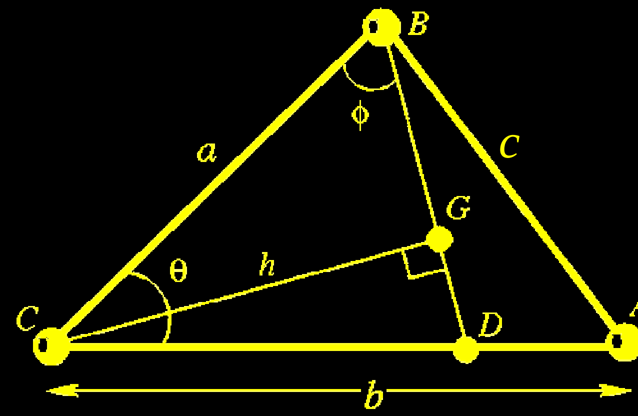

Update Procedure

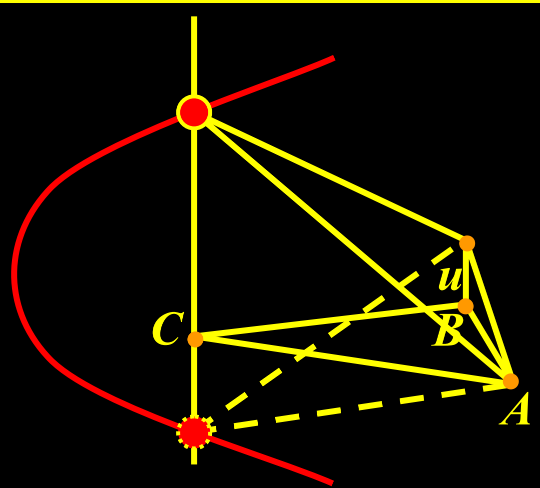

We end up with:

$$c^2 t^2 + 2bu(a^2 \cos\theta - b)t + b^2\left(u^2 - (a\sin\theta)^2\right) = 0$$

$$\alpha t^2 + \beta t + \chi = 0$$

t must satisfy *u<t*, and *h* in *ABC*.

The update procedure is

IF *(u<t)* AND *(a cos θ < b(t-u)/t < a/cosθ)*

THEN *T(C) = min {T(C),t+T(A)};*

ELSE *T(C)= min {T(C),b+T(A),a+T(B)}.*

Minimal Geodesics

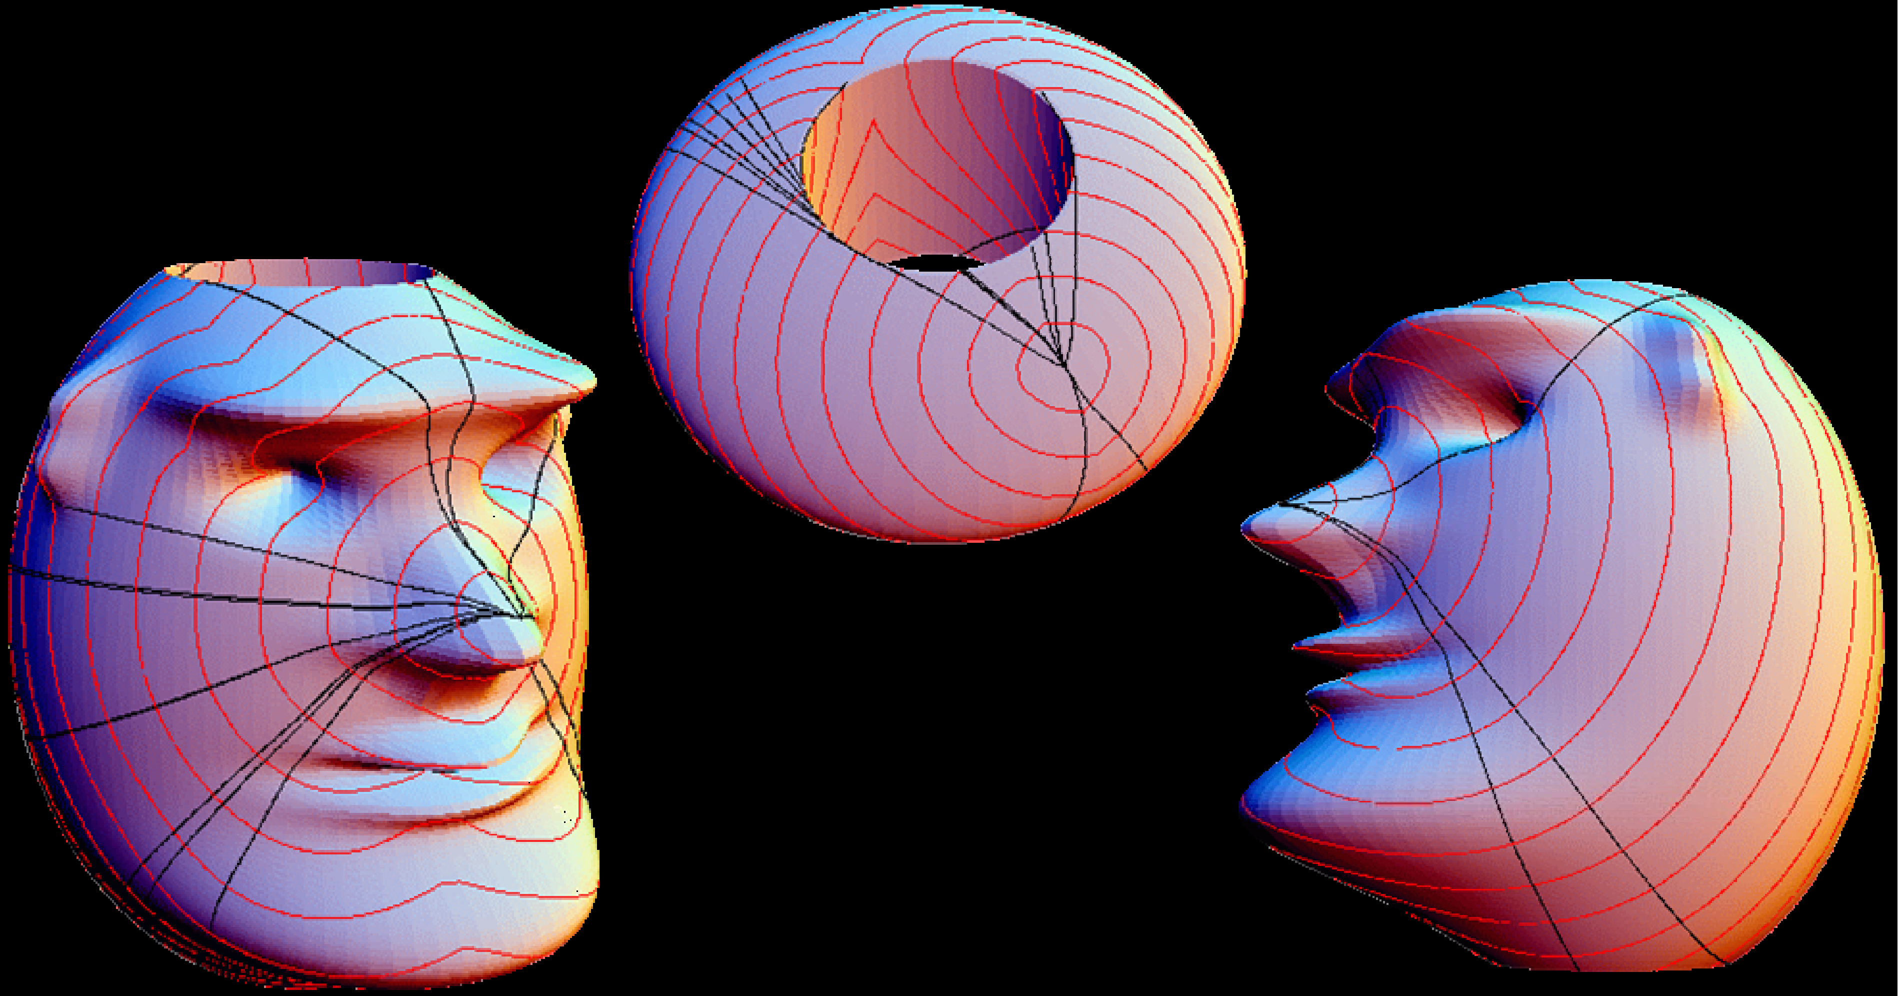

Kimmel and Sethian, PONAS 1998

More Applications

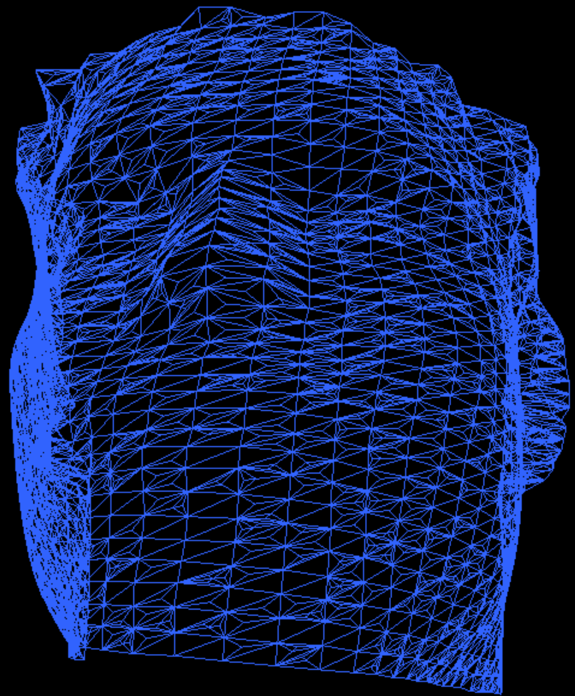

re-triangulation

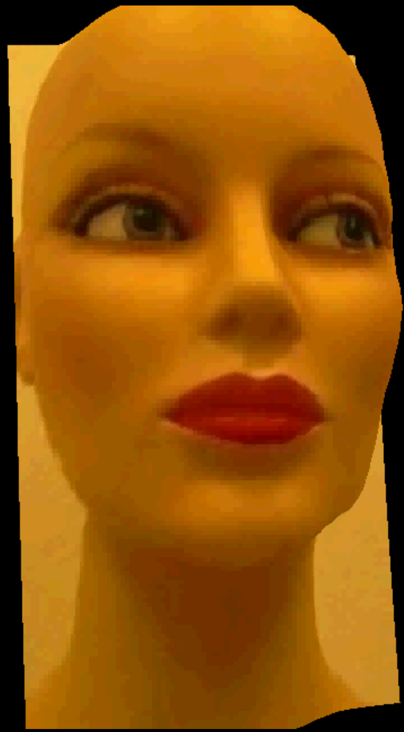

semi-manual segmentation

halftoning in 3D

Adi, Kimmel 2002

Released: Oct. 2003
Publisher: Springer

ISBN: 0387955623

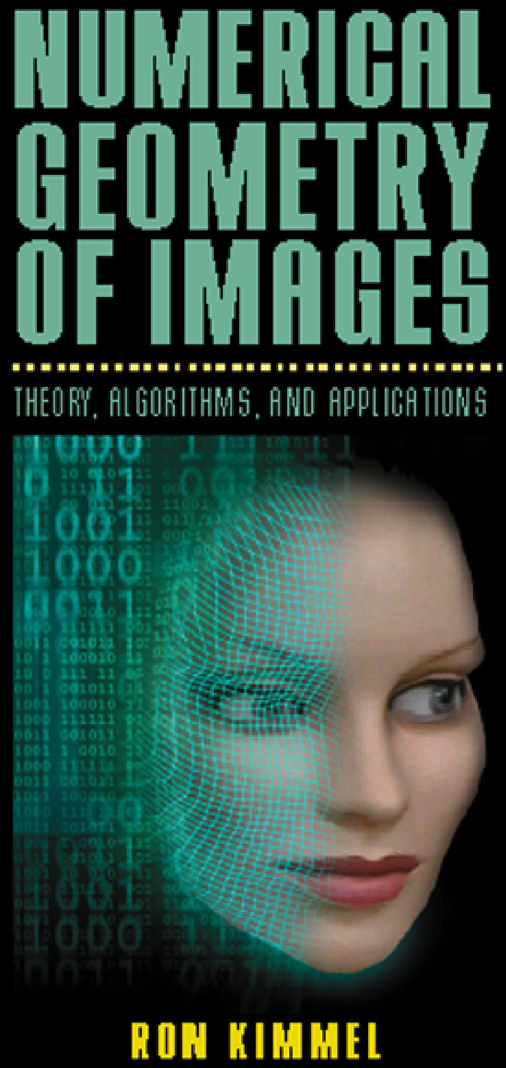

Tomorrow:
Flat embedding and isometric invariant signatures

Applications:
Lip Reading
Texture mapping
3D Face recognition

Front cover rendered by A&M Bronstein

Computer Science Department

Technion-Israel Institute of Technology

Matching Isometric Manifolds by Flat Embedding

Ron Kimmel

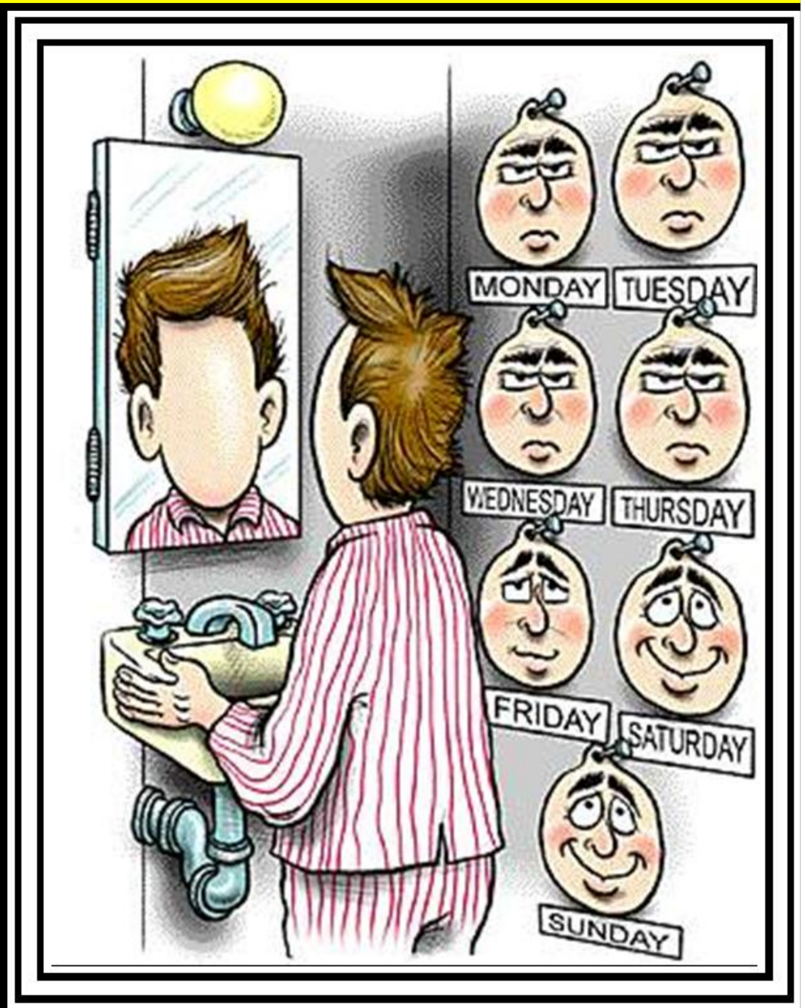

Geometric Image Processing Lab

Students

Michal Aharon
CS Technion

Ilya Blayvas
CS Technion

Alexander Bronstein
EE-CS Technion

Michael Bronstein
CS Technion

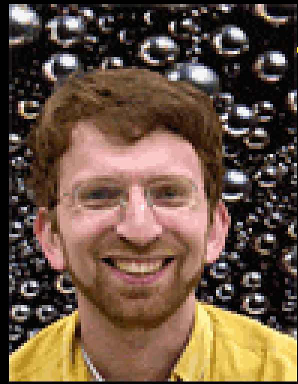

Alexander Brook
Math Technion

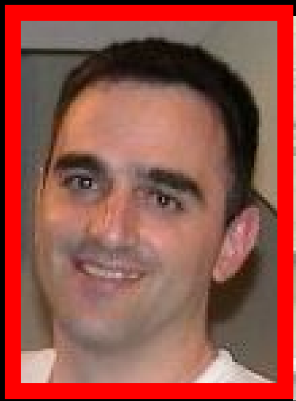

Asi Elad
IO-Image

Roman Goldenberg
CS Technion

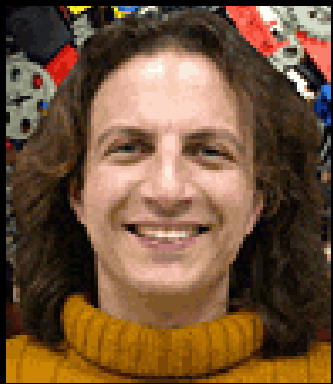

Eyal Gordon
CS Technion

Michal Holzman Gazit
EE-CS Technion

Alon Spira
CS Technion

Gil Zigelman
3DV Systems

Colleagues

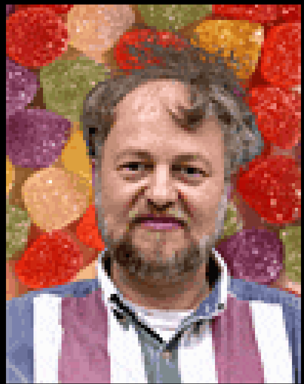

Alfred M. Bruckstein
CS Technion

Laurent Cohen
Paris Dauphine

Vicent Caselles
Math. Barcelona

Nahum Kiryati
Tel Aviv Univ.

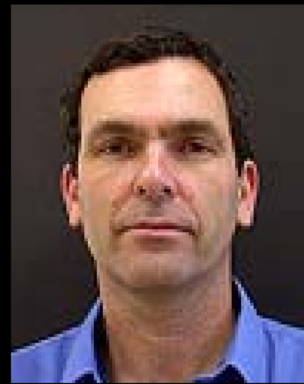

Ehud Rivlin
CS Technion

Michael Rudszky
CS Technion

Guillermo Sapiro
EE Minnesota

James Sethian
UC Berkeley

Nir Sochen
Tel Aviv Univ.

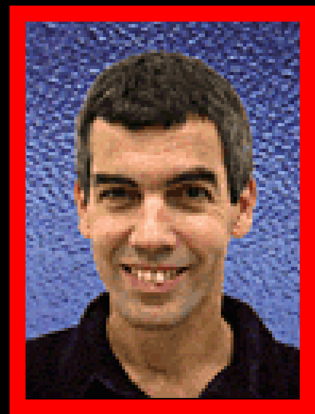

Irad Yavneh
CS Technion

Expressions are ~isometries

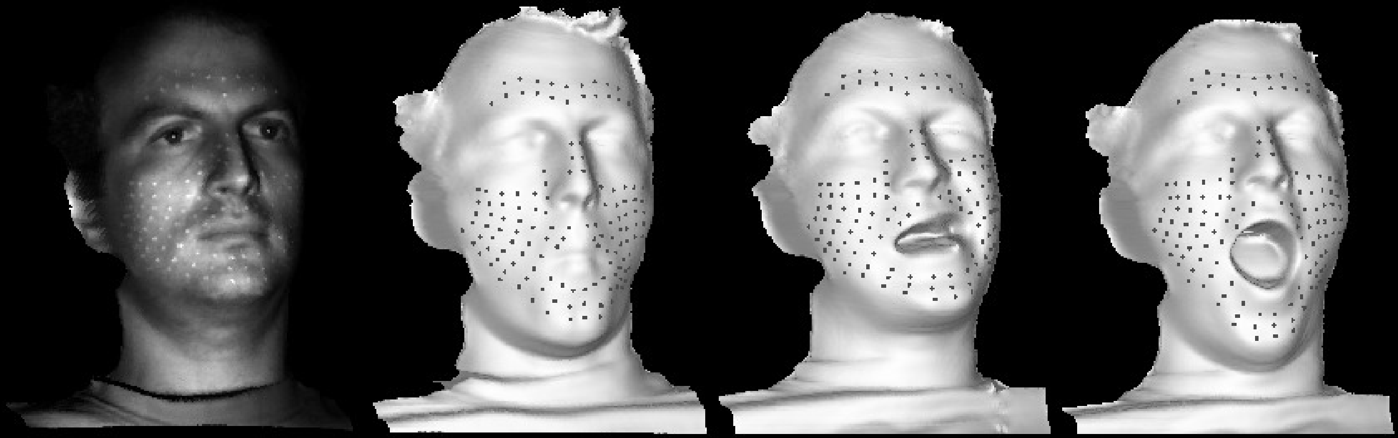

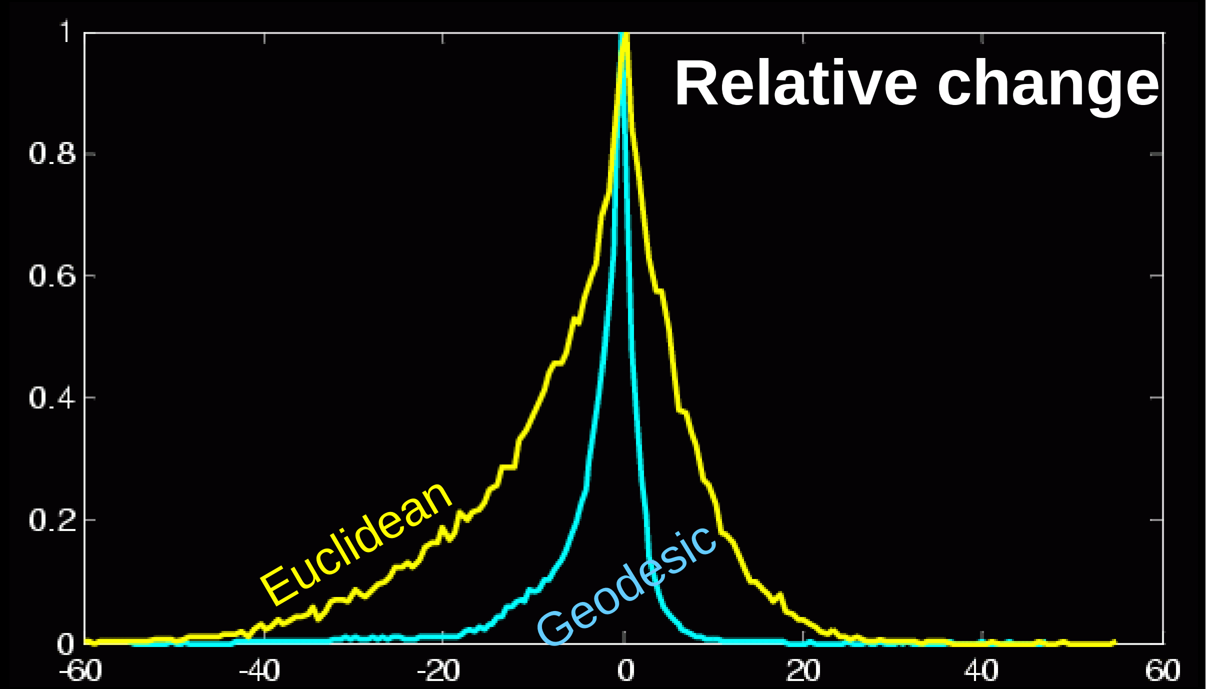

Bronstein Bronstein Kimmel *2005*

Open mouth & isometric expressions

On MDS,
Lip Reading,
Texture Mapping,
and
Isometric Signatures

How it all started... (my story)

1989	Schwartz Shaw Wolfson	TPAMI	2D flattening
1995	Kimmel Amir Bruckstein	TPAMI	geodesic distances
1998	Kimmel Sethian	PONAS	geodesic distances
2000	Zigelman Kimmel Kiryati	TVCG	texture mapping
2001	A. Elad Kimmel	TPAMI	isometric signatures
2003	Bronstein^2 Kimmel	AVBPA	3D face recognition
2003	Spira Kimmel	scale space	geodesic distances
2004	Bronstein^2 Spira Kimmel	ECCV	3D recognition without 3D
2005	Bronstein^2 Kimmel	scale space	S3 embedding
2005	Bronstein^2 Kimmel	IJCV	3D recognition
2005	Bronstein^2 Kimmel	2Bsub.	Isometry/S2 embedding

Multidimensional scaling

Multidimensional scaling

- MDS is a family of methods that map similarity measurements among objects, to points in a small dimensional Euclidean space.

- Enables to explore the geometric structure of the data.

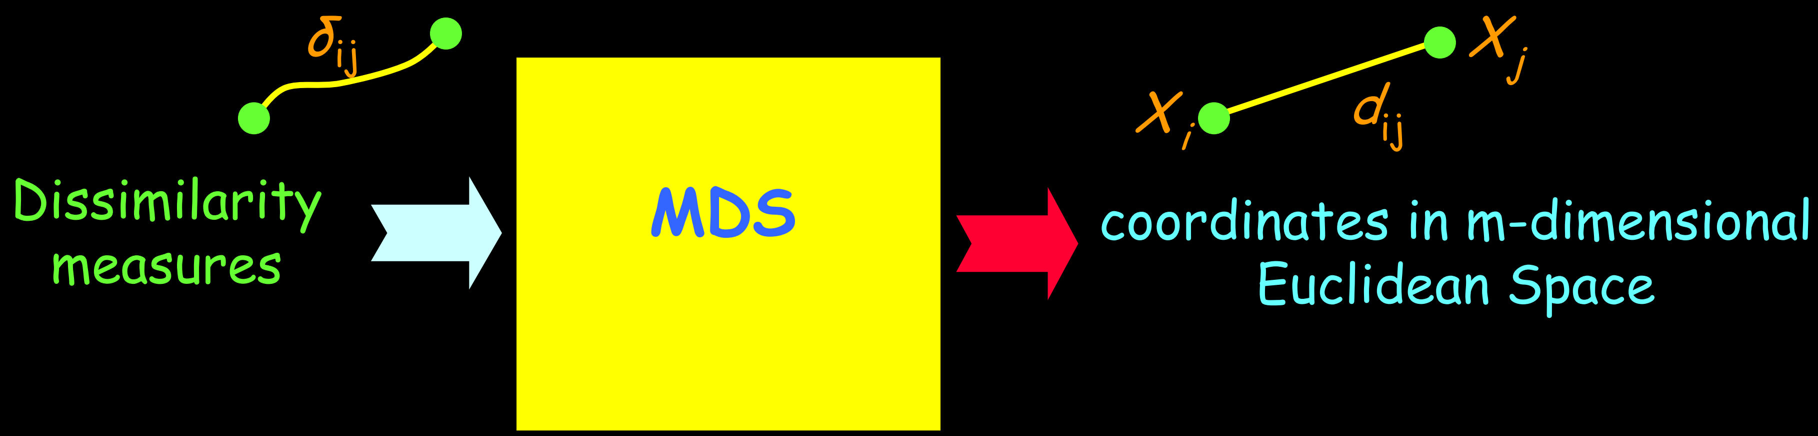

$$\textit{Stress} = \frac{\sum w_{ij} \left(\delta_{ij} - d_{ij}(X)\right)^2}{\sum w_{ij}\, \delta_{ij}^2}$$

A simple example

Rotation

Reflection

1	2	3	4	5	6	7	8	9	10
12.7	0	32.8	29.4	14.6	8.8	4.5	6.4	16.4	13.8
4.1	13.9	13.2	14.9	9.0	8.2	16.6	19.6	25.4	0

Young et al. 1930,
Torgerson&Gower 1952,1958,1966.

Classical scaling

Given n points in R^k, denote $\mathbf{p}_i = [x_i^1, x_i^2, \dots, x_i^k]^{\mathbf{T}}$

Define coordinates vector $\mathbf{P} = [\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_n]^{\mathbf{T}}$

The Euclidean distance between 2 points:

$$d_{ij}^2 = \left|\mathbf{p}_i - \mathbf{p}_j\right|^2 = \left|\mathbf{p}_i\right|^2 - 2\mathbf{p}_i\mathbf{p}_j^{\mathbf{T}} + \left|\mathbf{p}_j\right|^2$$

$$\mathbf{Q} = \begin{bmatrix} |\mathbf{p}_1|^2 & |\mathbf{p}_1|^2 & \cdots & |\mathbf{p}_1|^2 \\ |\mathbf{p}_2|^2 & |\mathbf{p}_2|^2 & \cdots & |\mathbf{p}_2|^2 \\ . & . & \cdots & . \\ |\mathbf{p}_n|^2 & |\mathbf{p}_n|^2 & \cdots & |\mathbf{p}_n|^2 \end{bmatrix}$$

$$\mathbf{D} = \mathbf{Q} - 2\mathbf{P}\mathbf{P}^{\mathbf{T}} + \mathbf{Q}^{\mathbf{T}}$$

$$\mathbf{JDJ} = \mathbf{J}(\mathbf{Q} + \mathbf{Q}^{\mathbf{T}} - 2\mathbf{P}\mathbf{P}^{\mathbf{T}})\mathbf{J} = \mathbf{0} + \mathbf{0} - 2\tilde{\mathbf{P}}\tilde{\mathbf{P}}^{\mathbf{T}}$$

$$\mathbf{J} = \mathbf{I} - \tfrac{1}{n}\mathbf{1}\mathbf{1}^T$$

$$-\tfrac{1}{2}\mathbf{JDJ} = \mathbf{U}\Lambda\mathbf{U}^{\mathbf{T}} \quad \Rightarrow \quad \tilde{\mathbf{P}}^{\mathbf{T}} = \Lambda^{\frac{1}{2}}\mathbf{U}^{\mathbf{T}}$$

Classical scaling

Matlab code
for 2D flattening

$$J = eye(n) - ones(n)./n;$$
$$B = -0.5 * J * D * J;$$
$$[U, L] = eigs(B, 2, 'LM');$$
$$newy = sqrt(L(1,1)).*U(:,1);$$
$$newx = sqrt(L(2,2)).*U(:,2);$$

Or a single command (Matlab 6.5)

$$Y = cmdscale(D);$$
$$newx = Y(1,:);$$
$$newy = Y(2,:);$$

Zigelman, Kimmel, Kiryati IEEE TVCG 2002

Analyzing and Synthesizing Lips Movements

Lip-reading is a difficult task.

Goal: Recognize a limited set of words.

Previous Lip synthesis:
video-rewrite (Bregler)
eigenfacemasks (Van Gool).

Previous Lip Reading:
HMM (Bregler, 1998) (the bartender problem),
PCA (Li et al. 1997) (eigensequences),
NN and SVM (Bregler, Duchnowski).

Michal Aharon, Kimmel 2004

Main Idea

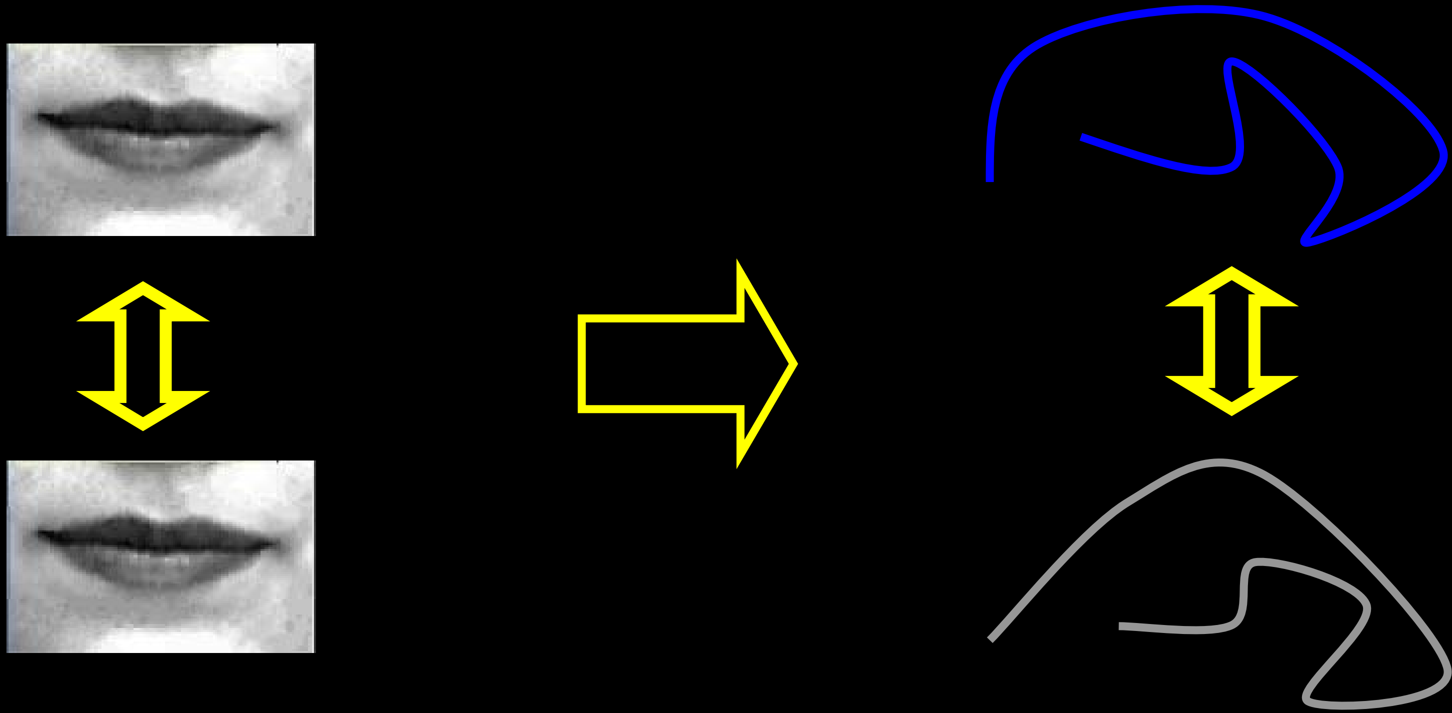

Michal Aharon, Kimmel 2004

Main Idea

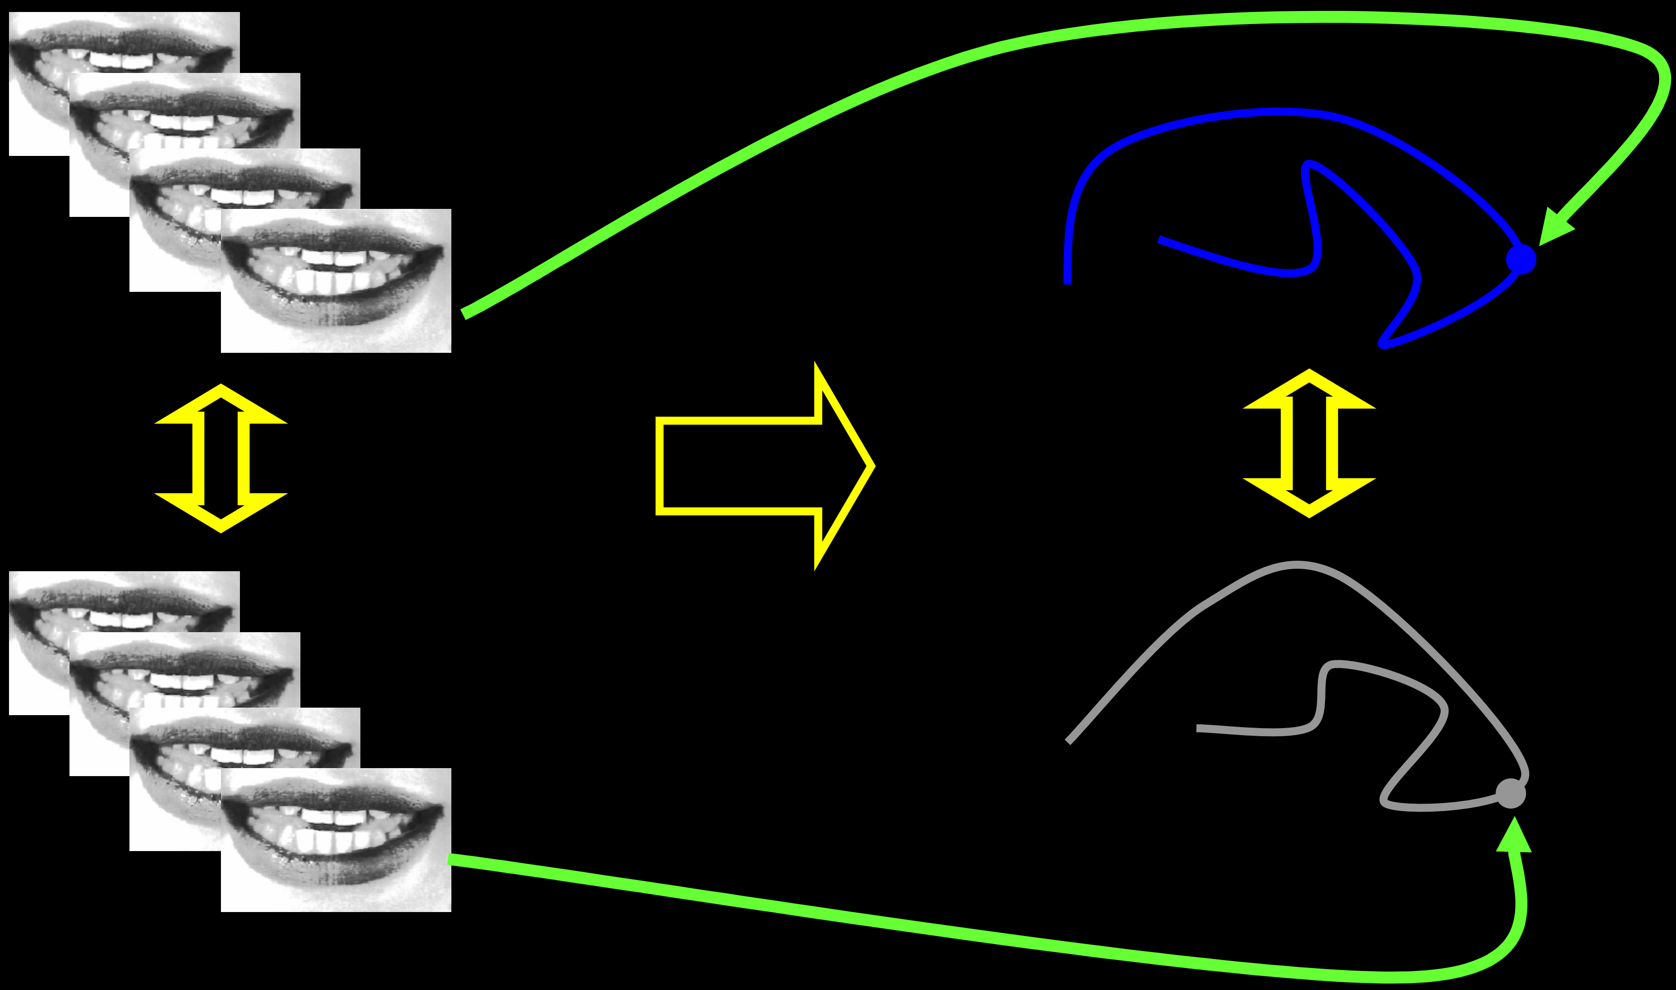

Michal Aharon, Kimmel 2004

Framework

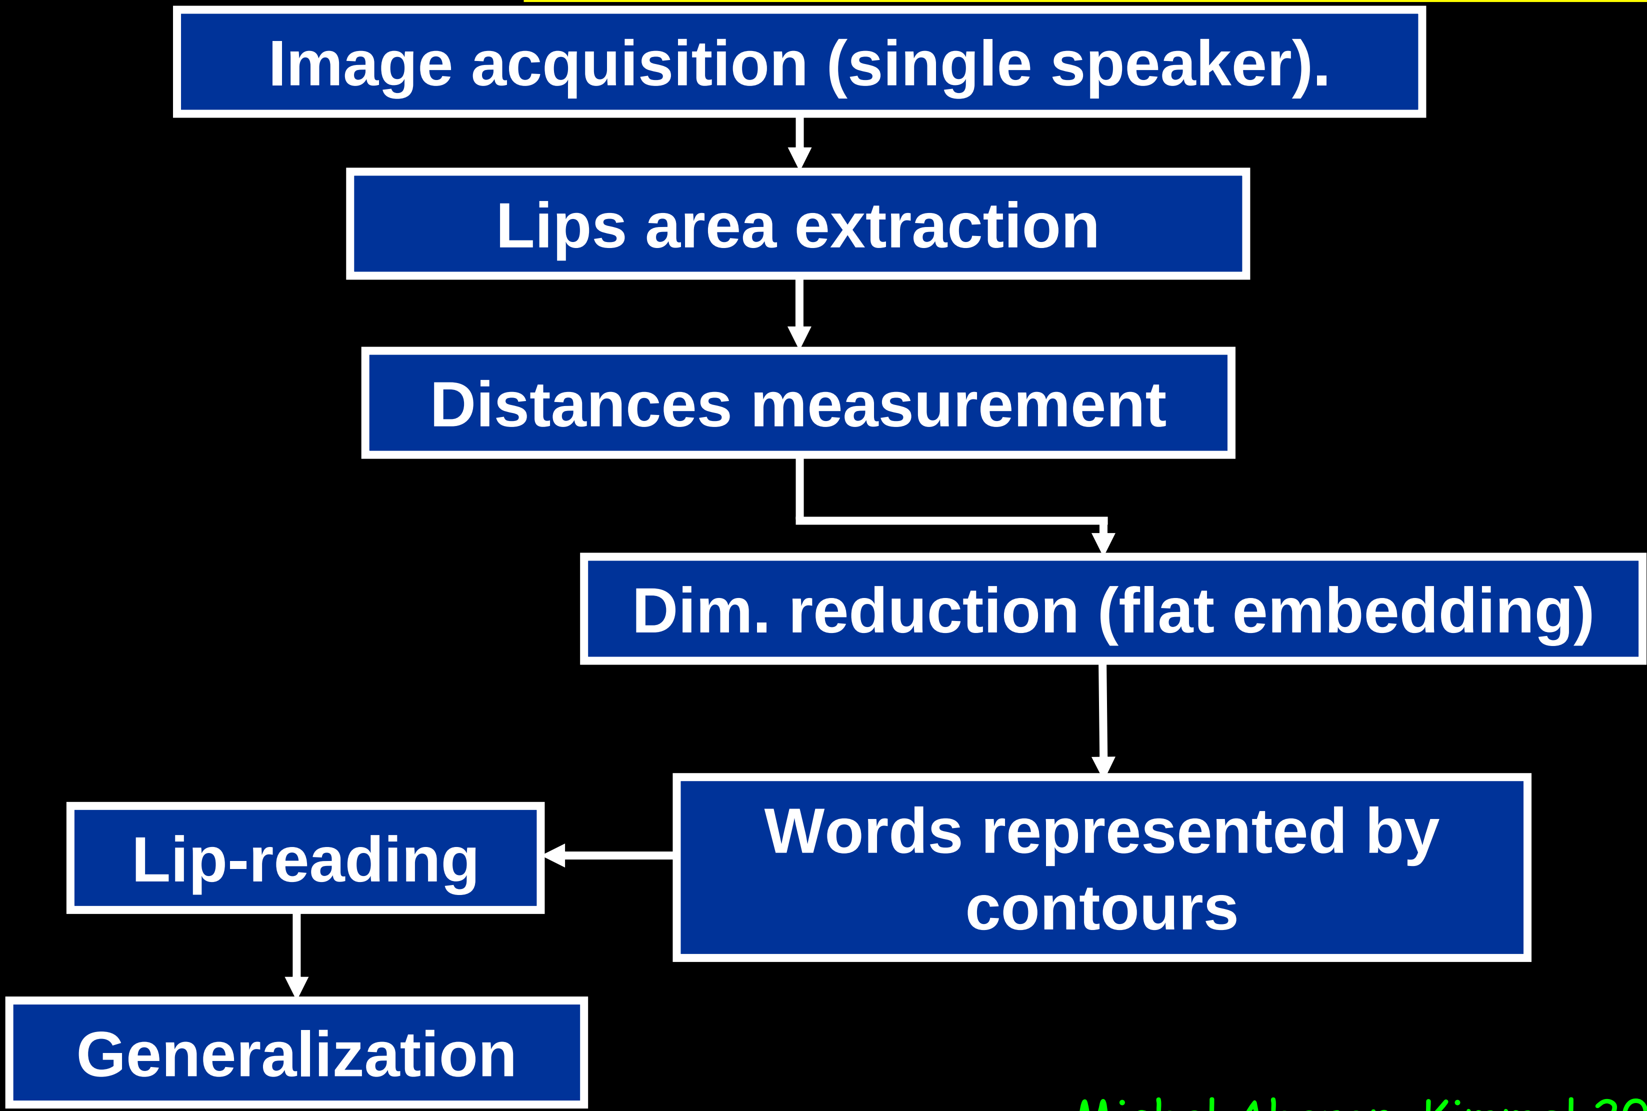

Michal Aharon, Kimmel 2004

Framework

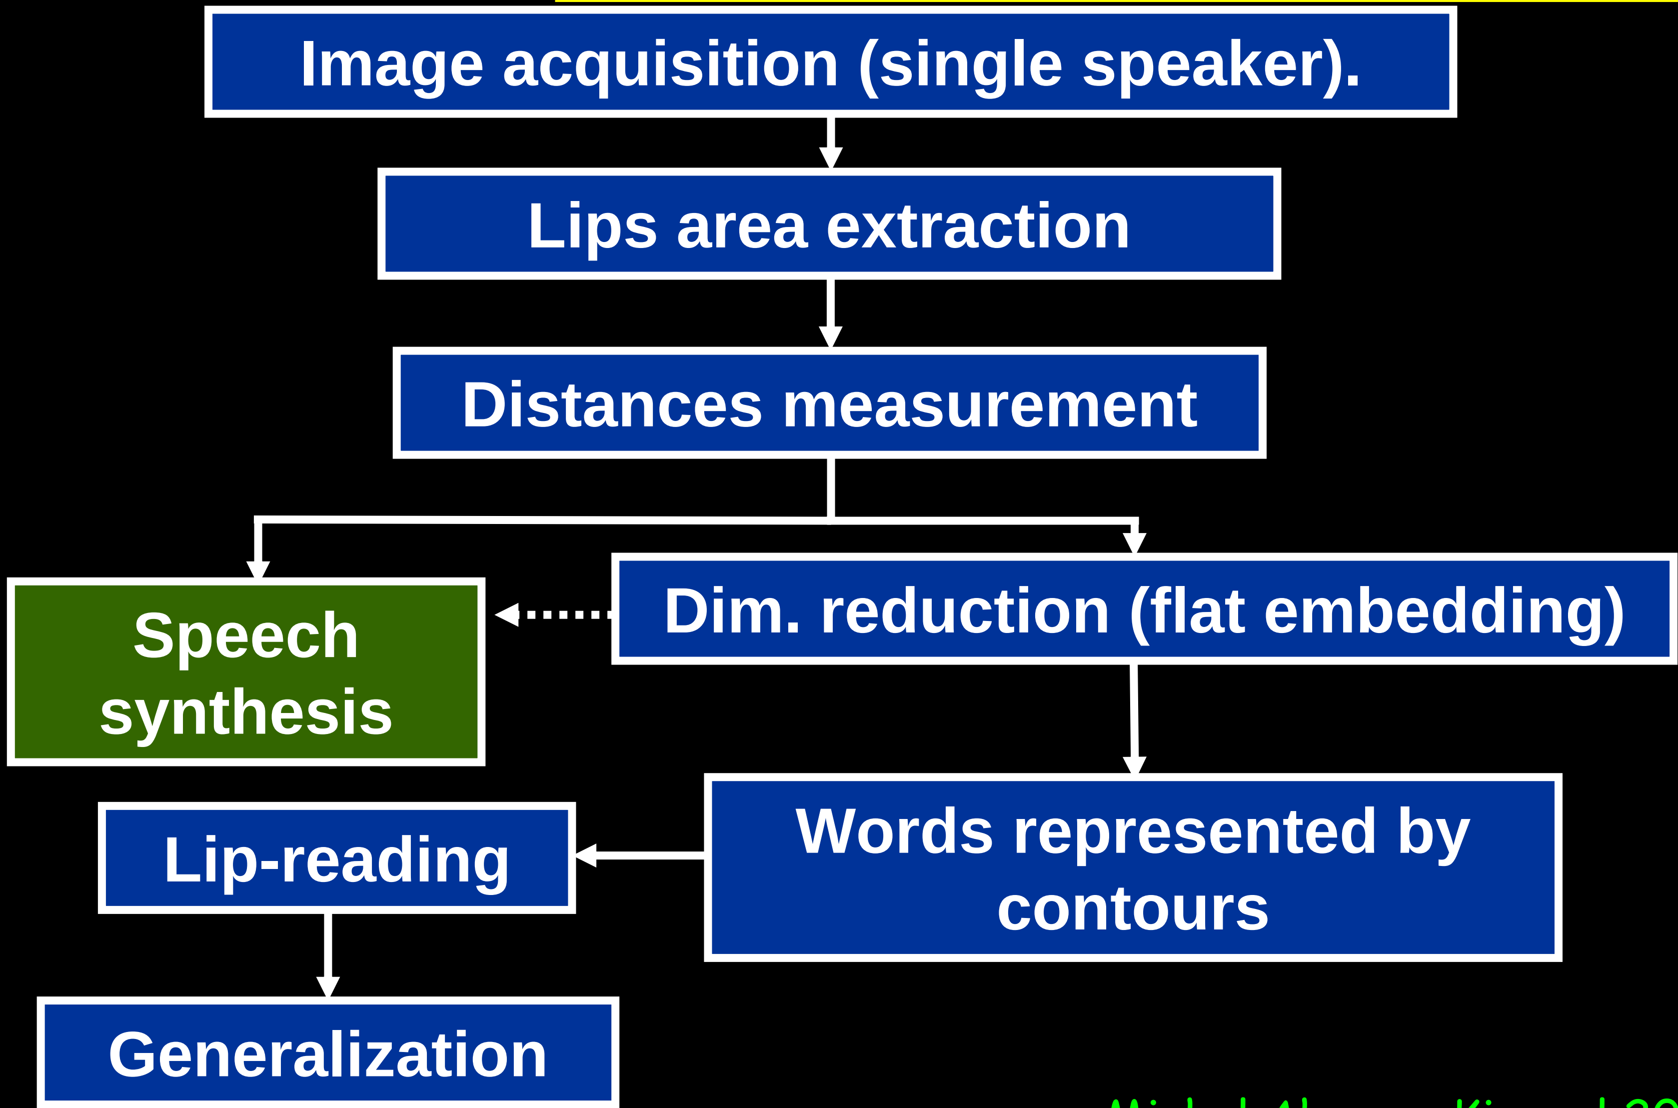

Michal Aharon, Kimmel 2004

Image acquisition (single speaker).

Lips Area extraction

- nose used for alignment
- images aligned using affine model
- The mouth area section is extracted as gray level images.

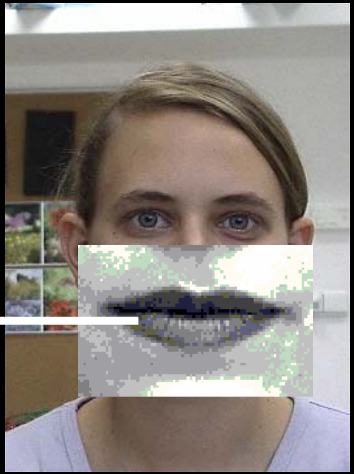

Distances measurement

- Variation on Jacobs, Belhumeur, Basri (1998)

$$E(I,J)=\iint I\cdot J\cdot\left|\nabla\left(\frac{I}{J}\right)\right|\cdot\left|\nabla\left(\frac{J}{I}\right)\right|dxdy$$

- ~ Invariant to changes in illumination.
- Symmetric consideration.
- Avoid singularities in dark areas.

Comparing JBB to other measures

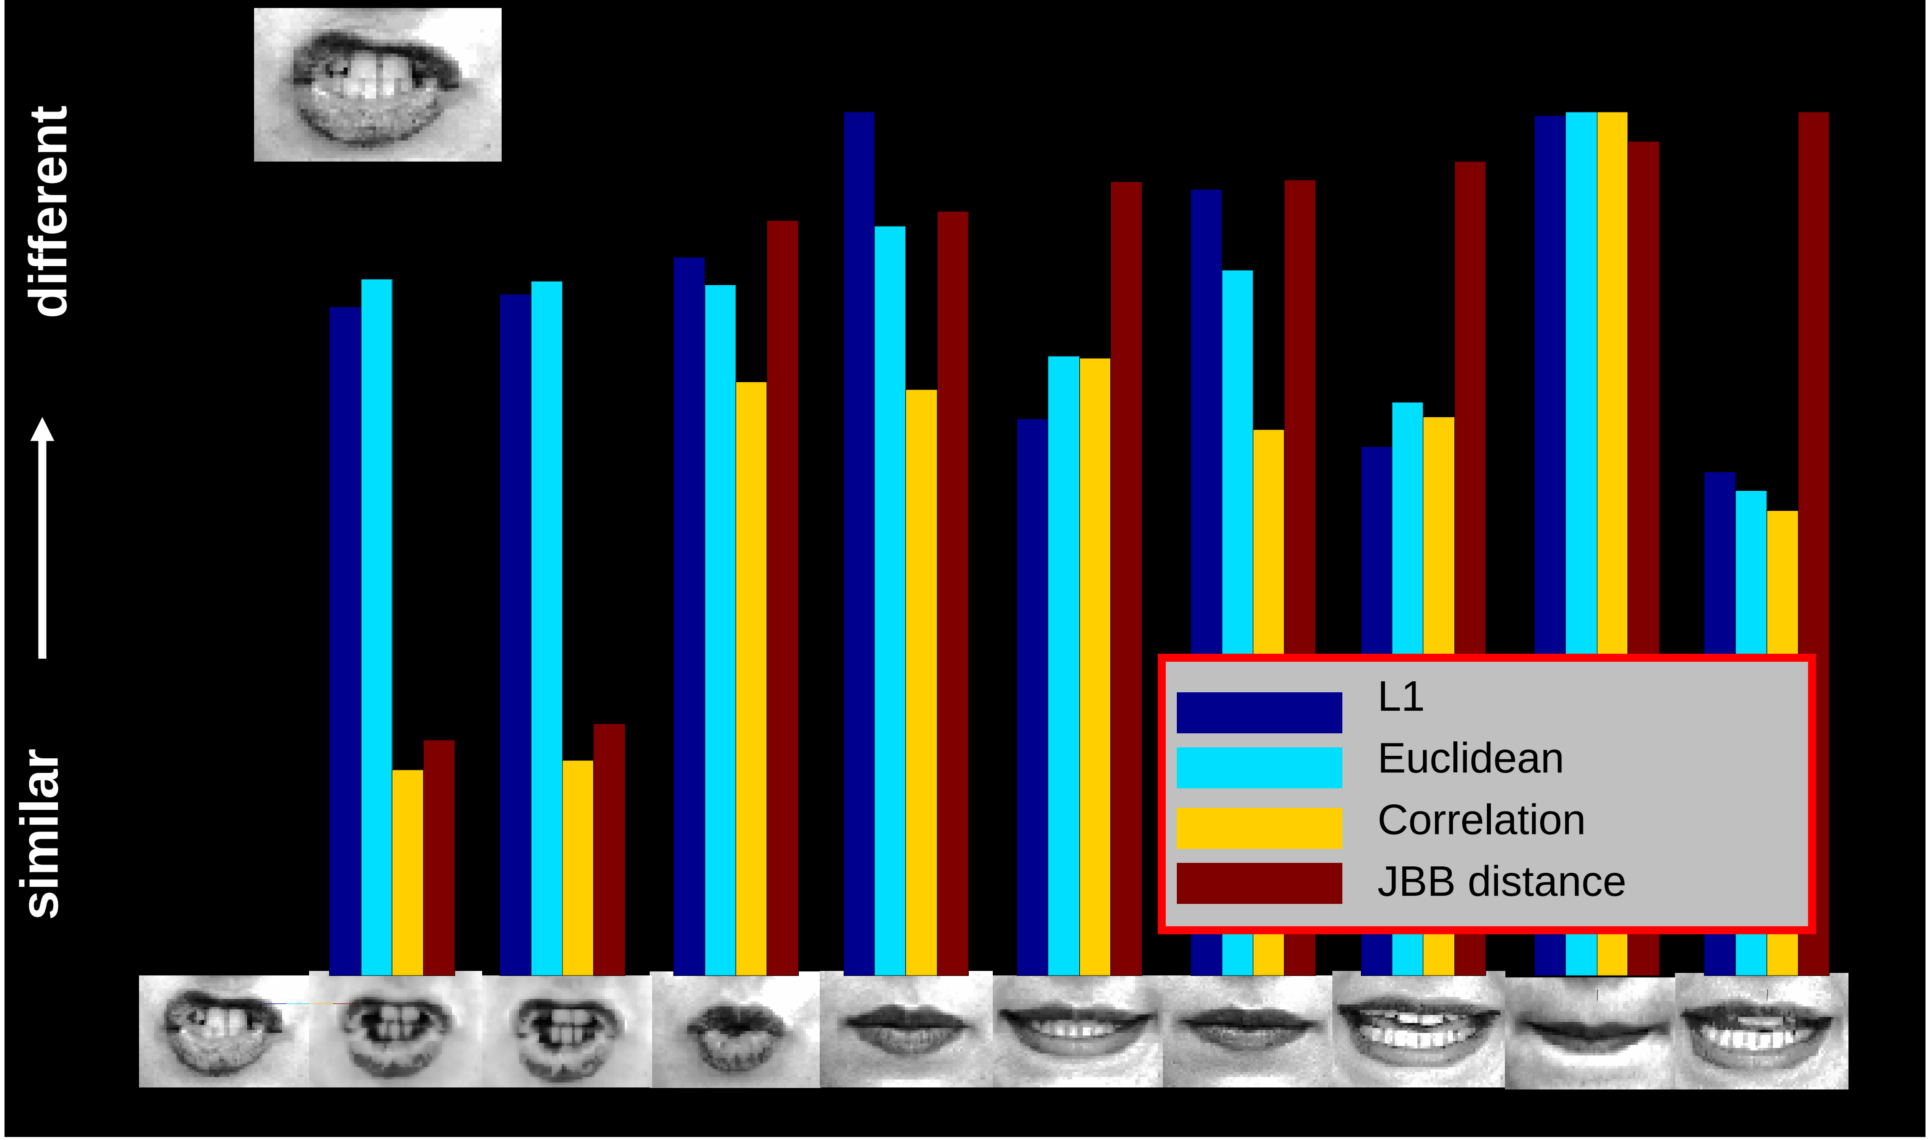

Lip representation space

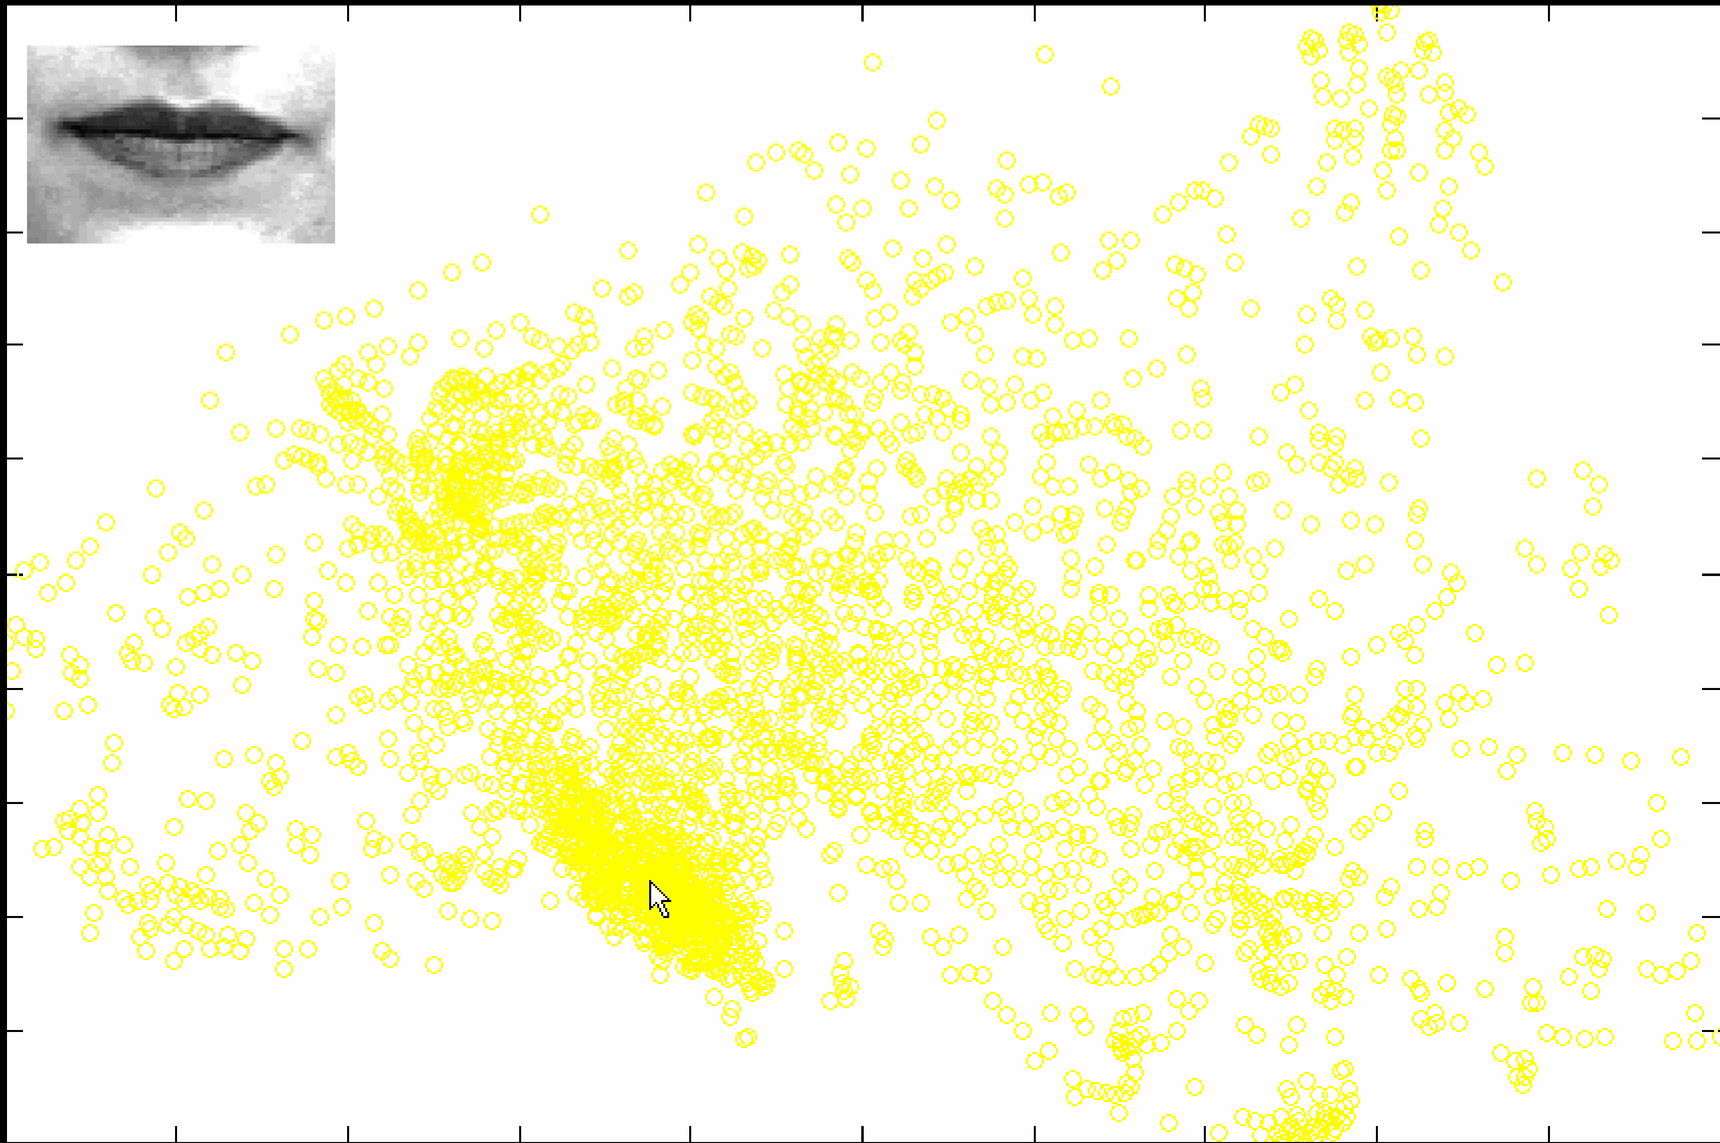

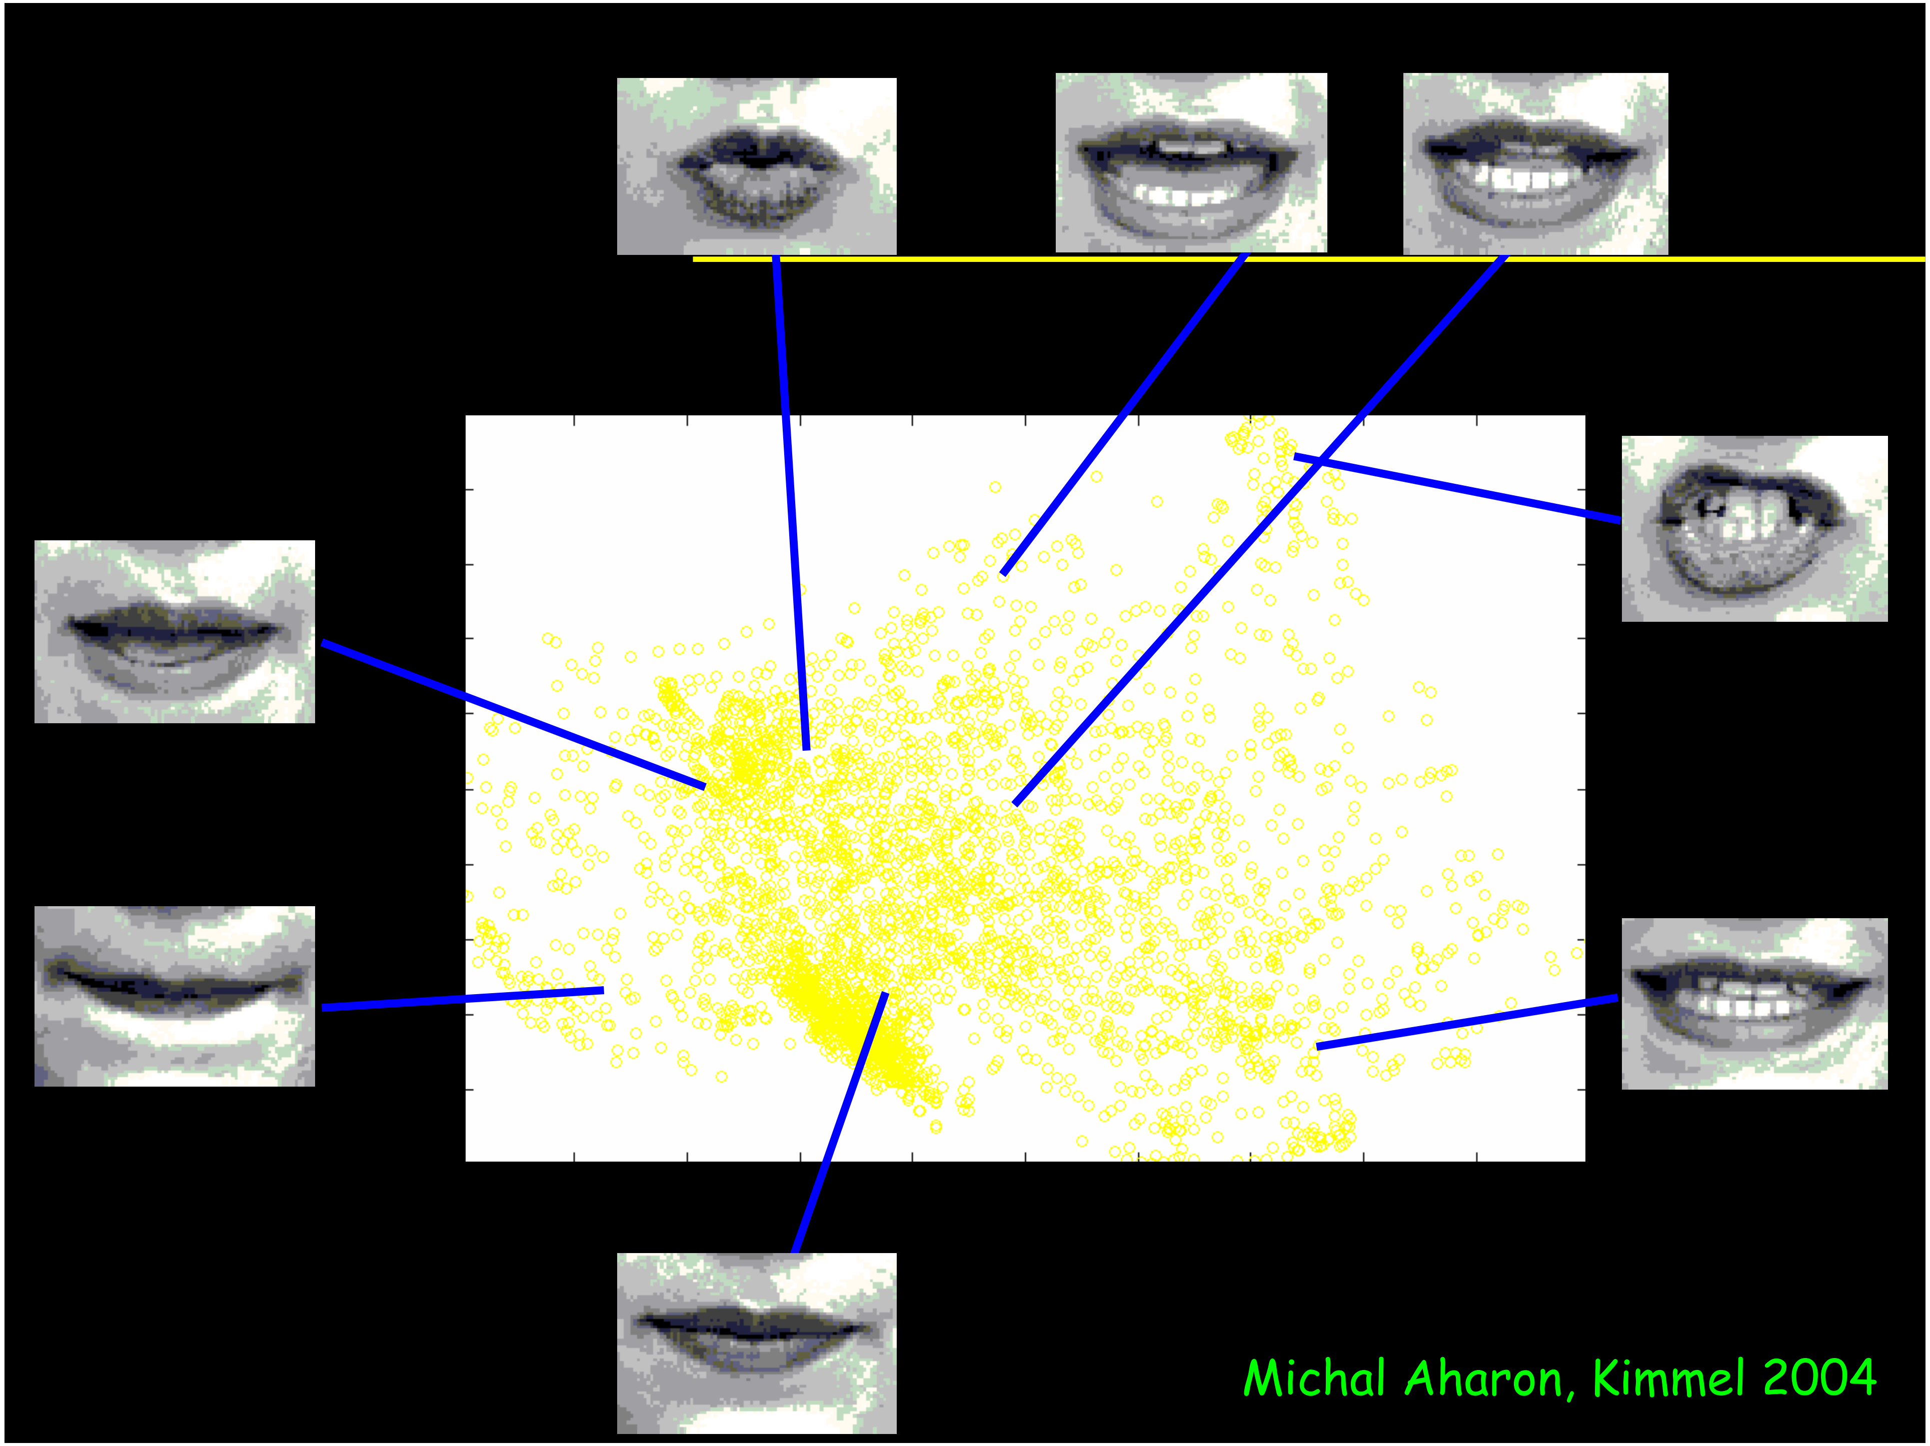
Michal Aharon, Kimmel 2004

Lip representation space

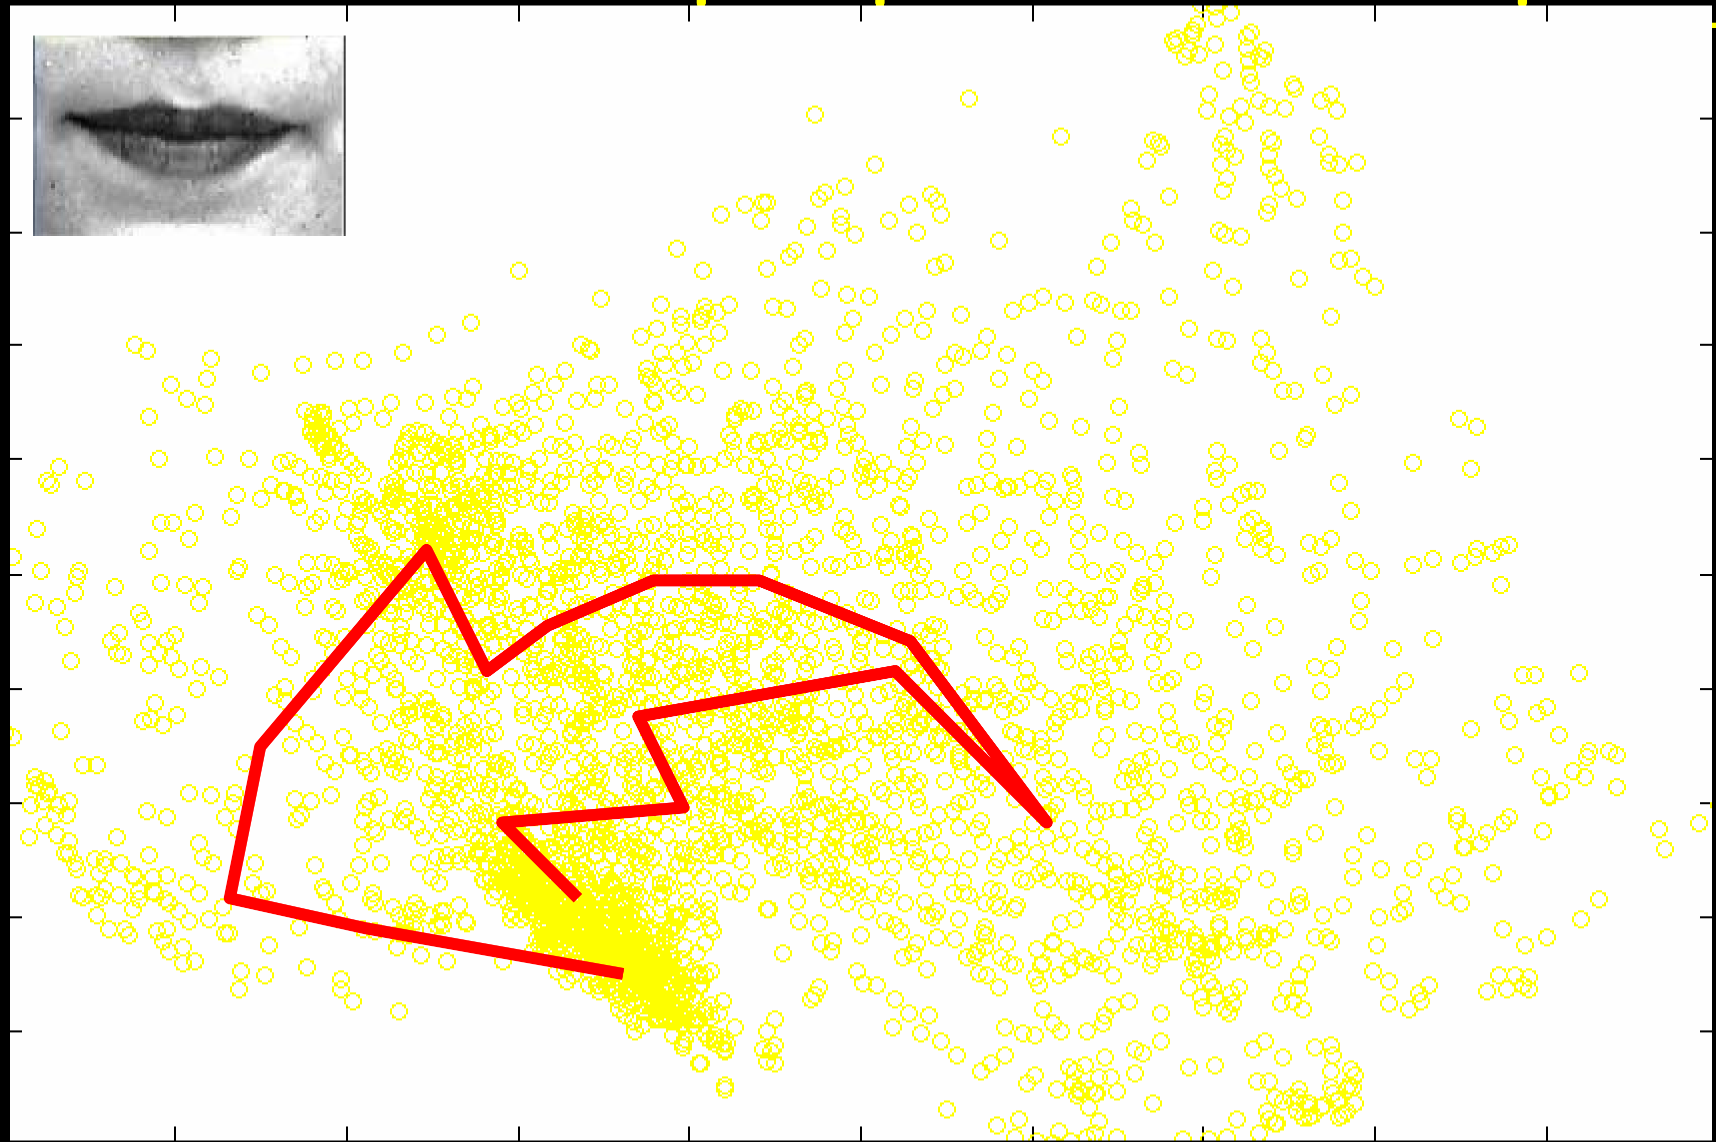

Speech Synthesis

- Linear interpolation in intensity space is not natural.
- The transition between syllables should be embedded in the lip representation space.

Michal Aharon, Kimmel 2004

Visual Articulation Signature

- *Visual articulation signature* (VAS) – the series of mouth configurations that occur in order for a sound to be vocalized.

Synthesis - SHU+KI

Synthesis - SHU+KI

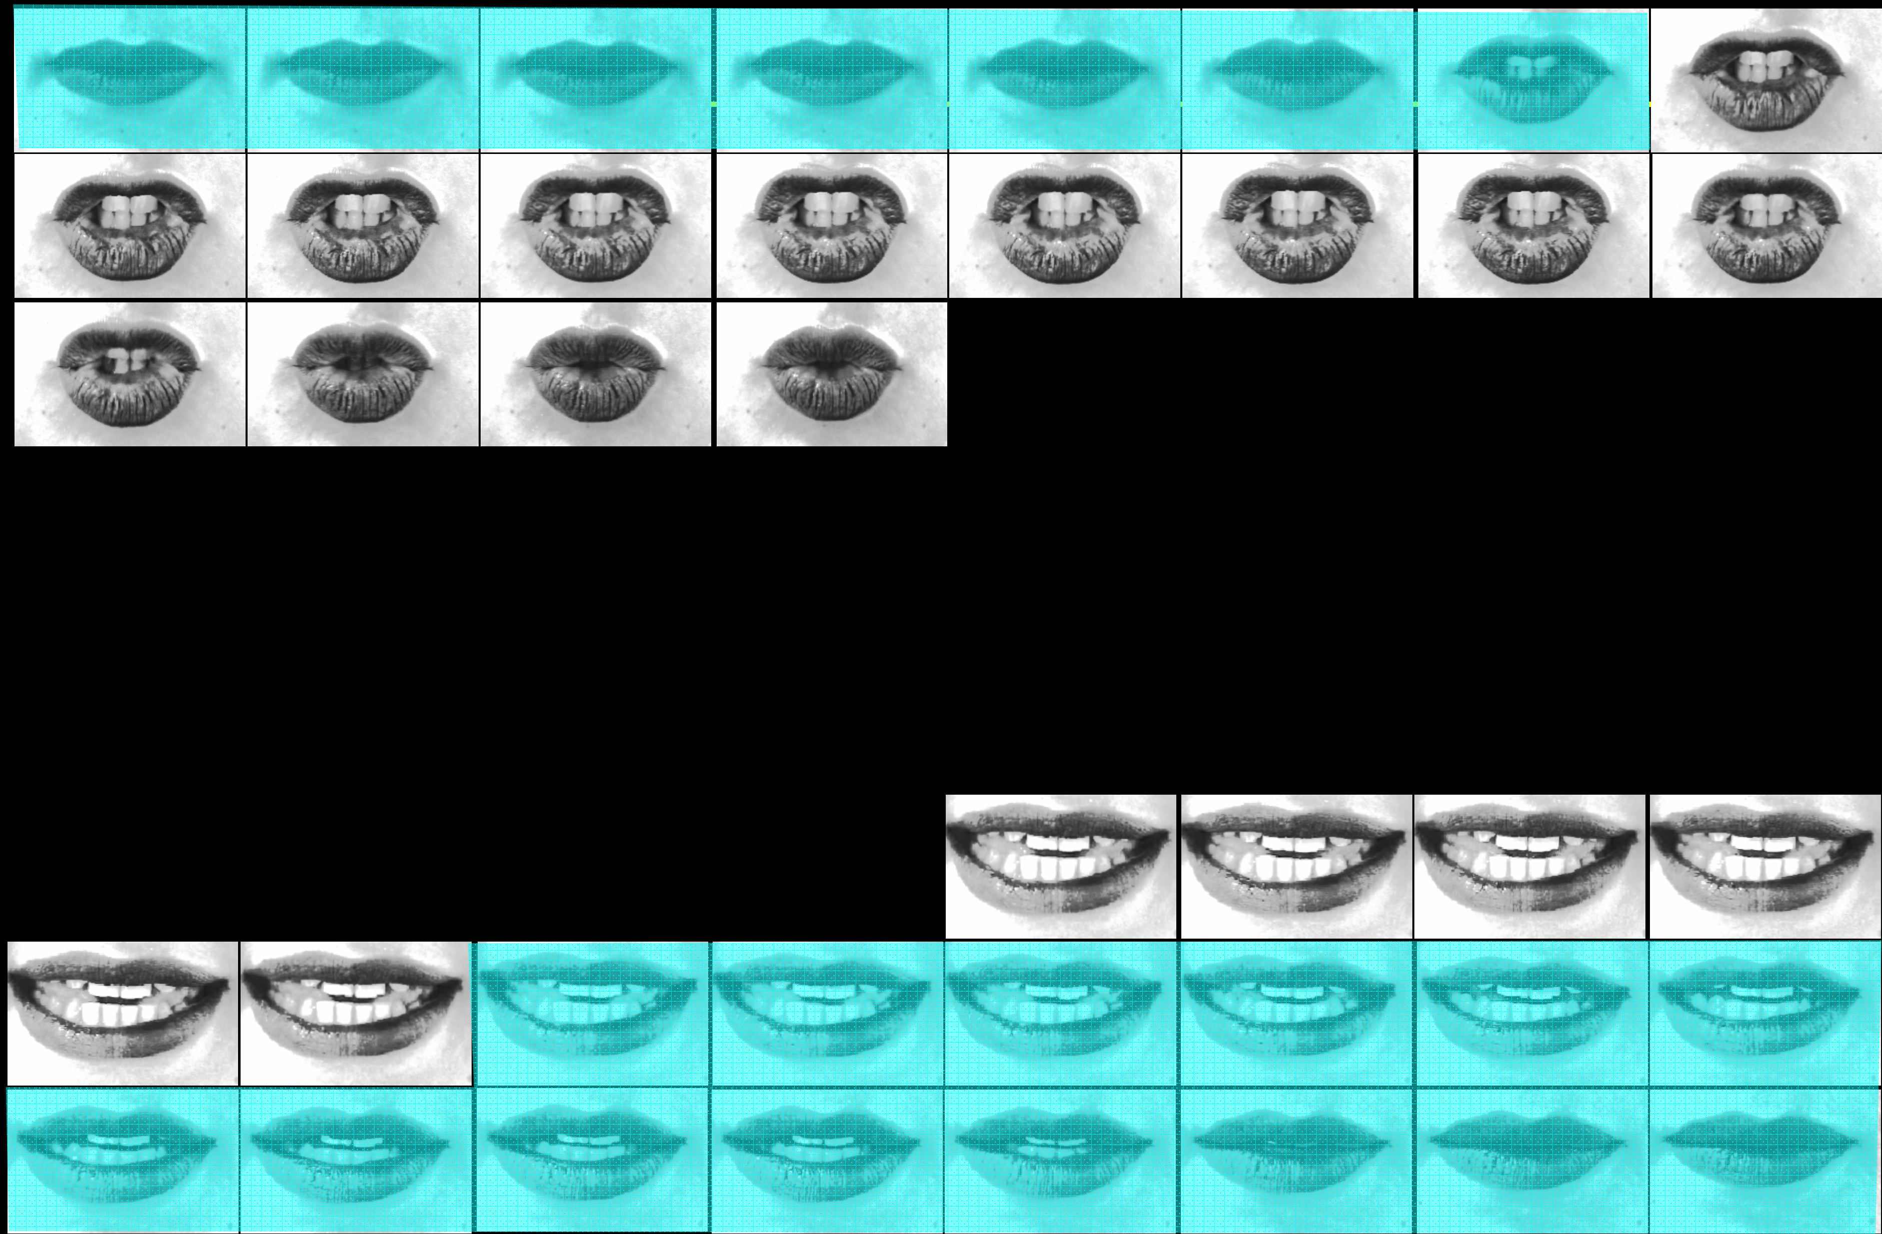

Synthesis - SHU+KI

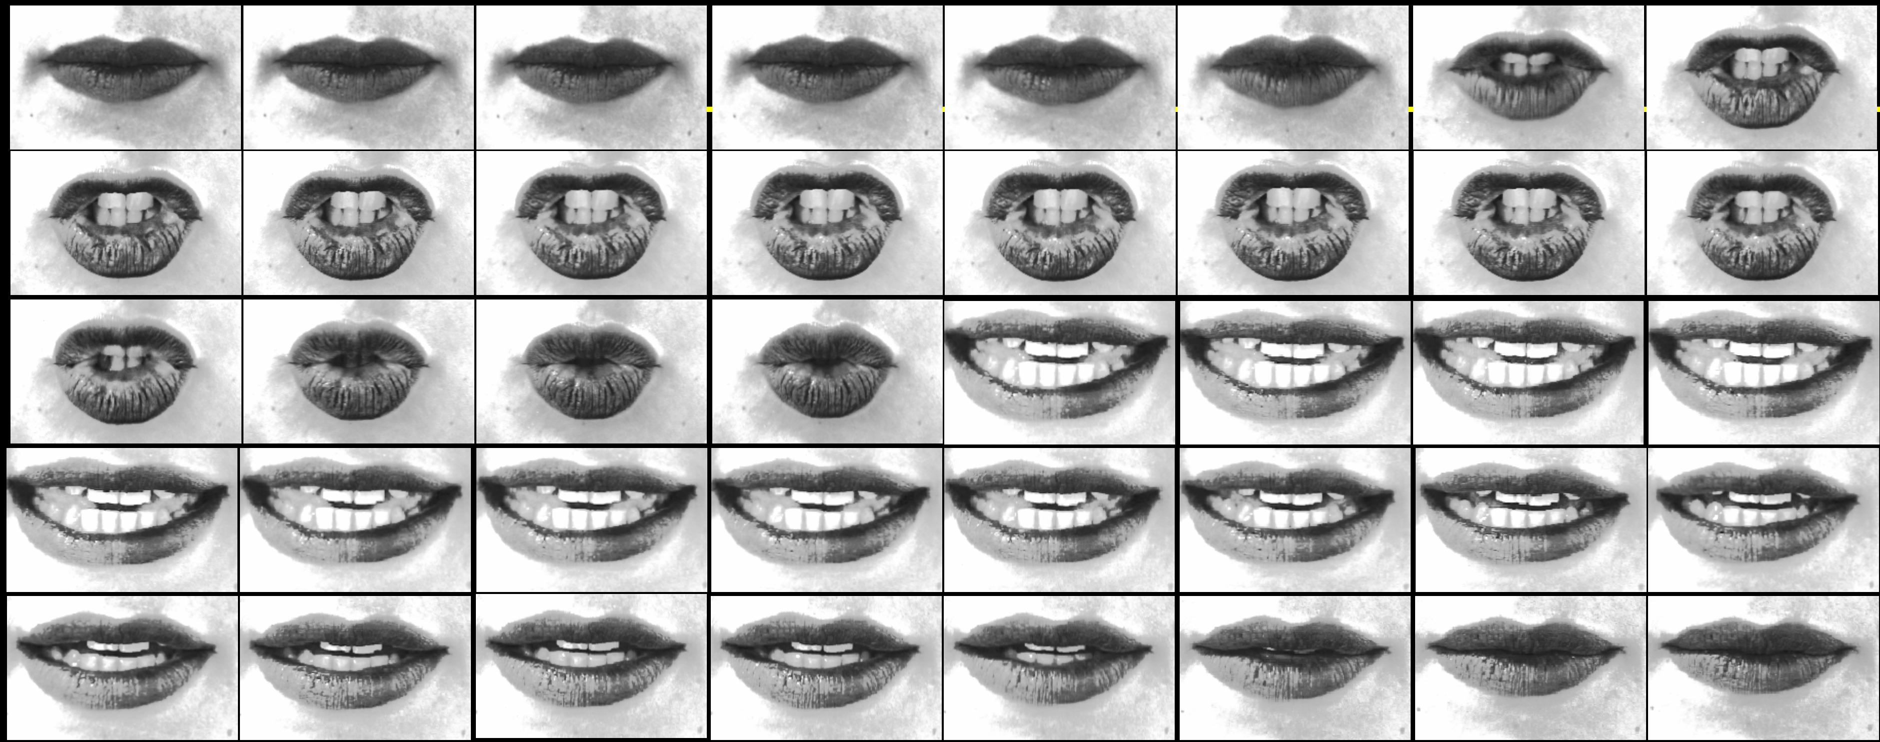

Synthesis - Example

- Concatenate the VAS of the syllables that make the word.
- Smooth the transition between each successive VAS.


Shu


Ki


Shuki – simple concatenation


Shuki– smooth transition

Synthesis - Example

- Concatenate the VAS of the syllables that make the word.
- Smooth the transition between each successive VAS.

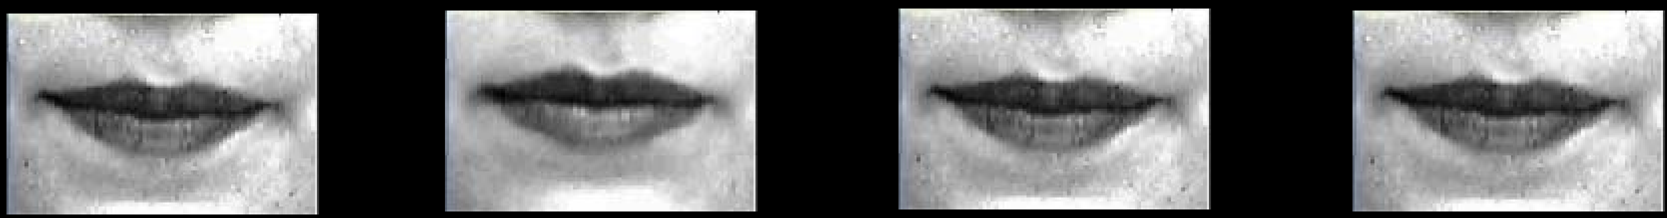

Mi

La

MiLa – simple concatenation

MiLa – smooth transition

Lip-reading Test

16 different words

"Coffee" "Cola" "Champagne"

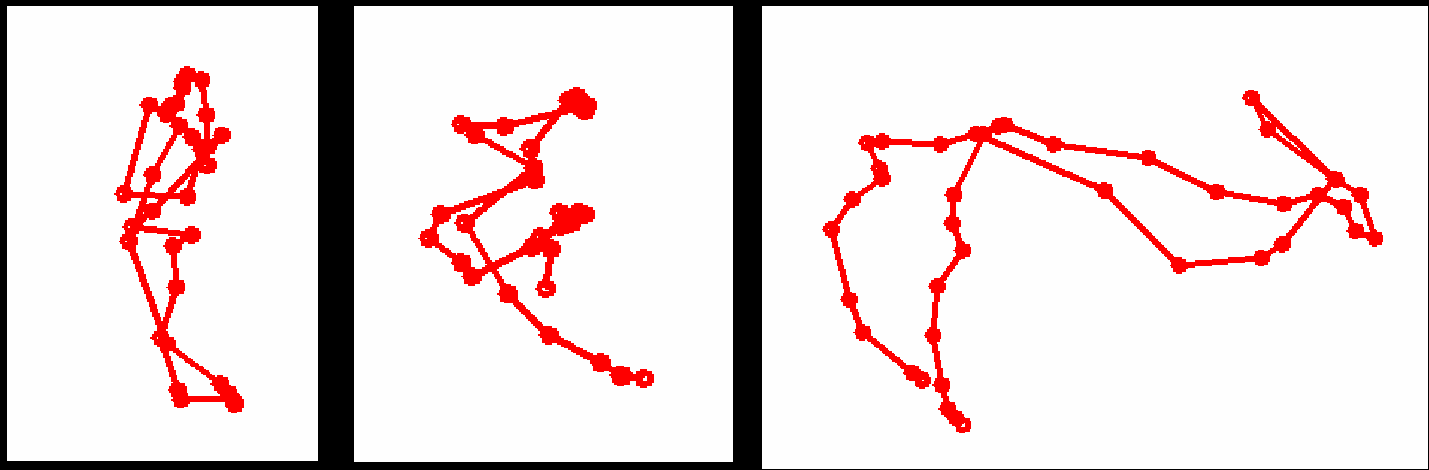

...

The lips area is extracted and a new contour is computed

The word contour is then matched to all others.

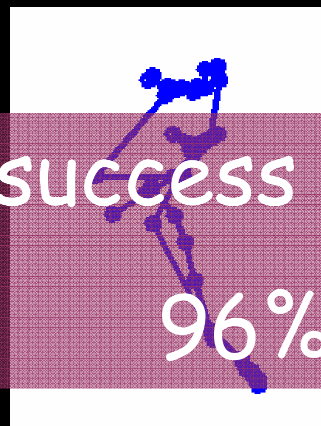

success rate

96%

Generalization: Images as anchors

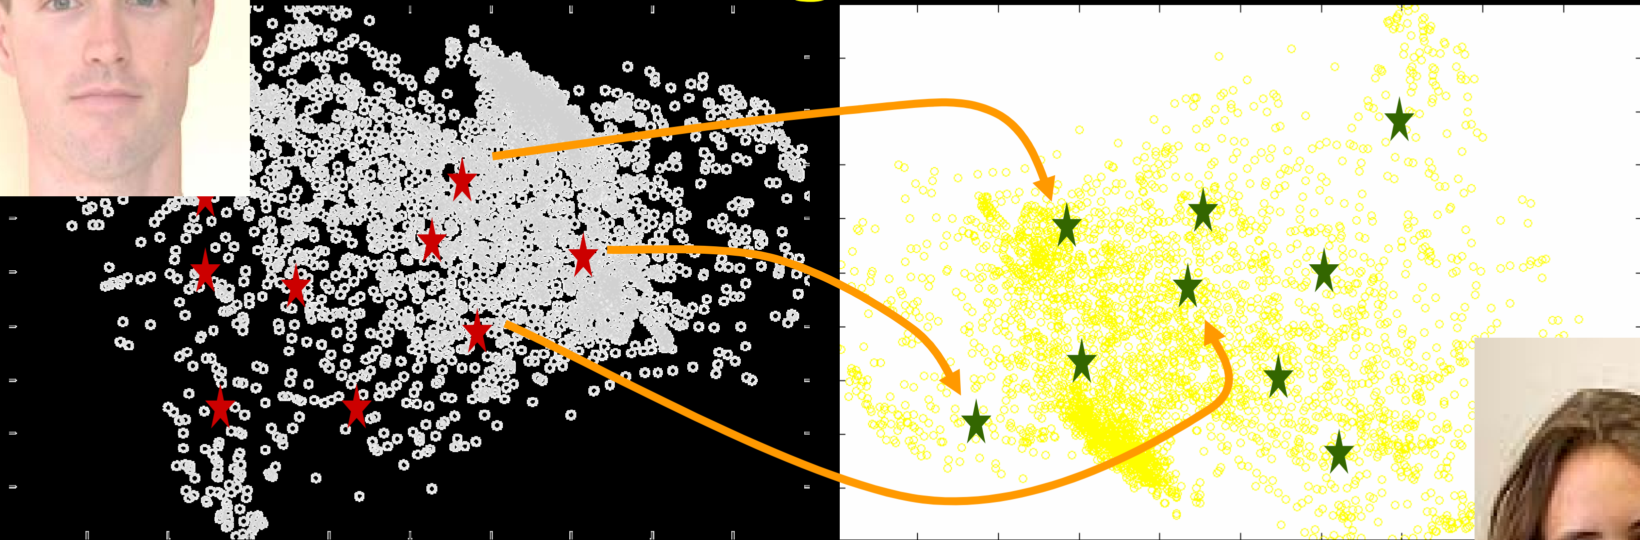

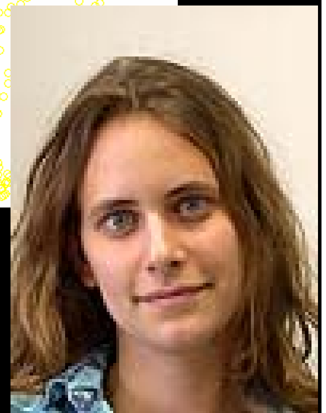

1st hit recognition rate 70%, 2nd hit rate 80%.

Distance Maps and Minimal Geodesics

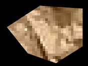

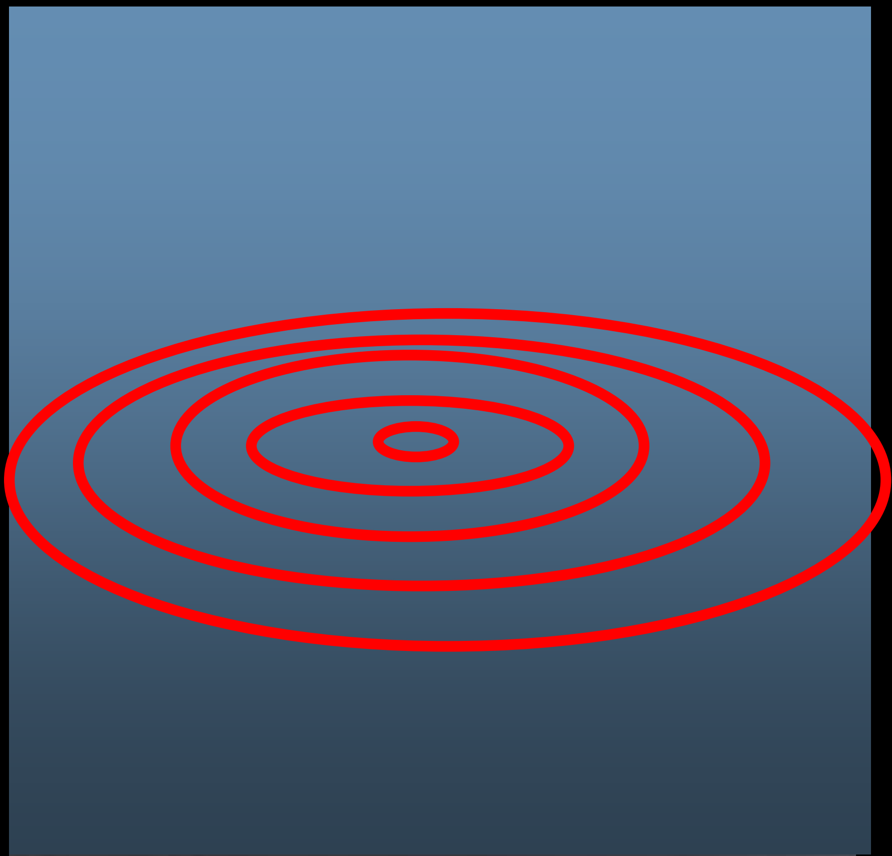

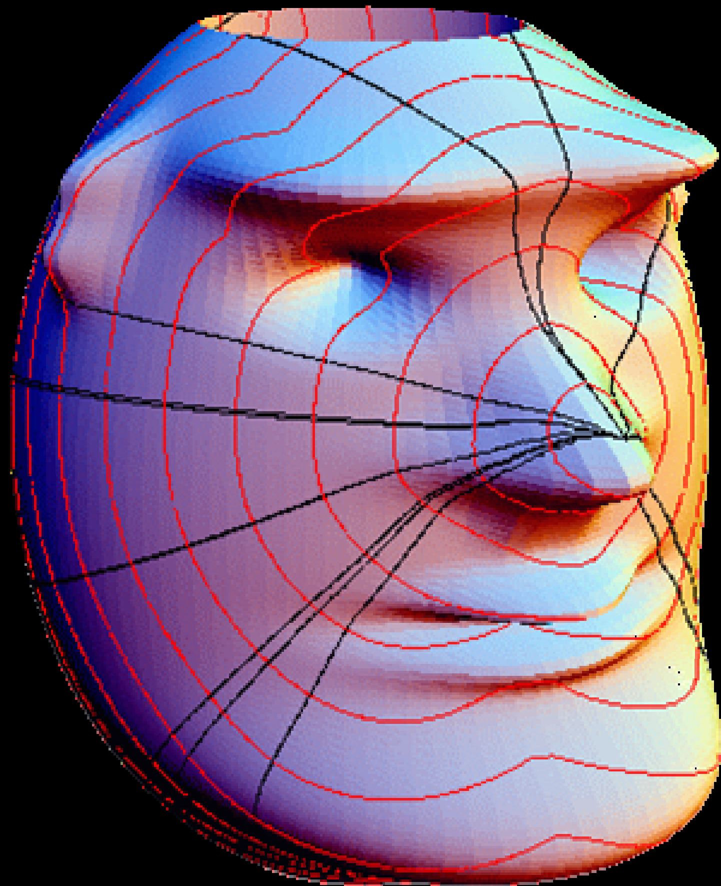

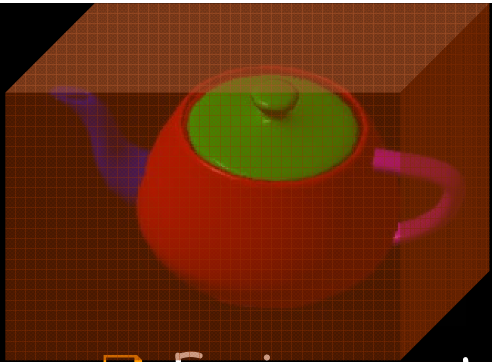

Texture Mapping

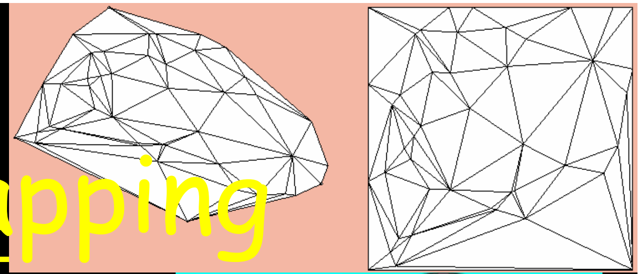

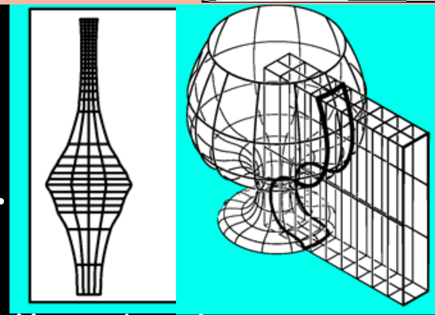

- Environment mapping: Blinn, Newell (76).
- Environment mapping: Greene, Bier and Sloan (86).
- Free-form surfaces: Arad and Elber (97).
- Polyhedral surfaces: Floater (96, 98), Levy and Mallet (98).
- Multi-dimensional scaling: Schwartz, Shaw and Wolfson (89).

Difficulties:

- Need for user intervention.
- Local and global distortions.
- Restrictive boundary conditions.
- High computational complexity.

Flattening via MDS

- Compute geodesic distances between pairs of points.
- Construct a square distance matrix of geodesic distances^2.
- Find the coordinates in the plane via multi-dimensional scaling.

  The simplest is `classical scaling'.
- Use the flattened coordinates for texturing the surface, while preserving the texture features.

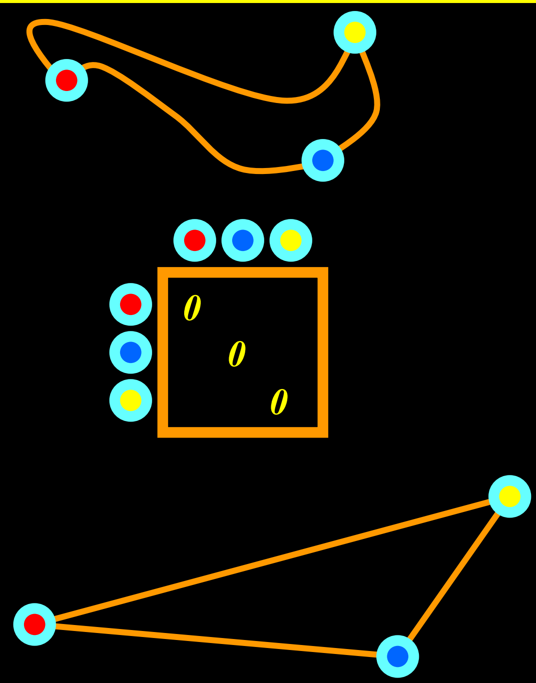

Zigelman, Kimmel, Kiryati, **IEEE T. on Visualization and Computer Graphics** 2002.

Flattening

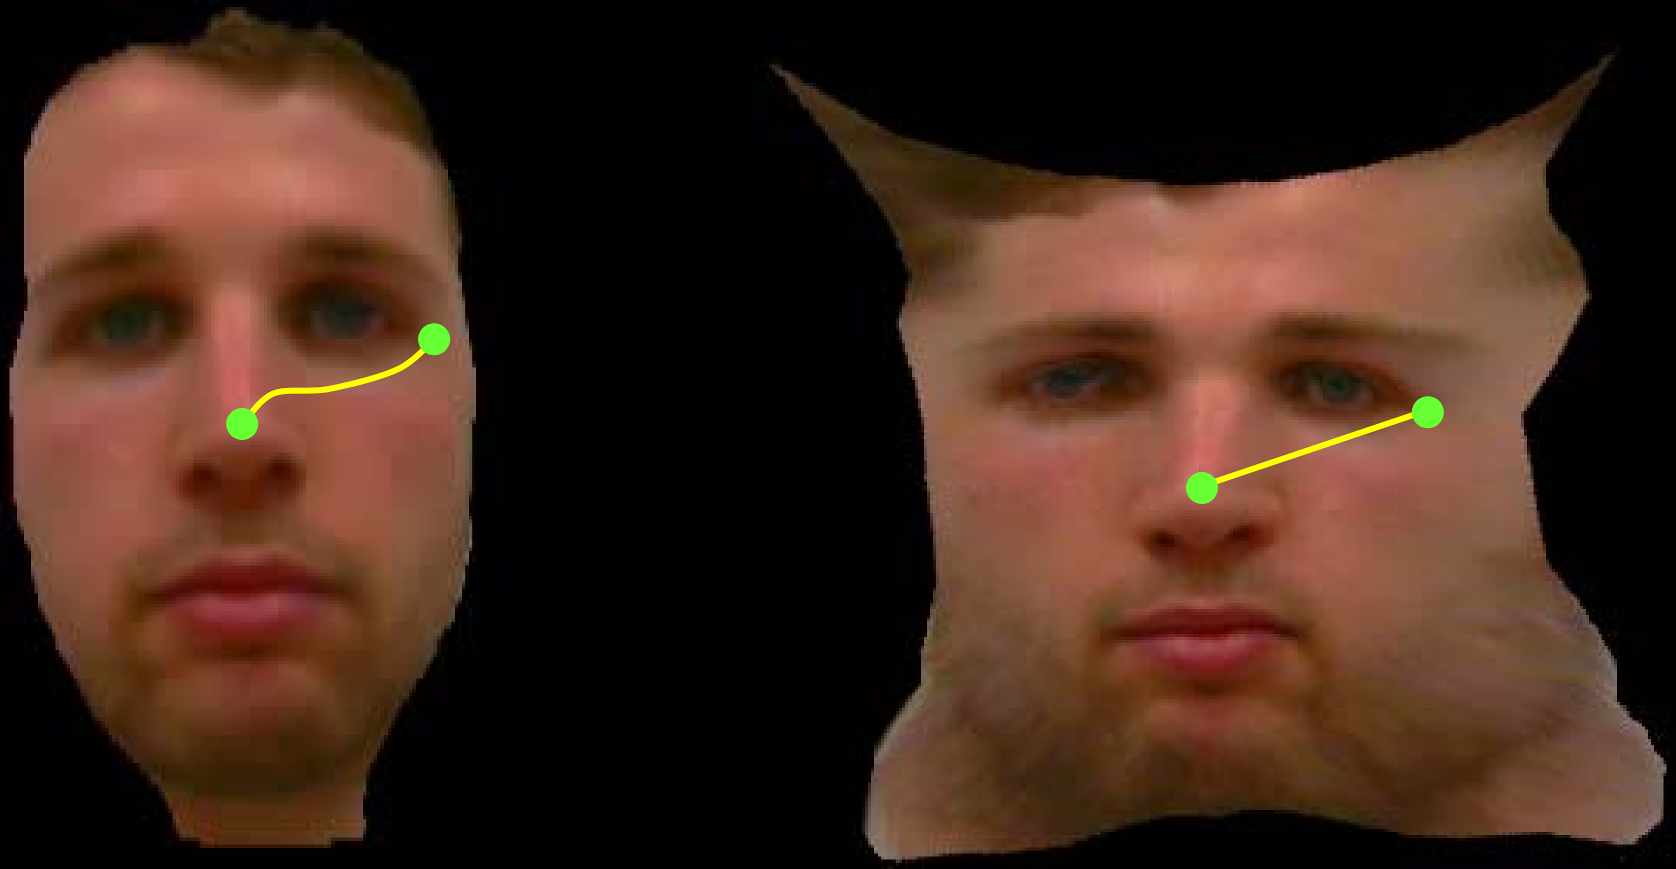

Gil Zigelman

Flattening

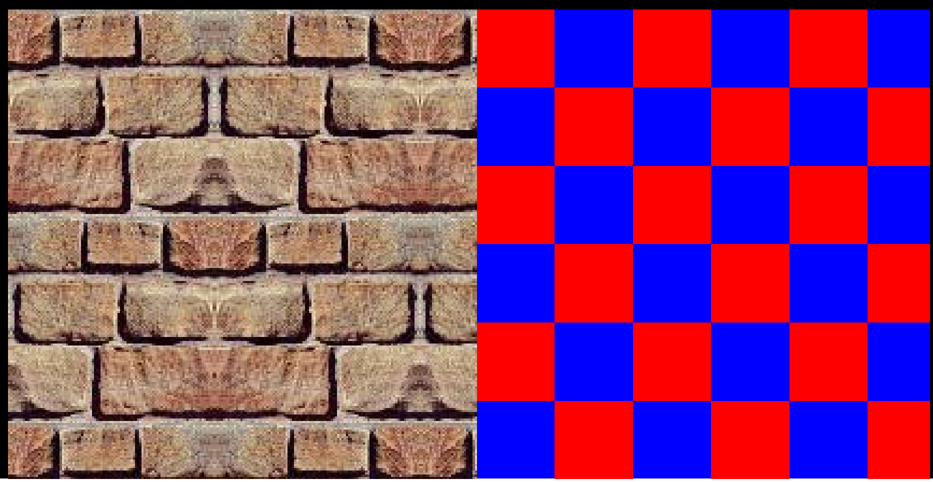

Bending invariant signatures

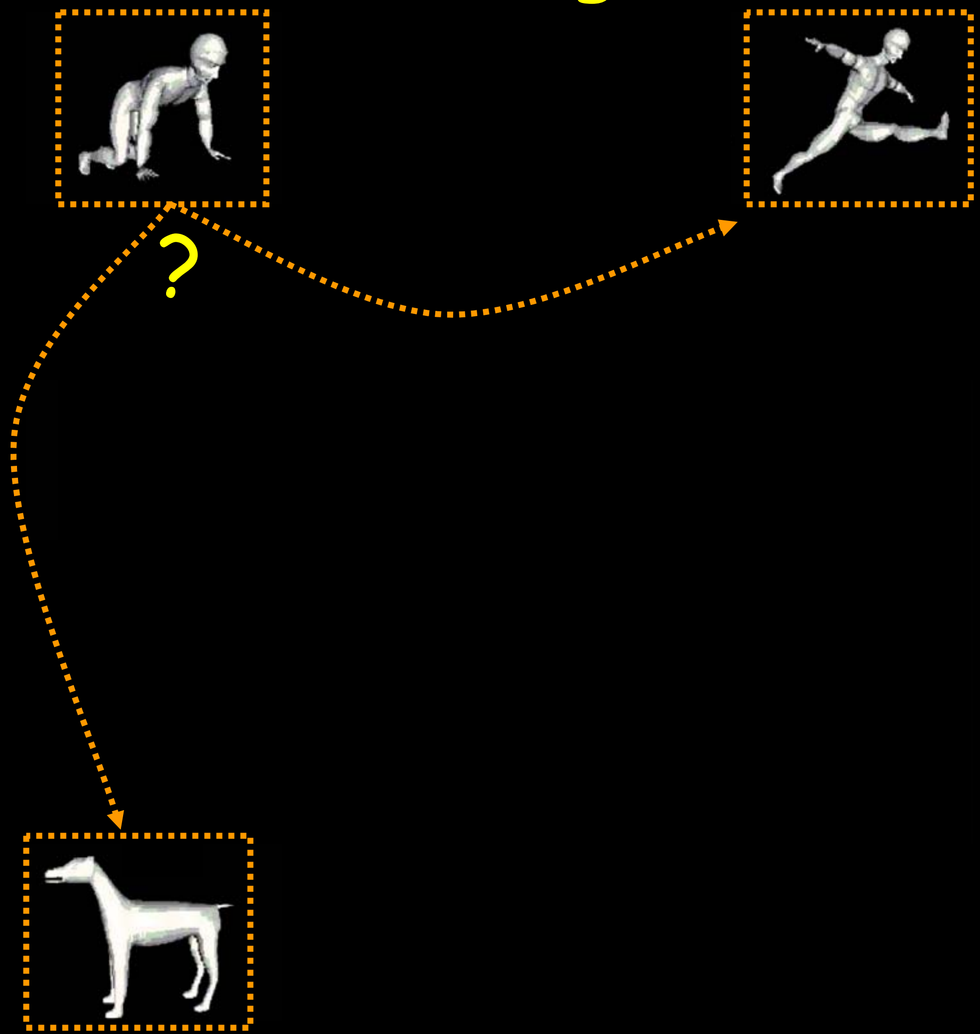

Asi Elad and Kimmel, *CVPR'2001/PAMI'2003*

Bending invariant signatures

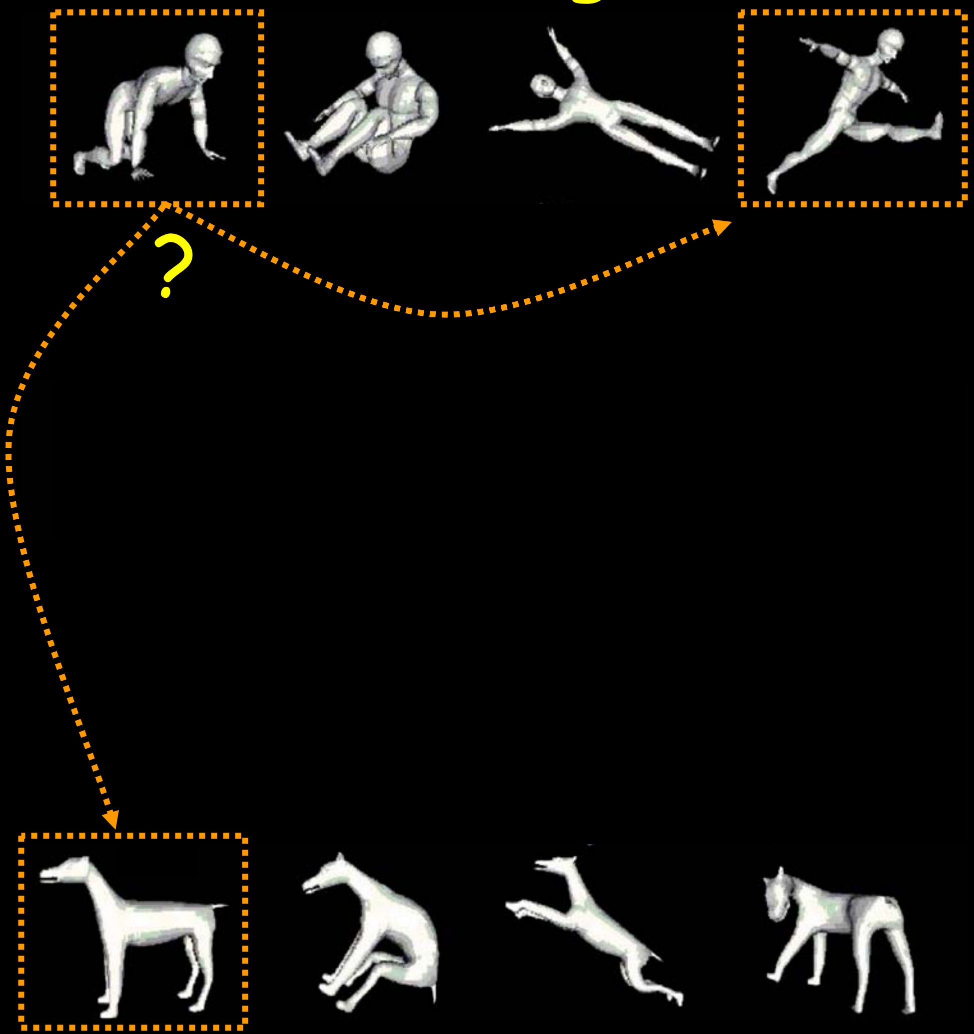

Elad and Kimmel, *CVPR'2001/PAMI'2003*

Bending invariant signatures

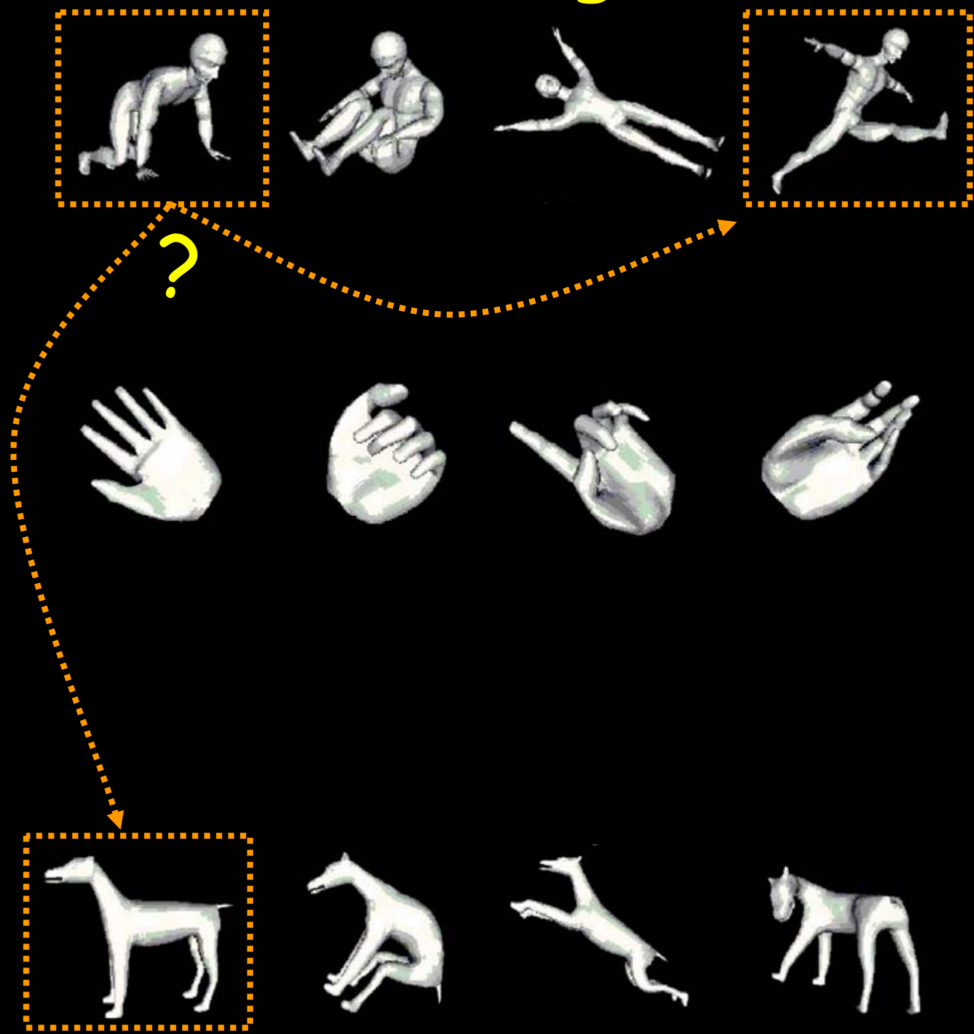

Elad and Kimmel, *CVPR'2001/PAMI'2003*

Bending Invariant Signatures

Elad and Kimmel, *CVPR'2001/PAMI'2003*

Bending Invariant Signatures

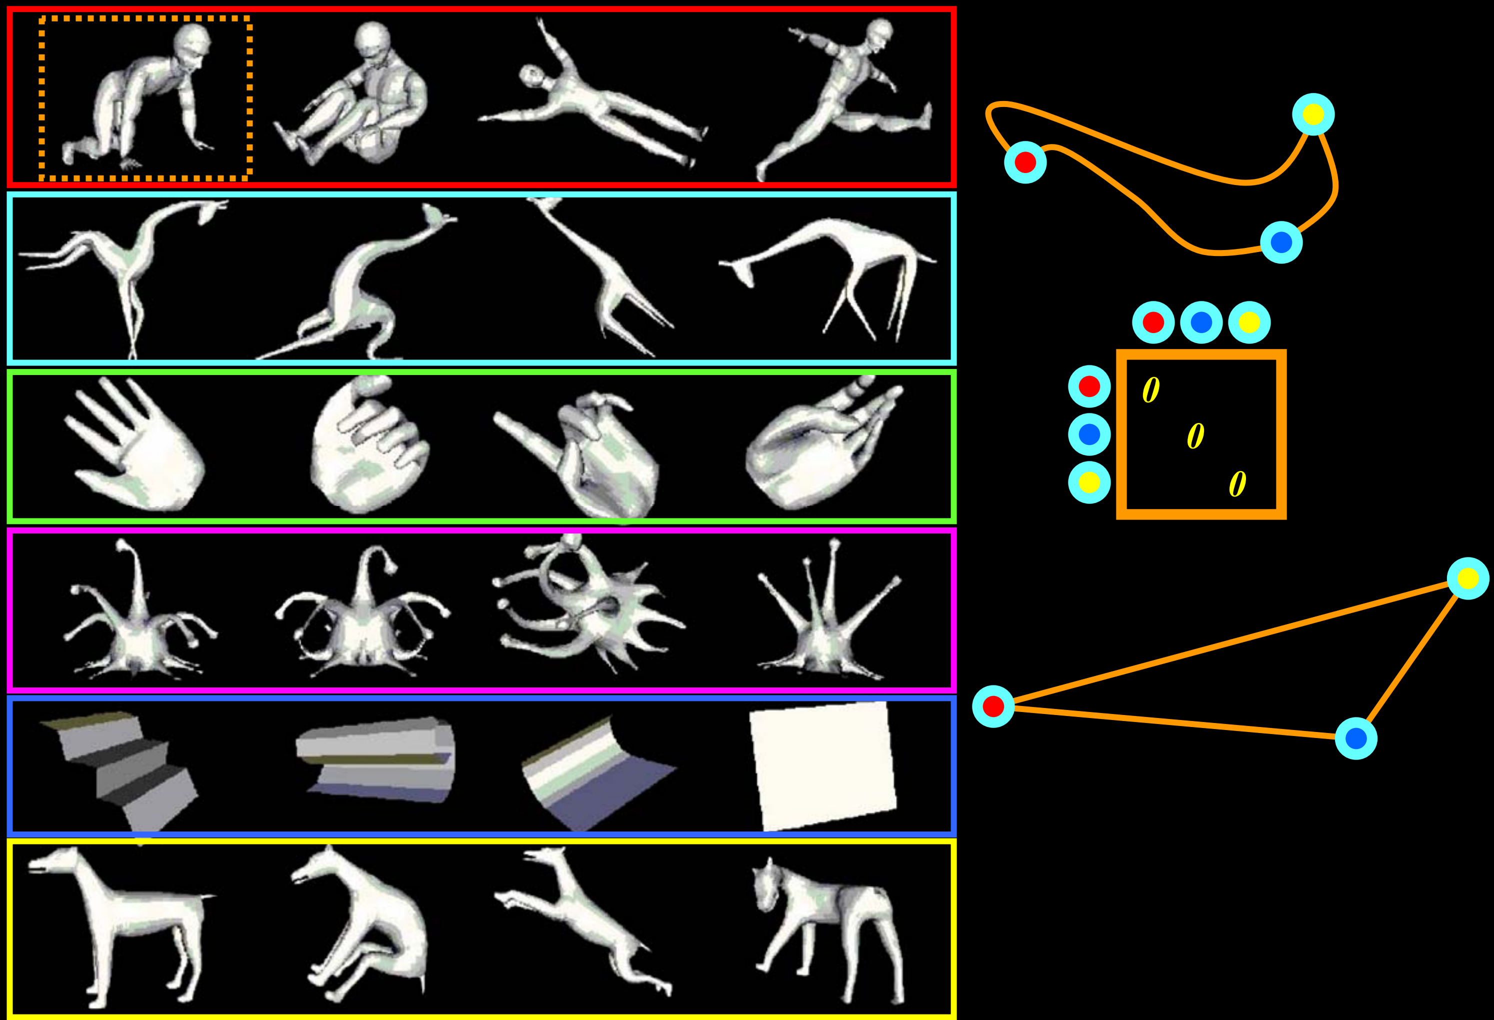

Elad and Kimmel, *CVPR'2001/PAMI'2003*

Bending invariant signatures

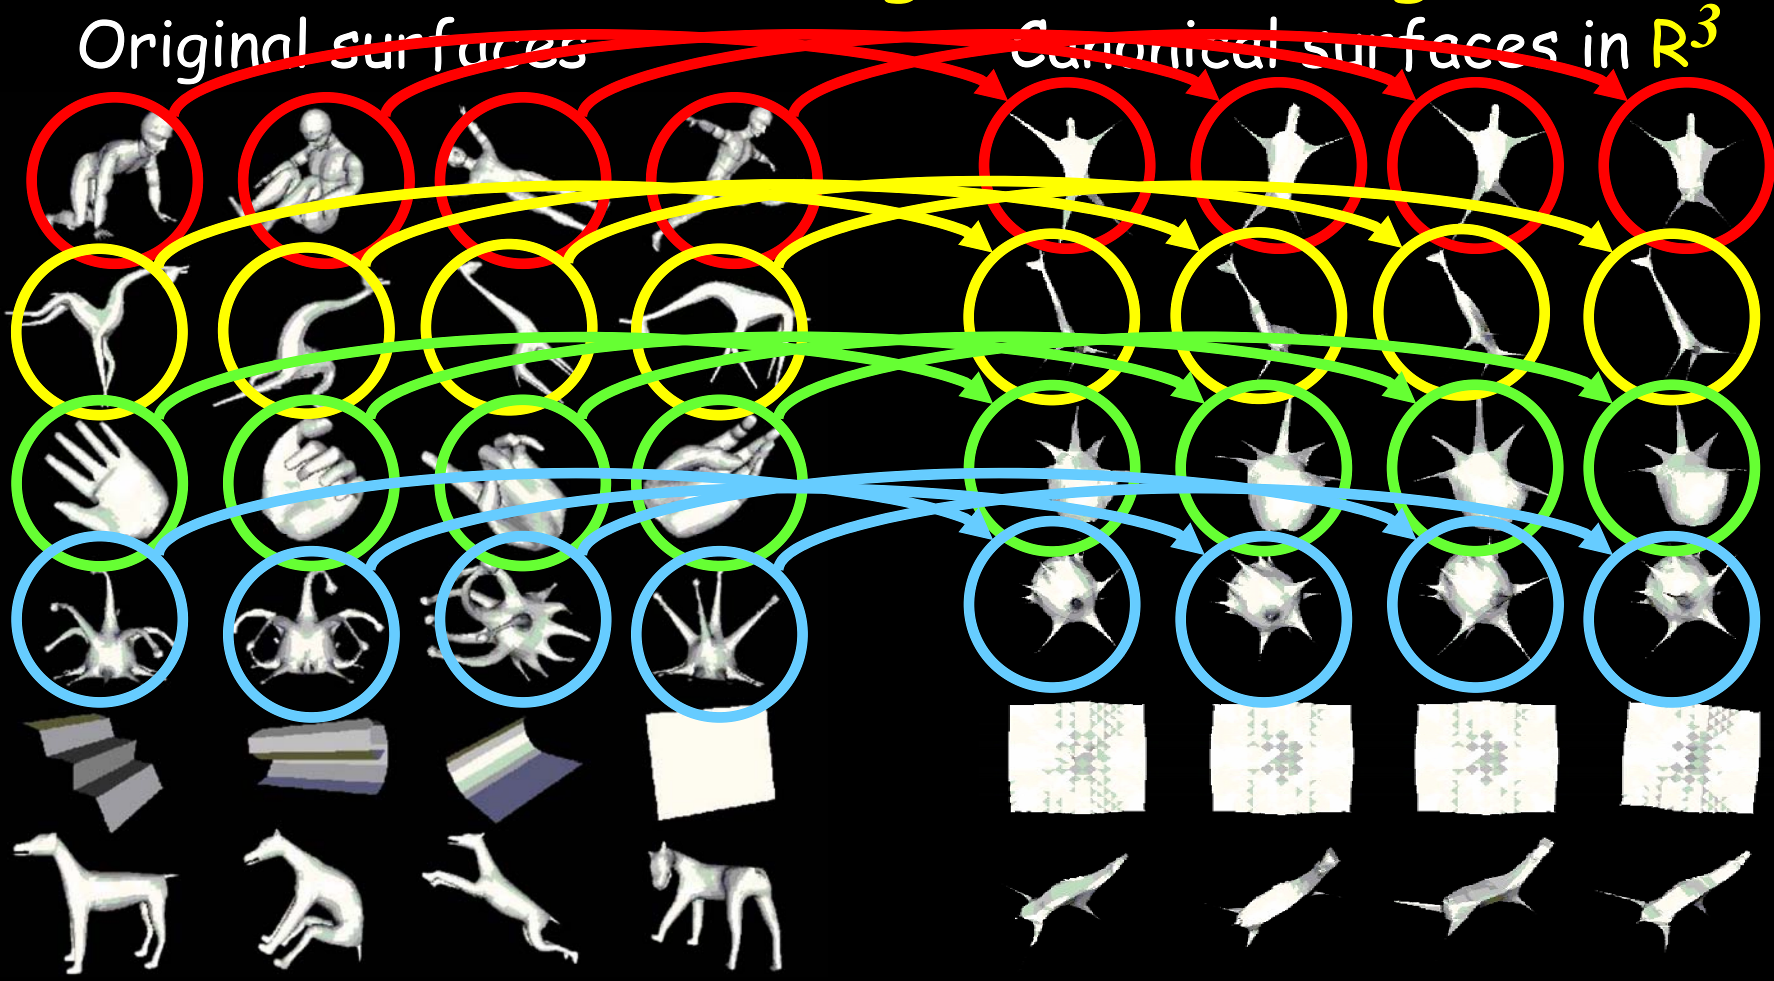

Elad and Kimmel, *CVPR'2001/PAMI'2003*

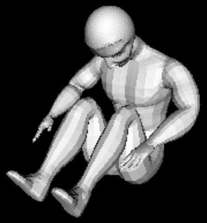 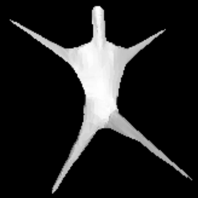

Classification

Original surfaces — Our signatures

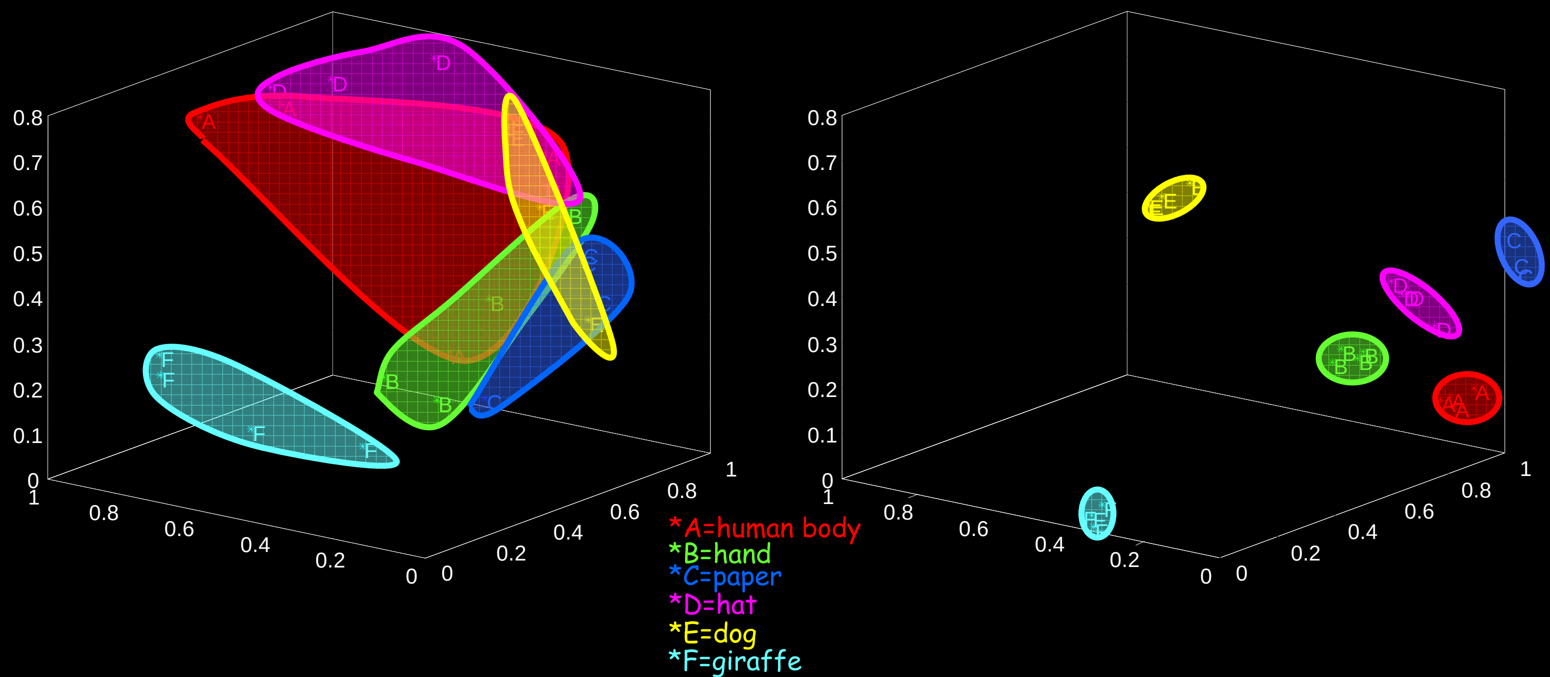

Elad Kimmel *2001*

Face Recognition

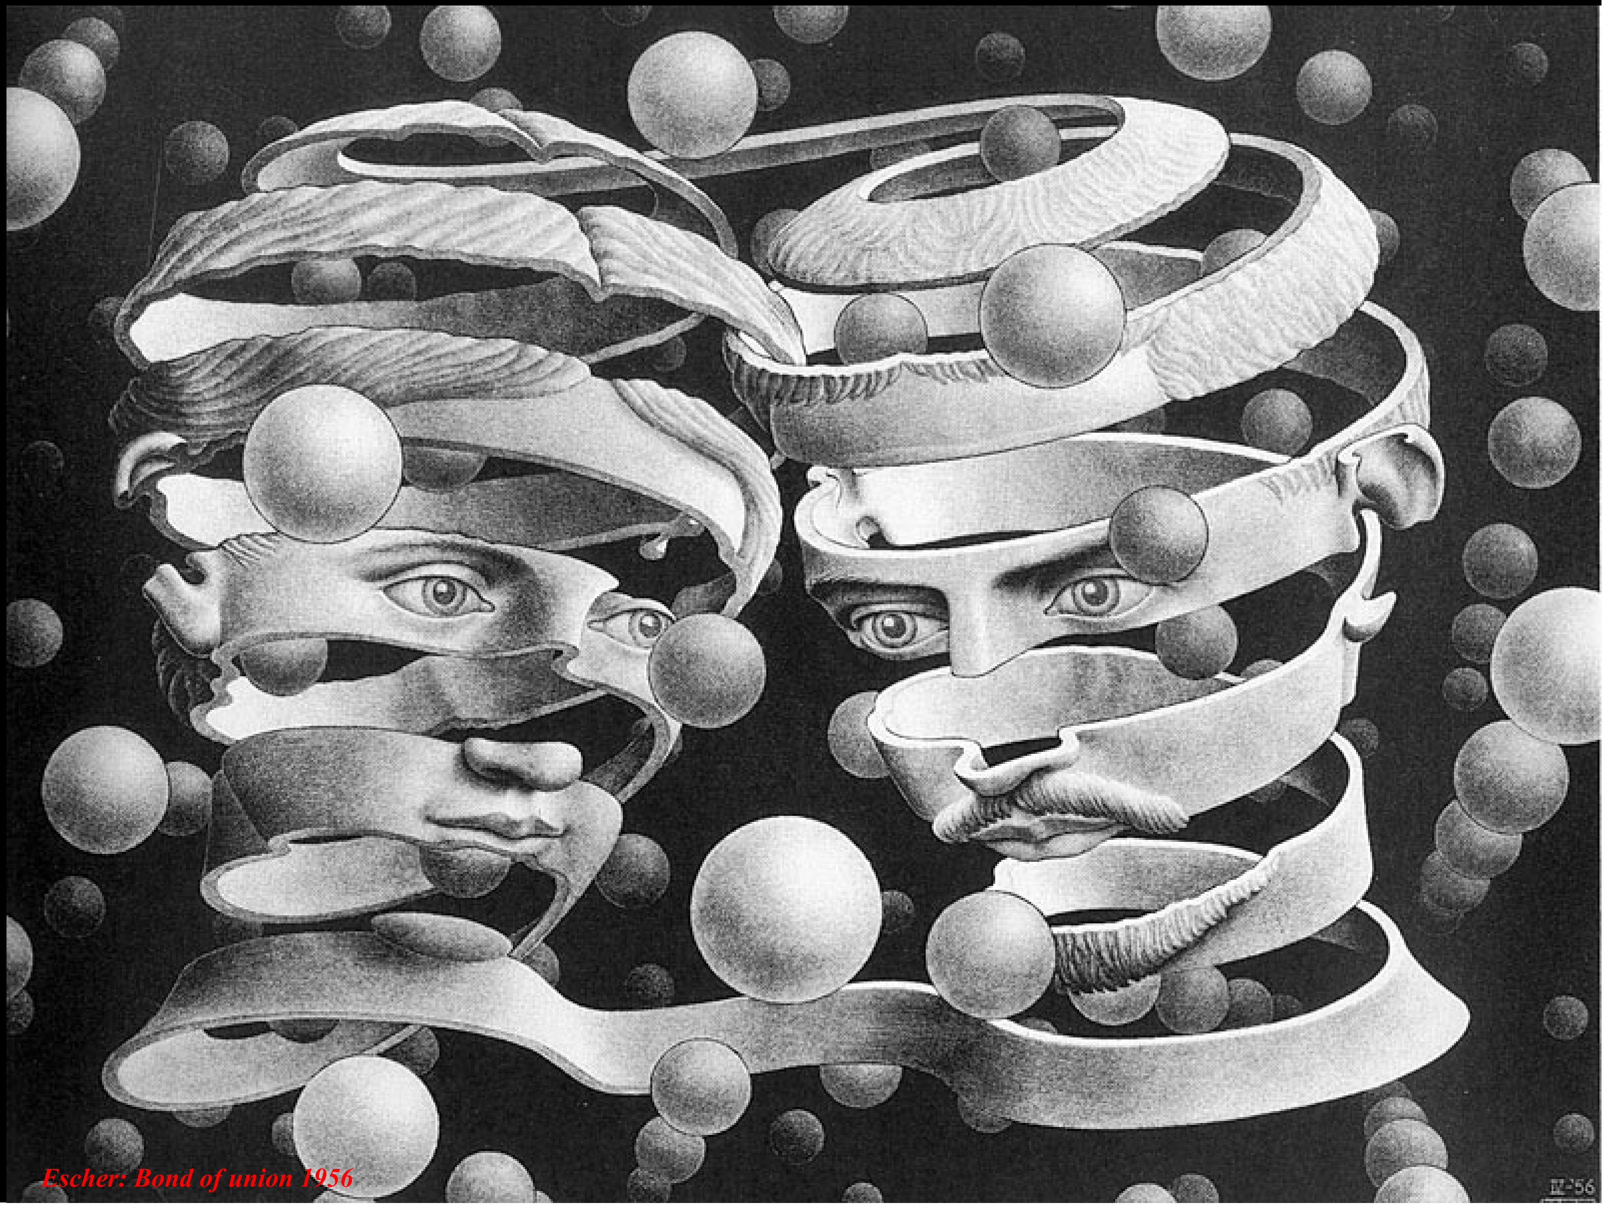

Escher: Bond of union 1956

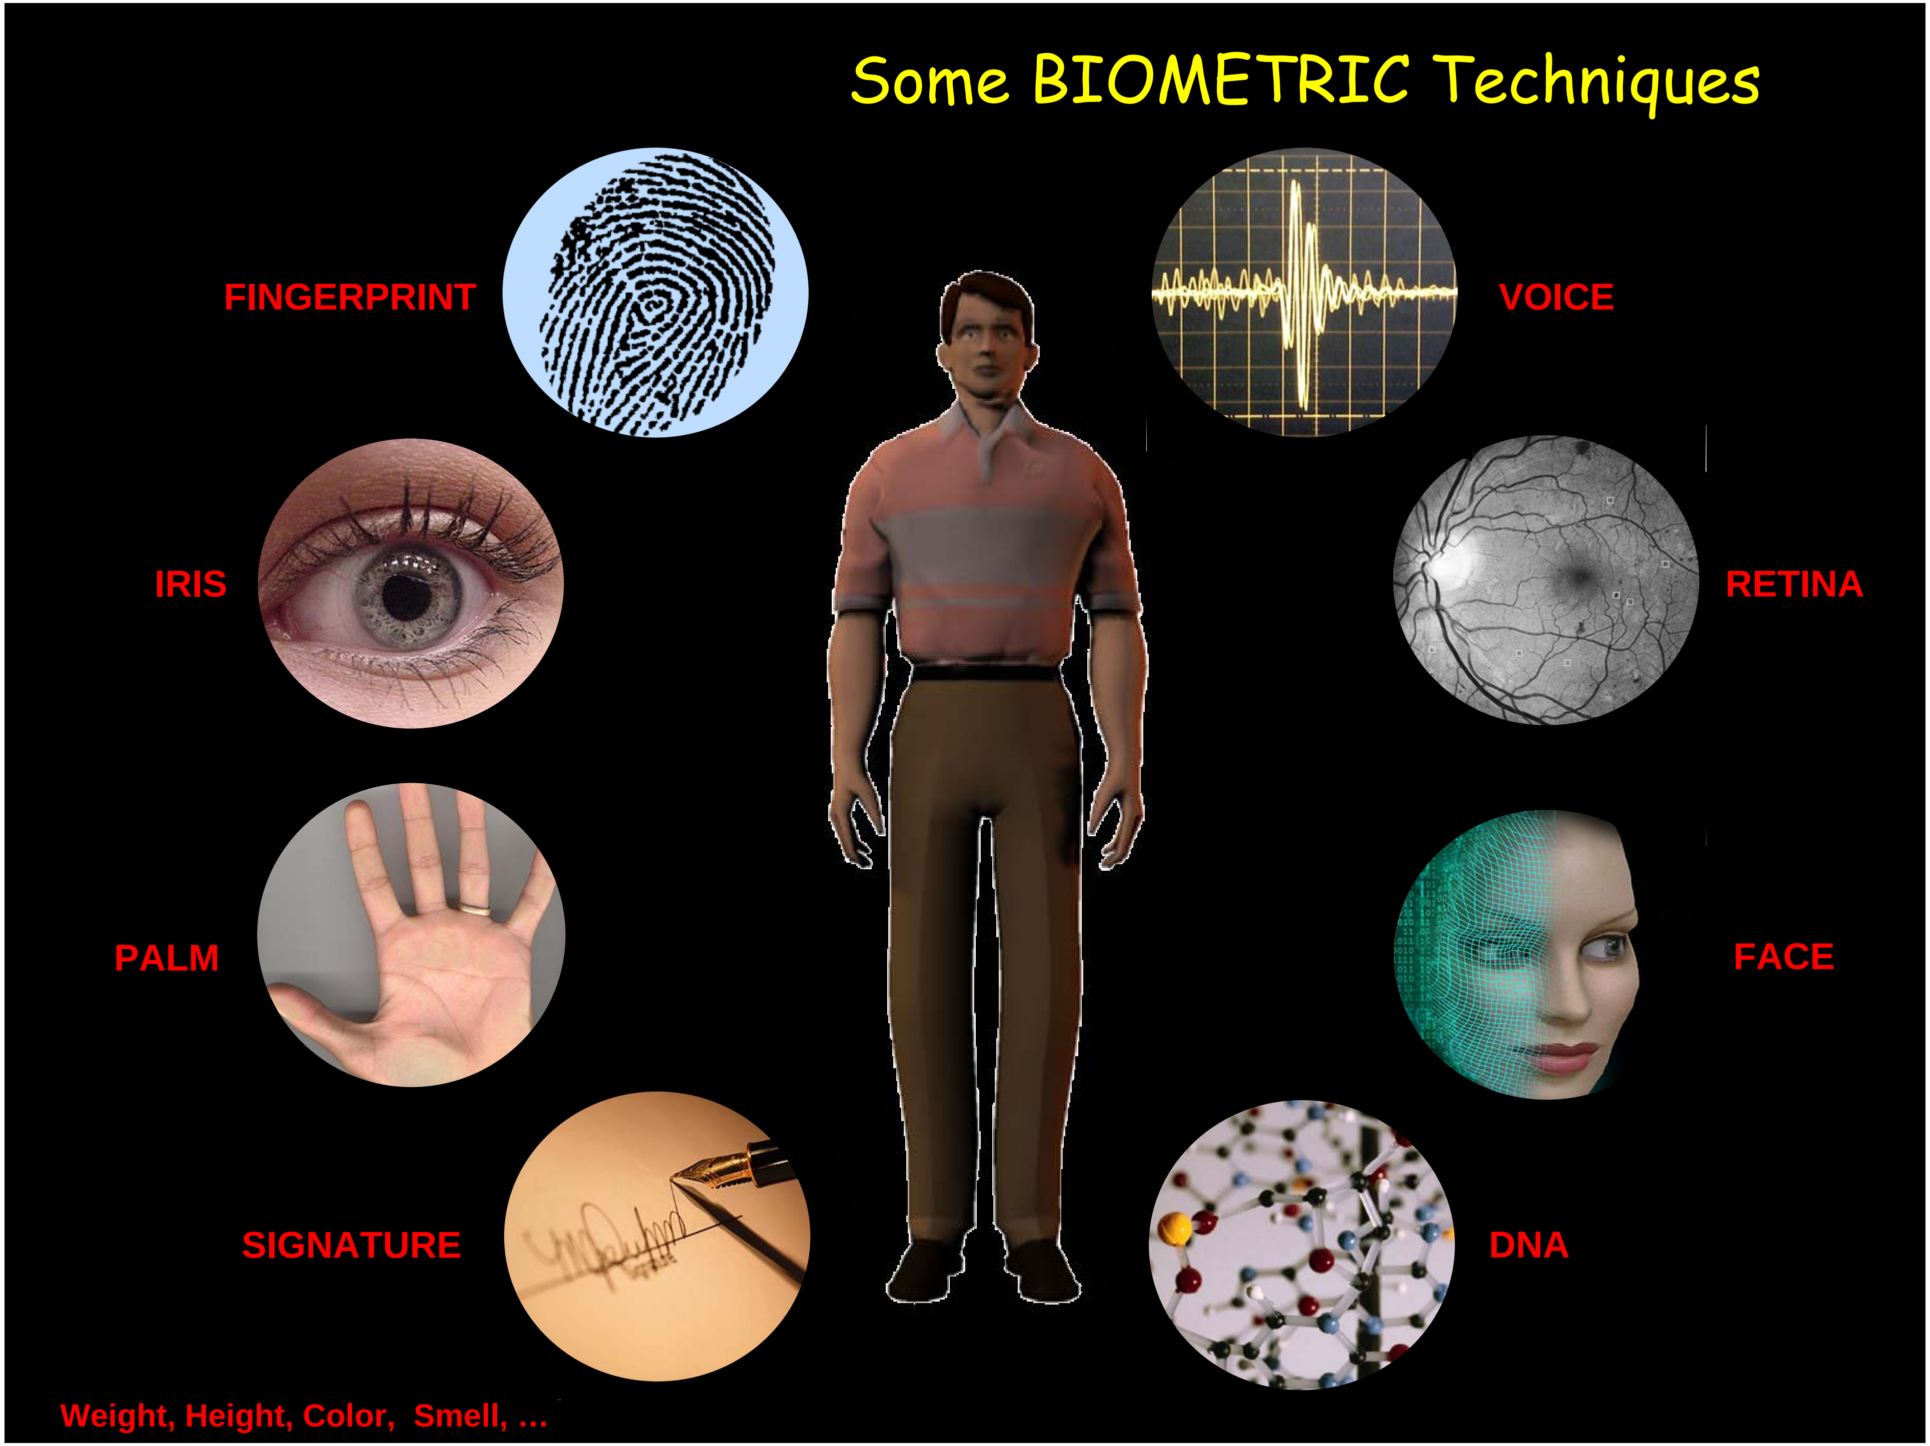
Some BIOMETRIC Techniques
FINGERPRINT
VOICE
IRIS
RETINA
PALM
FACE
SIGNATURE
DNA
Weight, Height, Color, Smell, …

Twins test

who is Michael?

FACE RECOGNITION: PROBLEMS

False acceptance: accept impostors as authenticated persons.

False rejection: fail to recognize an authenticated person.

ILLUMINATION and MAKEUP

FACIAL EXPRESSIONS

image = light reflected from the face, different illuminations yield different images, thus recognized as different subjects.

Modern face recognition algorithms unable to deal with facial expressions.

SOLUTION: THREE DIMENSIONAL FACE RECOGNITION

2D face recognition

CONs

- Less accurate than other biometrics
- Sensitive to environment conditions, postures, and expressions.
- Sensitivity to fooling (makeup)

PROs

- No direct contact
- Friendly to users (indirect contact)
- Passive monitoring (surveillance)
- Low cost

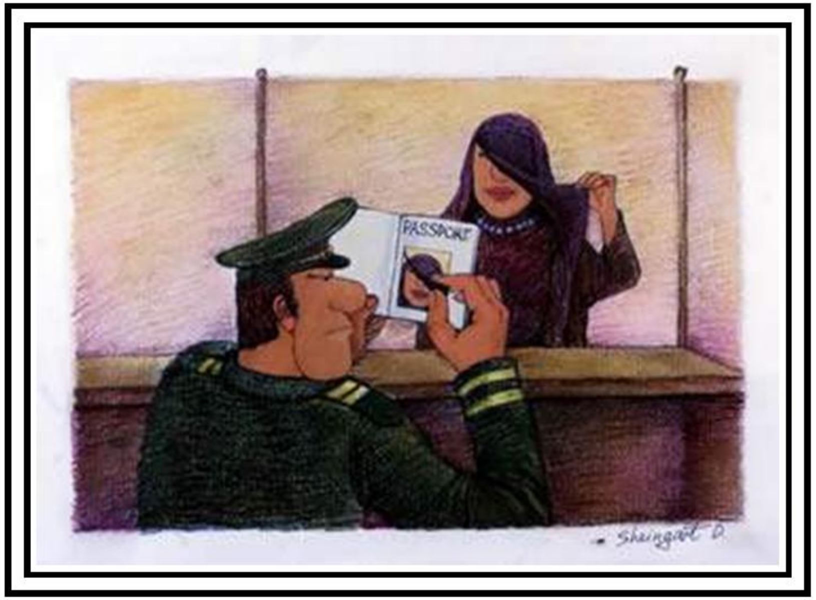

Why 3D?

2D information is sometimes misleading...

Why 3D?

...and affected by different factors

Can a 2D system tell that these images are all of the same subject?

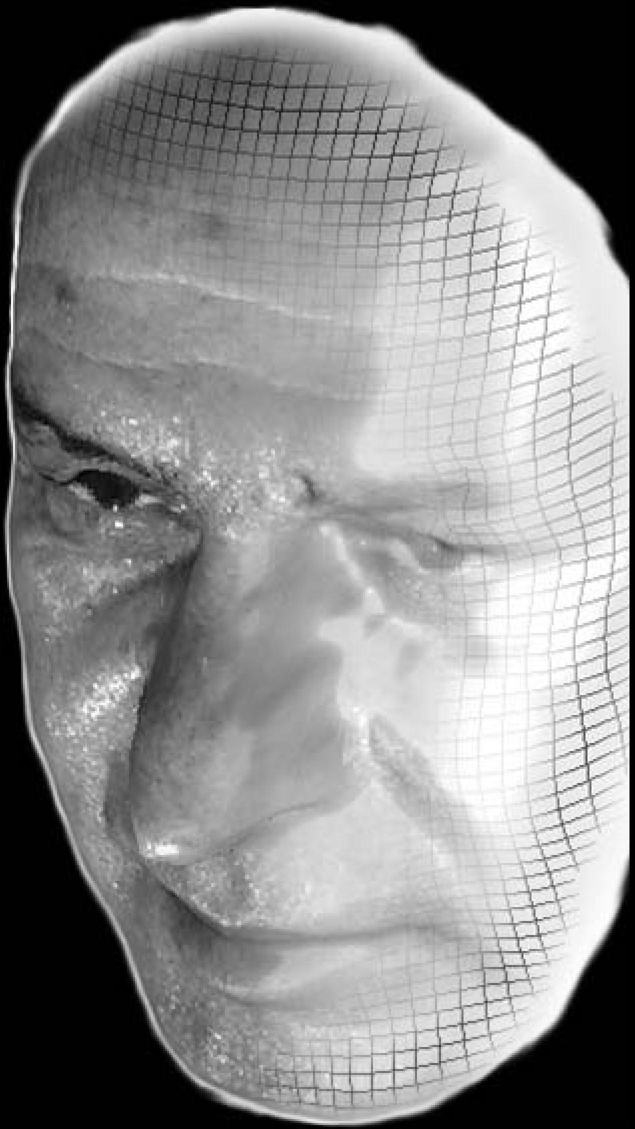

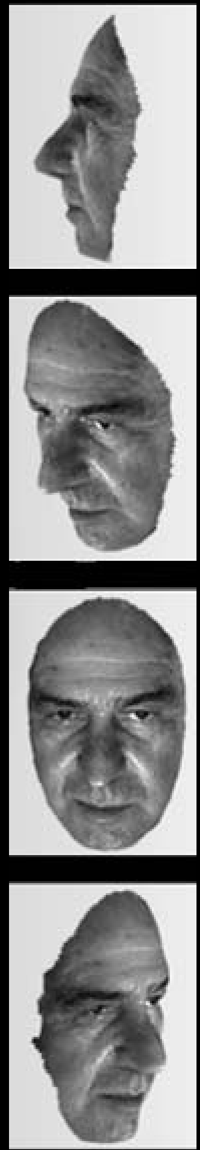

Coded light range video camera

N patterns allow angular (~depth) resolution of 2^N.

ex-minister of science, Matan Vilnai.

Our new scanner works at less than 150msec,
Components cost ~3k$

Coded light range video camera

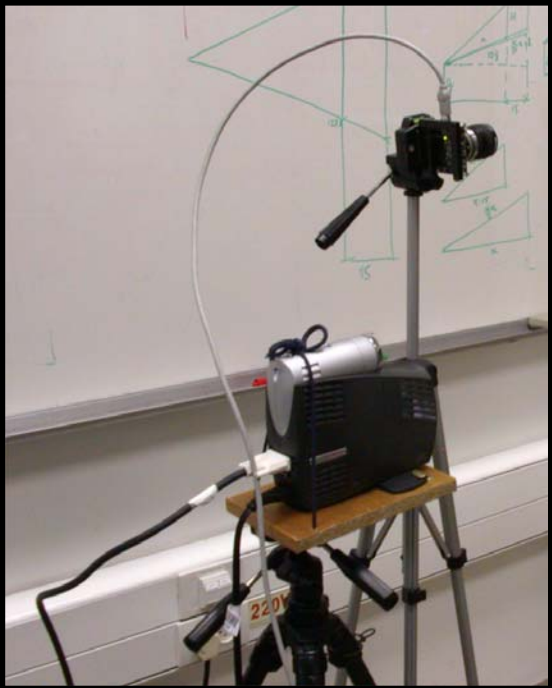

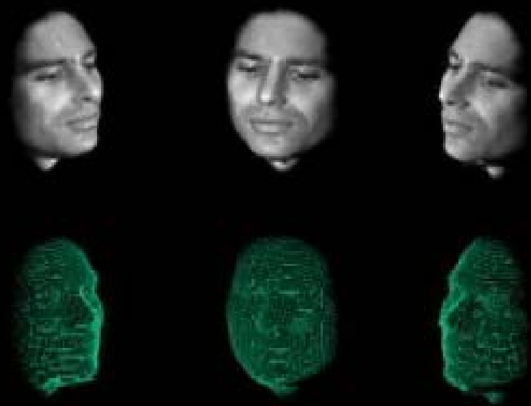

3D recognition via geometric invariants

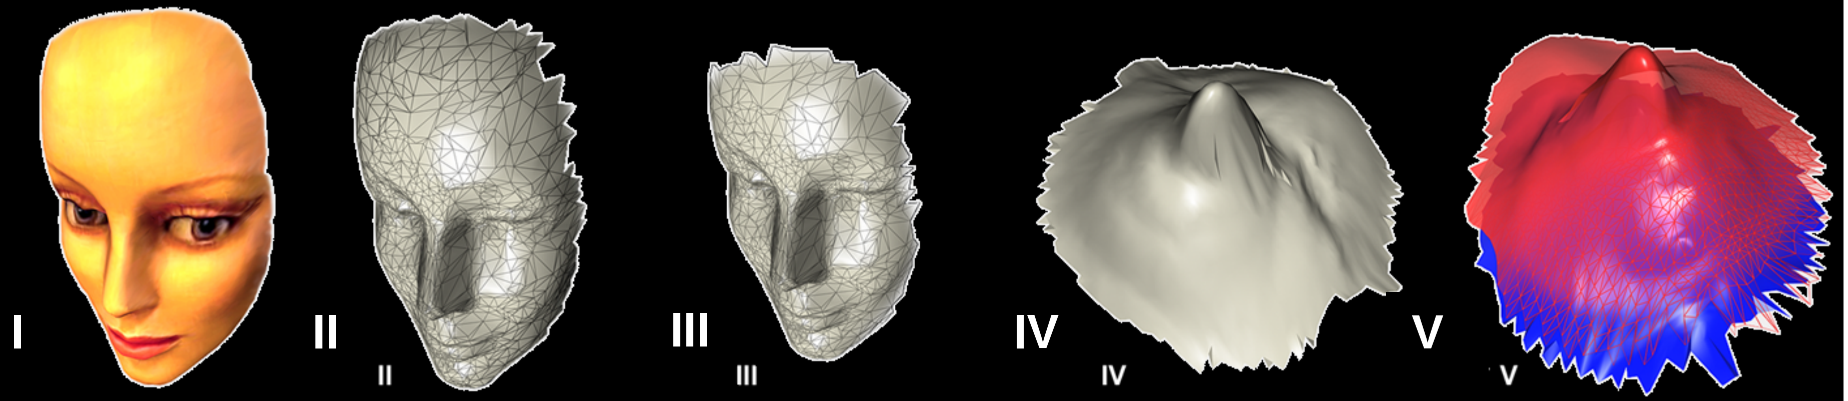

- Range camera acquires facial surface (**I**).
- The surface is smoothed (**II**), subsampled and cropped (**III**).
- Fast marching computes geodesic distances on the surface.
- Facial surface is flattened via MDS (**IV**).
- Rigid surface matching using the canonical surfaces (**V**).

A. Bronstein, M. Bronstein and R. Kimmel, “3D face recognition using geometric invariants“, 2003

Eigenforms

- The training set include: flattened texture + canonical image.
- Applying eigendecomposition to the two sets, we get two eigenspaces.
- The resulting eigenvectors are our eigenforms.

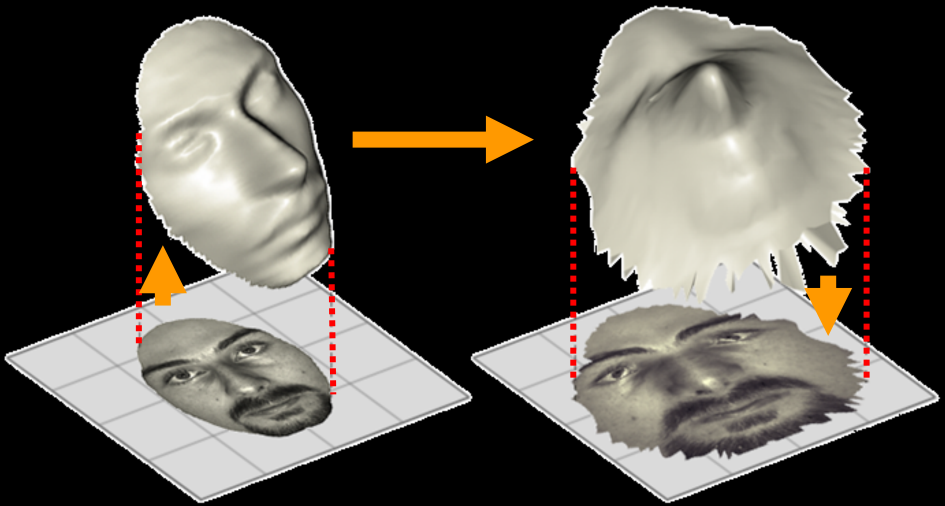

flattened texture

canonical image

A. Bronstein, M. Bronstein and R. Kimmel, “Expression-invariant 3D face recognition“, **AVBPA**, 2003

3D FACE RECOGNITION STAGES

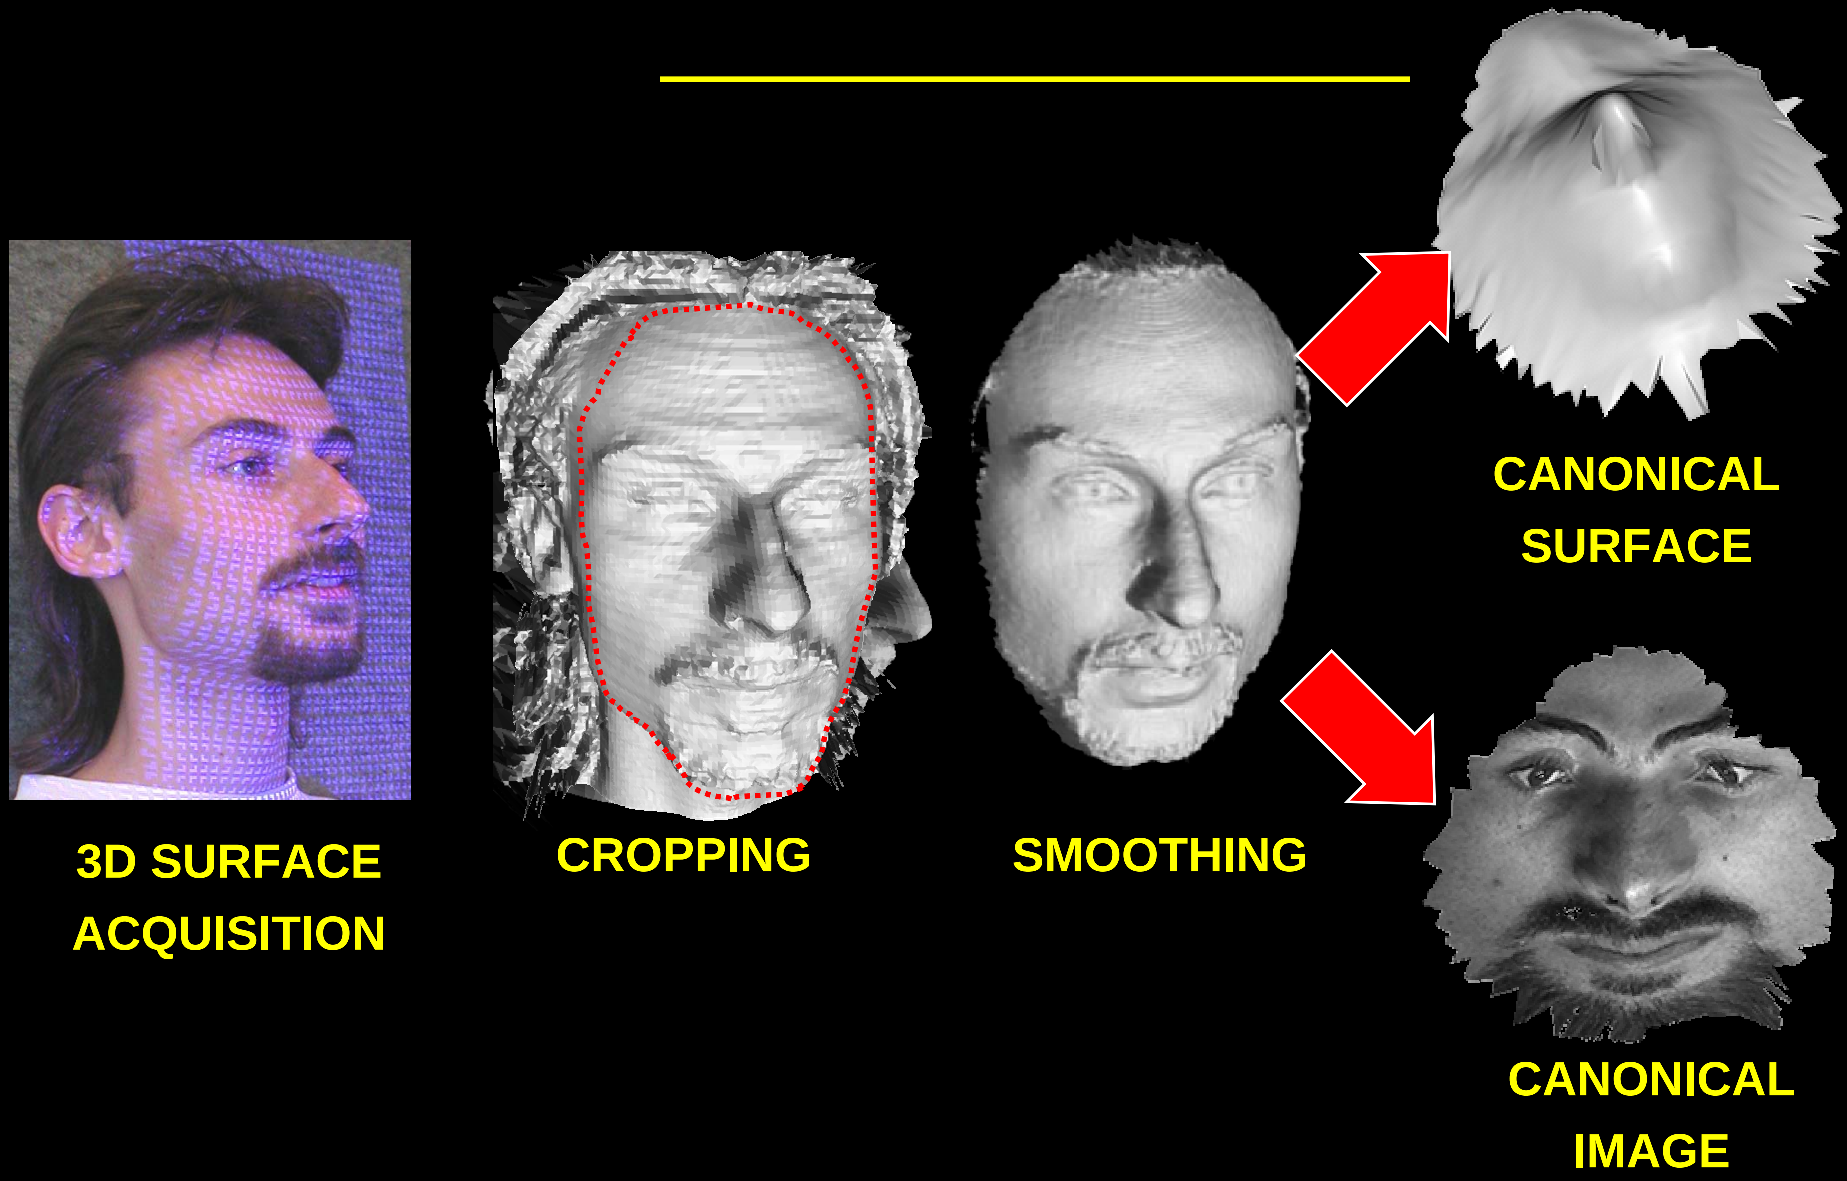

TWINS TEST I

- Recognizing **twins**, a challenging test for face recognition.

SURFACE MATCHING

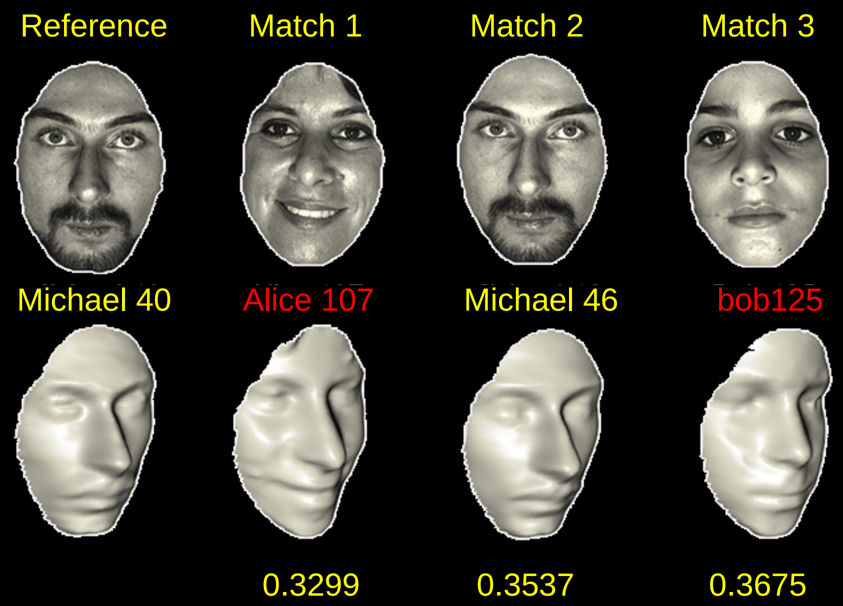

Facial surfaces as rigid objects -> inaccurate.

CANONICAL FORM MATCHING

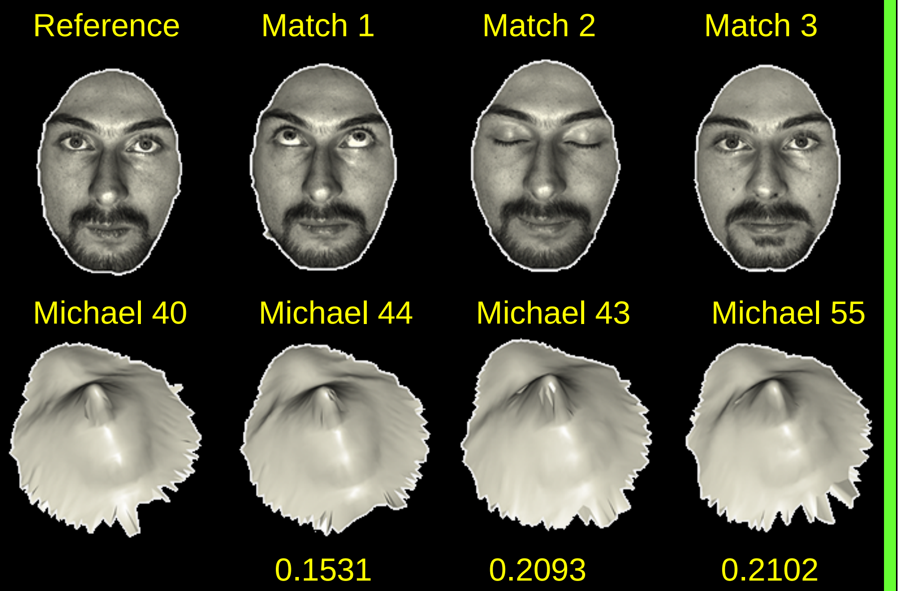

Canonical forms tell apart identical twins.

A. Bronstein, M. Bronstein and R. Kimmel, 2003

Twins test II

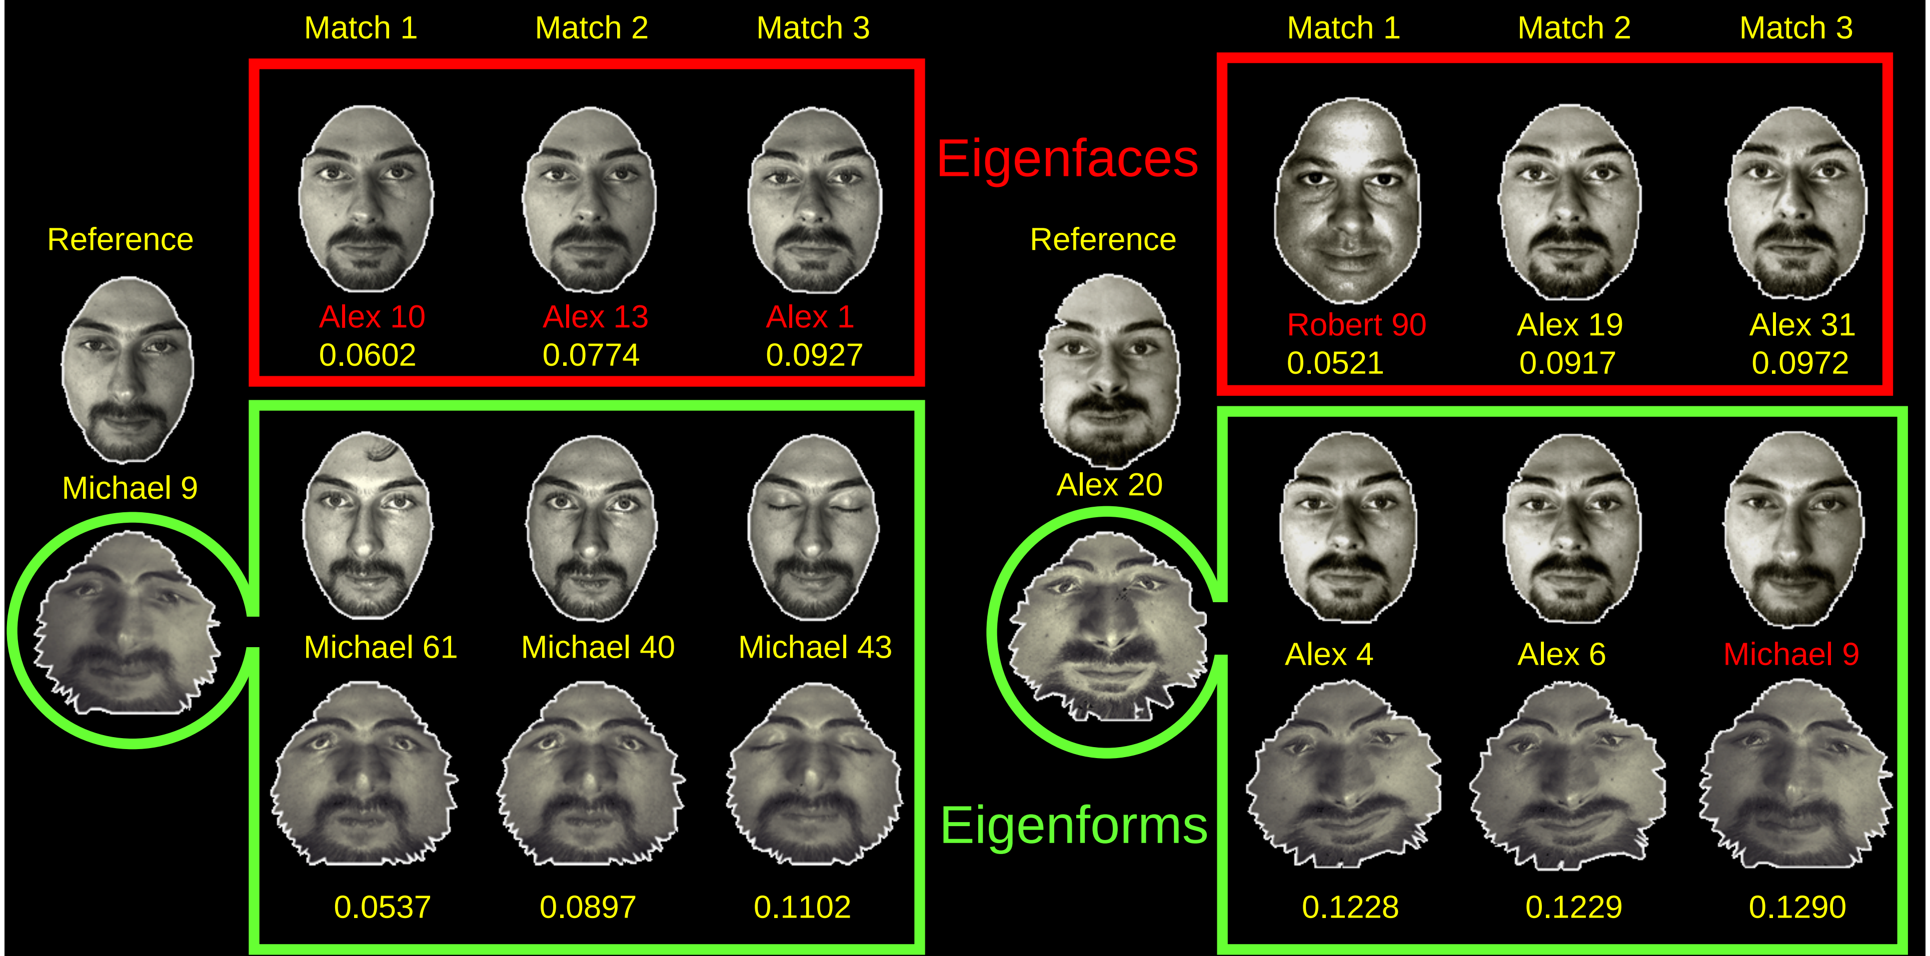

A. Bronstein, M. Bronstein and R. Kimmel, "3D face recognition using geometric invariants", 2003

Invariance to the effects of pose on the intensity

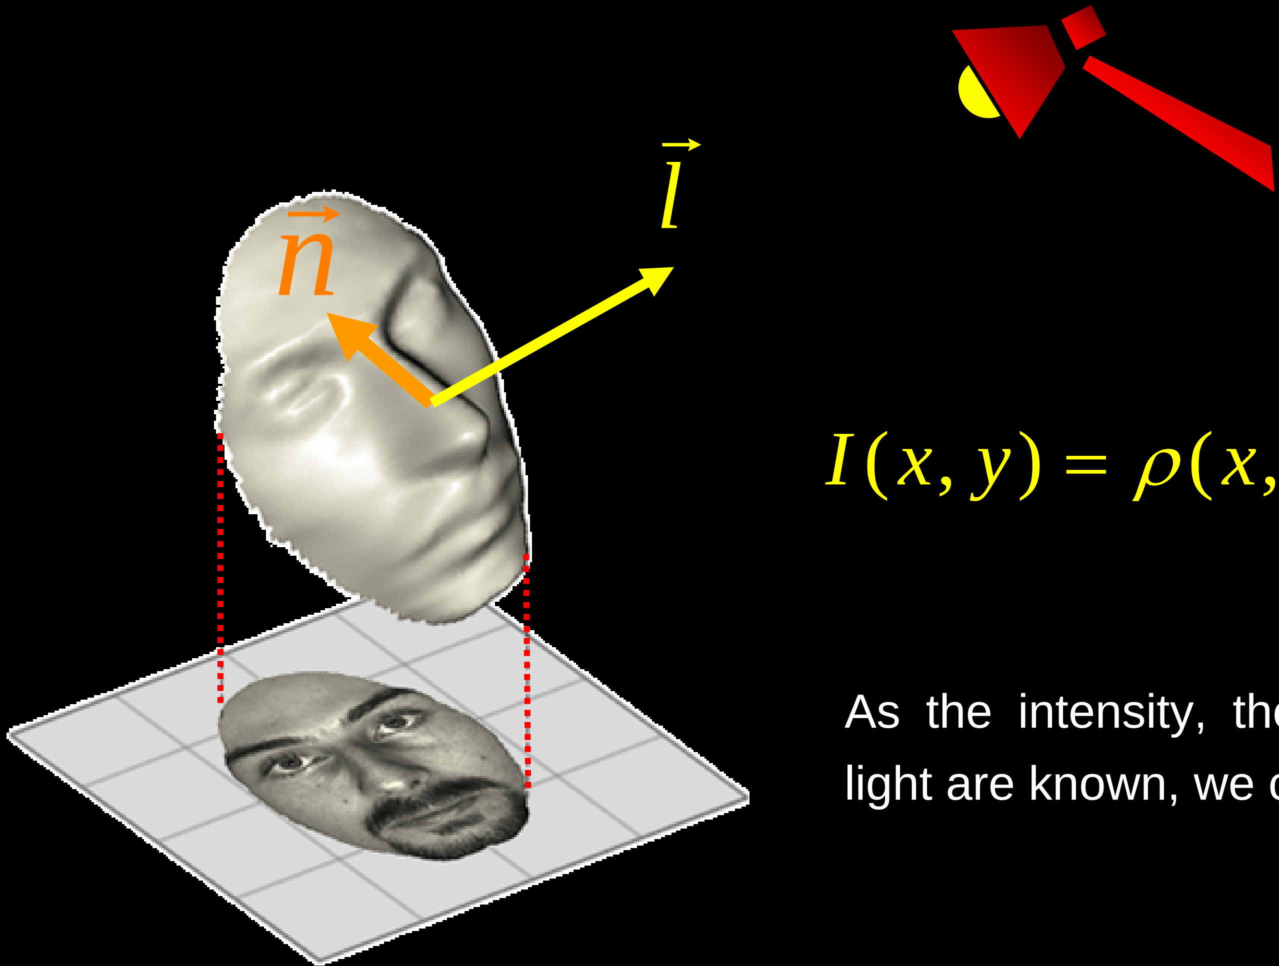

$$I(x,y) = \rho(x,y)\left\langle \vec{l}, \vec{n}(x,y) \right\rangle$$

As the intensity, the normal (surface), and light are known, we can solve for the albedo

A. Bronstein, M. Bronstein, E. Gordon, and R. Kimmel, 2003

Do we need the facial surface?

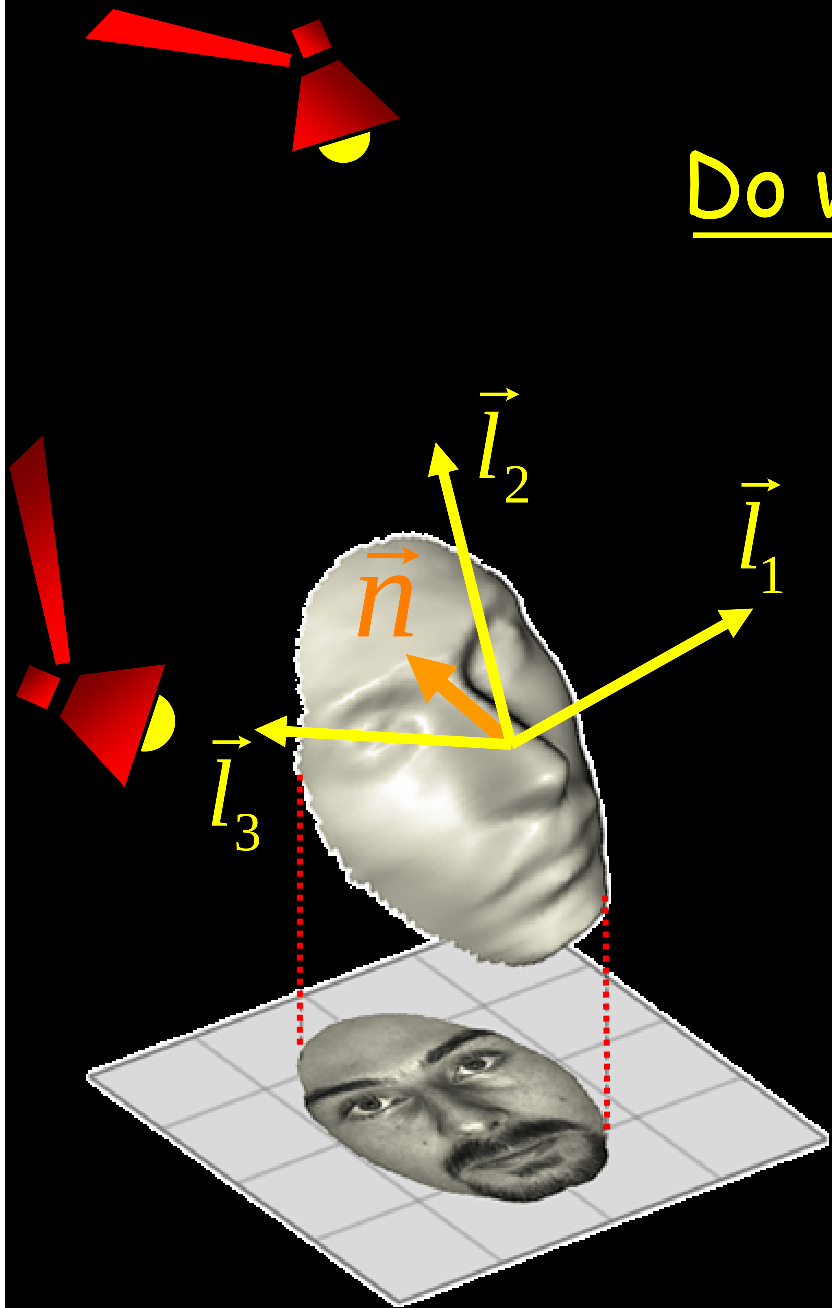

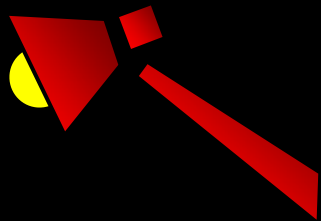

$$I_1(x,y) = \rho(x,y)\left\langle \vec{l}_1, \vec{n}(x,y) \right\rangle$$

$$I_2(x,y) = \rho(x,y)\left\langle \vec{l}_2, \vec{n}(x,y) \right\rangle$$

$$I_3(x,y) = \rho(x,y)\left\langle \vec{l}_3, \vec{n}(x,y) \right\rangle$$

Photometric Stereo:

- Compute the normal from 3 images:
  same camera different light sources.
- The normal is enough for all computations.
- No need to integrate the surface
  using Poisson solvers.

A. Bronstein, M Bronstein, R. Kimmel and A. Spira., ECCV 2004

In the news

- CNN-news, W-NBC, Rueters, Washington-Post, Channel-2, Haaretz, Maariv, Yahoo-news, SIAM news, and more than 60 other TV and newspapers all over the world…

CNN.com/TECHNOLOGY

SEARCH The Web CNN.com Search ENHANCED BY Google

Home Page
World
U.S.
Weather
Business at CNNMoney
Sports at SI.com
Politics
Law
Technology
Science & Space
Health
Entertainment
Travel
Education
Special Reports

Visit the international edition. Click here

SERVICES
Video
Newswatch
E-Mail Services
CNN To Go

SEARCH
Web CNN.com
Search ENHANCED BY Google

Twins crack face recognition puzzle

Monday, March 10, 2003 Posted: 9:51 AM EST (1451 GMT)

HAIFA, Israel (Reuters) -- For a fleeting moment, Mohamed Atta appeared on an airport security camera minutes before he boarded one of the planes which crashed into the World Trade Center on September 11, 2001.

Was there any way the camera or its operator would have been able to identify Atta as a suspect before he hijacked and flew the first of two planes into the twin towers?

Israelis Michael and Alex Bronstein think they have the answer.

Story Tools
SAVE THIS
EMAIL THIS
PRINT THIS
MOST POPULAR

The computer whiz-kids -- 22-year-old identical twins almost impossible to tell apart -- have applied a new technology to recognizing faces in a way that may yet revolutionize international security.

"I said it to them as a joke: if you succeed in building a system that can distinguish between the two of you, you'll get (a grade of) 100," said the twins' professor, Ron Kimmel of the Technion Institute in Haifa.

"They succeeded and got 100. They are brilliant."

The technology scans and maps the human face as a three-dimensional surface, providing a far more accurate reference for identifying a person than current systems, most of which rely on two-dimensional images, Kimmel said.

The product can potentially meet a wide range of security needs in a world shaken by the September 11 attacks and a series of bombings blamed on Osama Bin Laden's al Qaeda network, of which Atta was a suspected member.

Kimmel and one of his former pupils, Assi Elad, had already developed the algorithms used as building-blocks for the face-recognition system. The Bronstein twins constructed a 3-D scanner, together with engineer Eyal Gordon

Why flat embedding?

- Is there another space that better captures the face intrinsic geometry yet still provides the convenient comparison property of a flat embedding space?

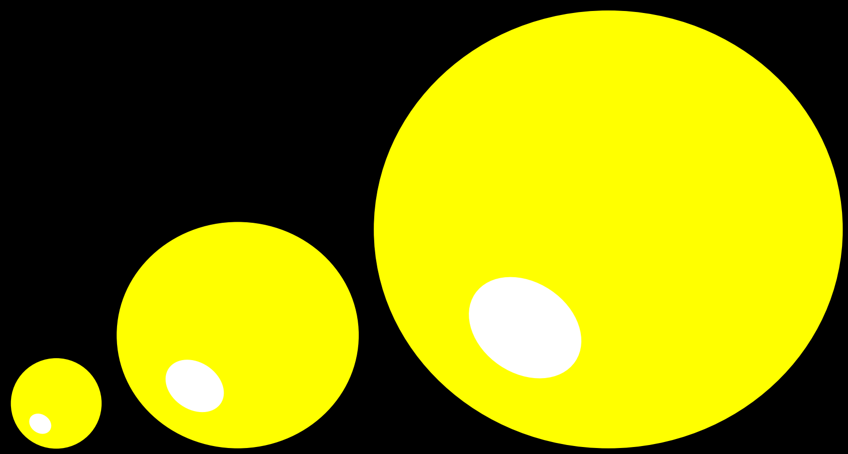

$$S^n$$

Embedding in S^2

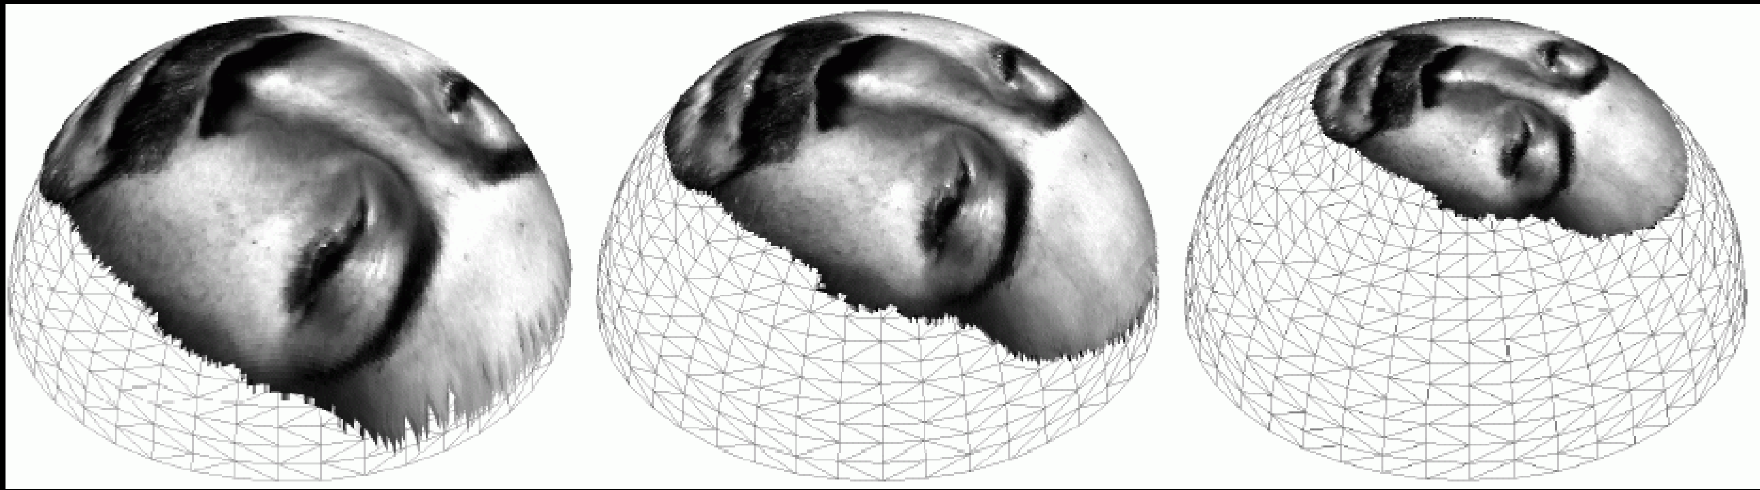

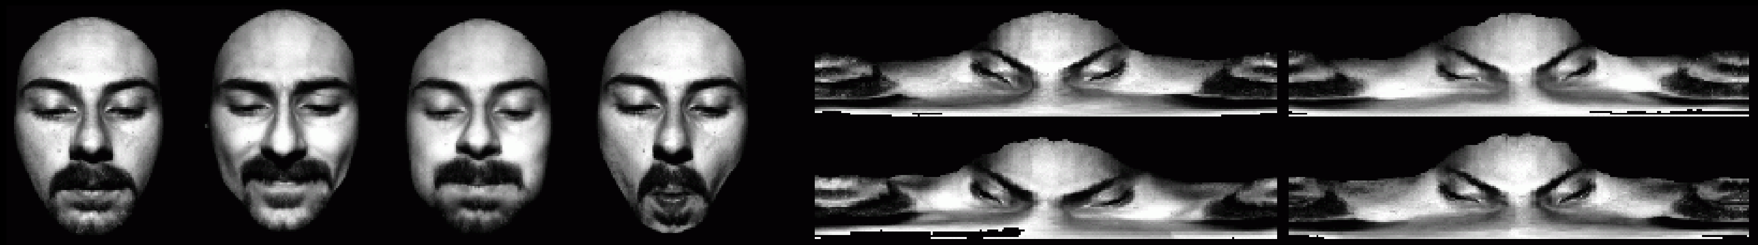

Spherical embedding error

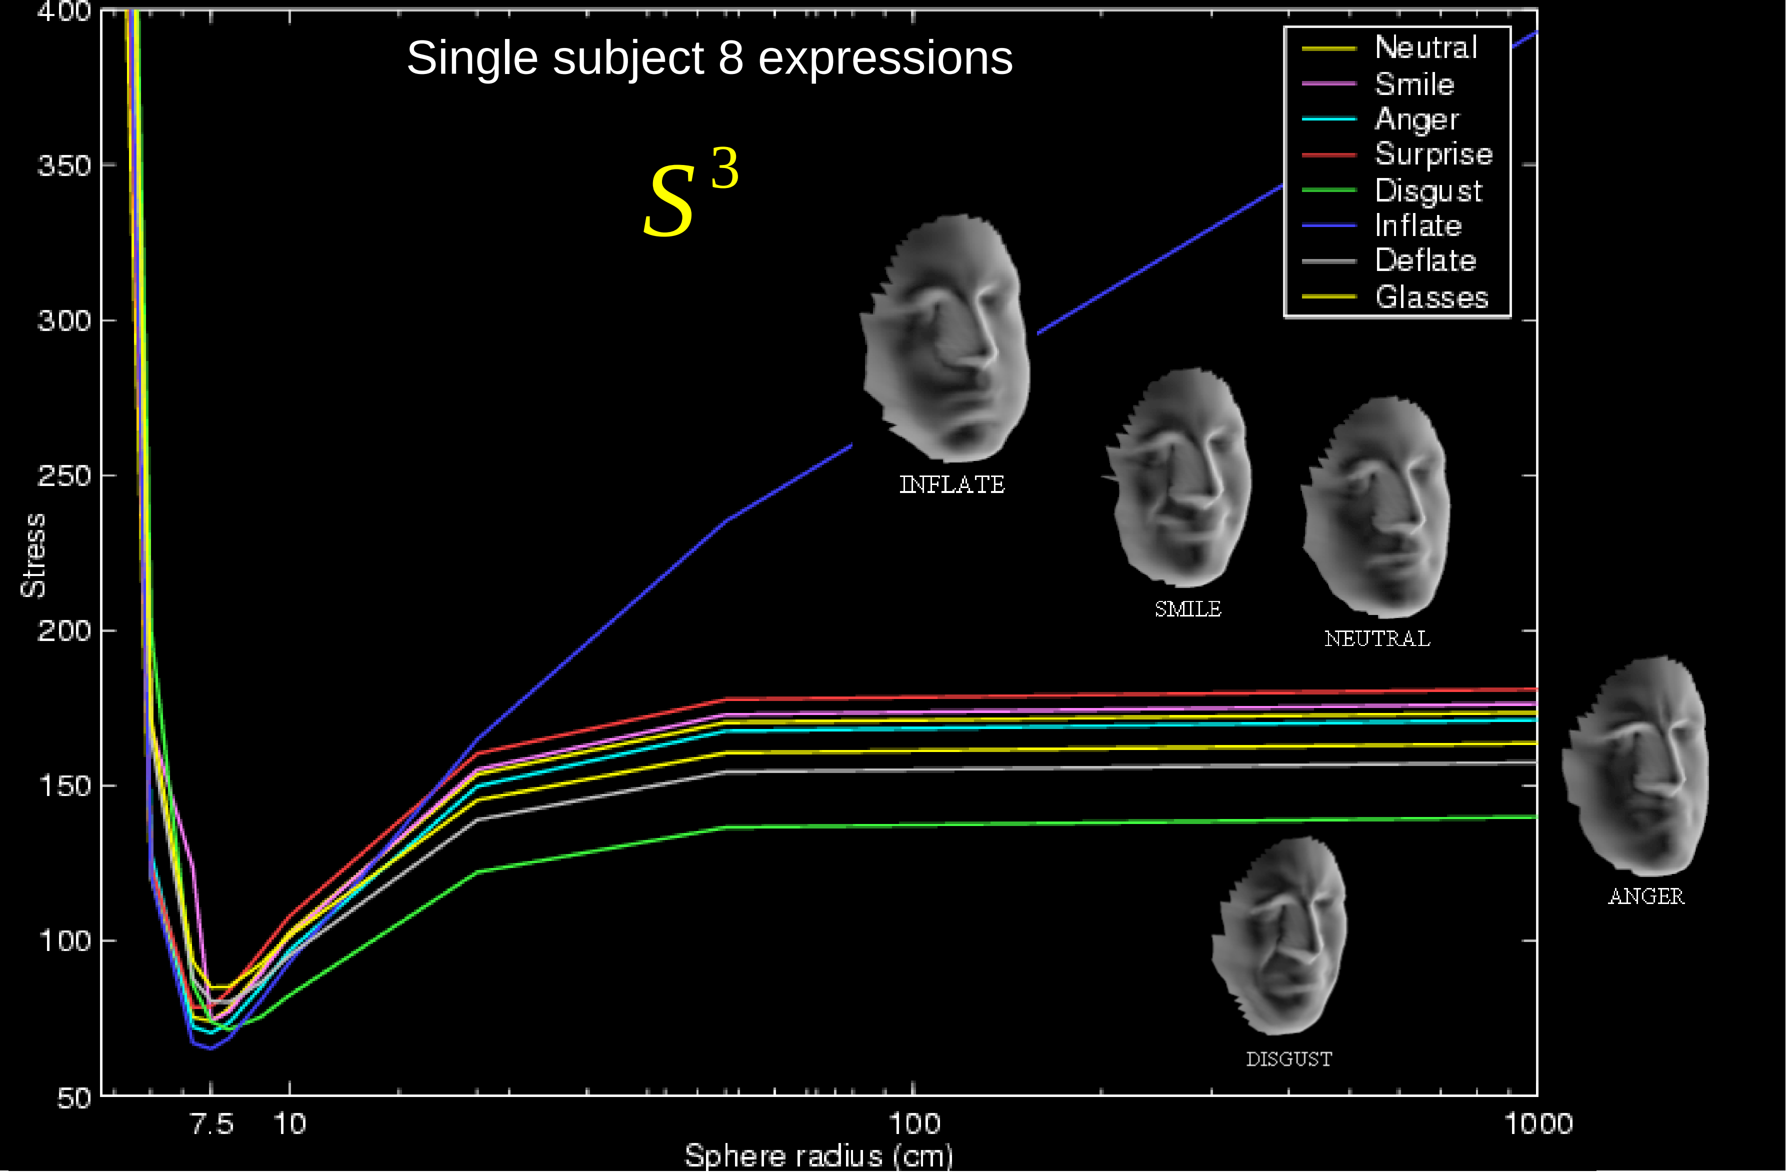

Spherical embedding error

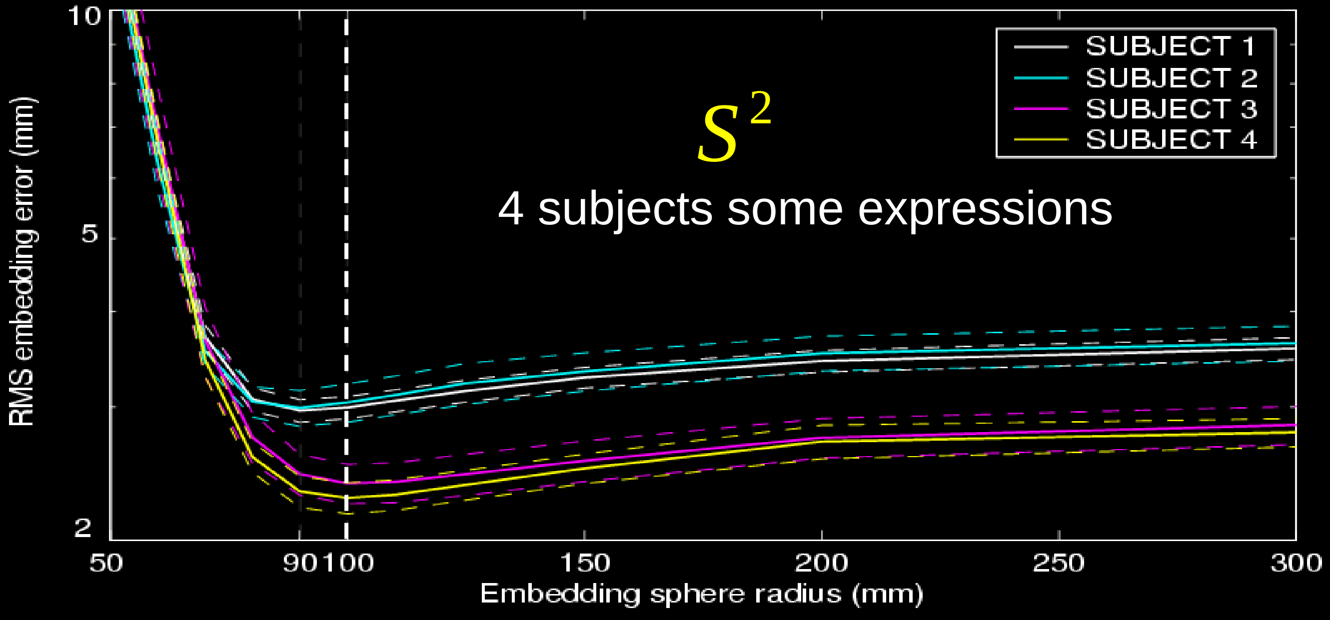

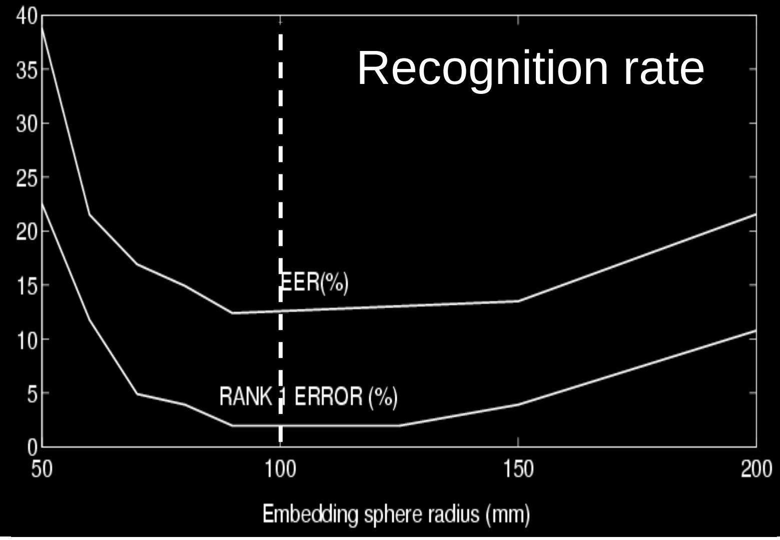

FLAT EMBEDDING SAMPLED SURFACES

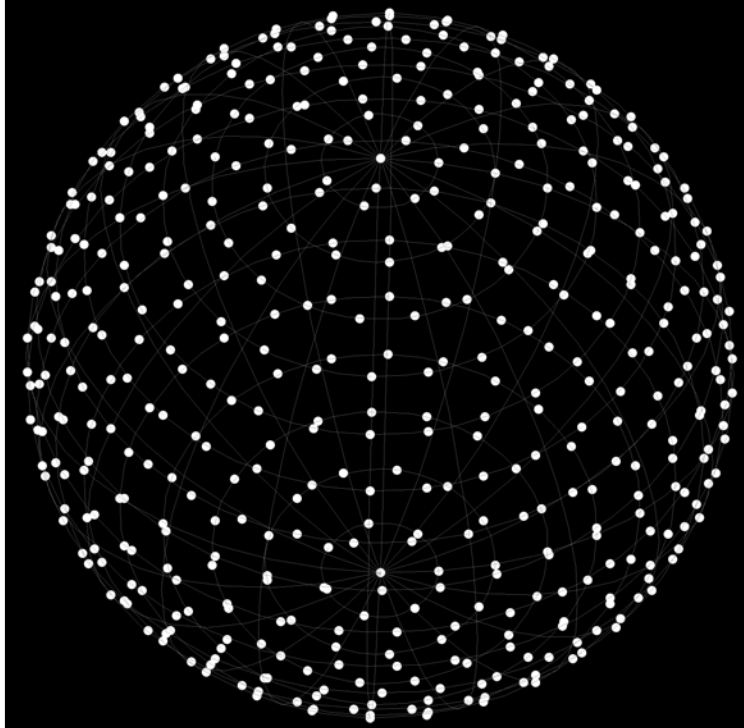

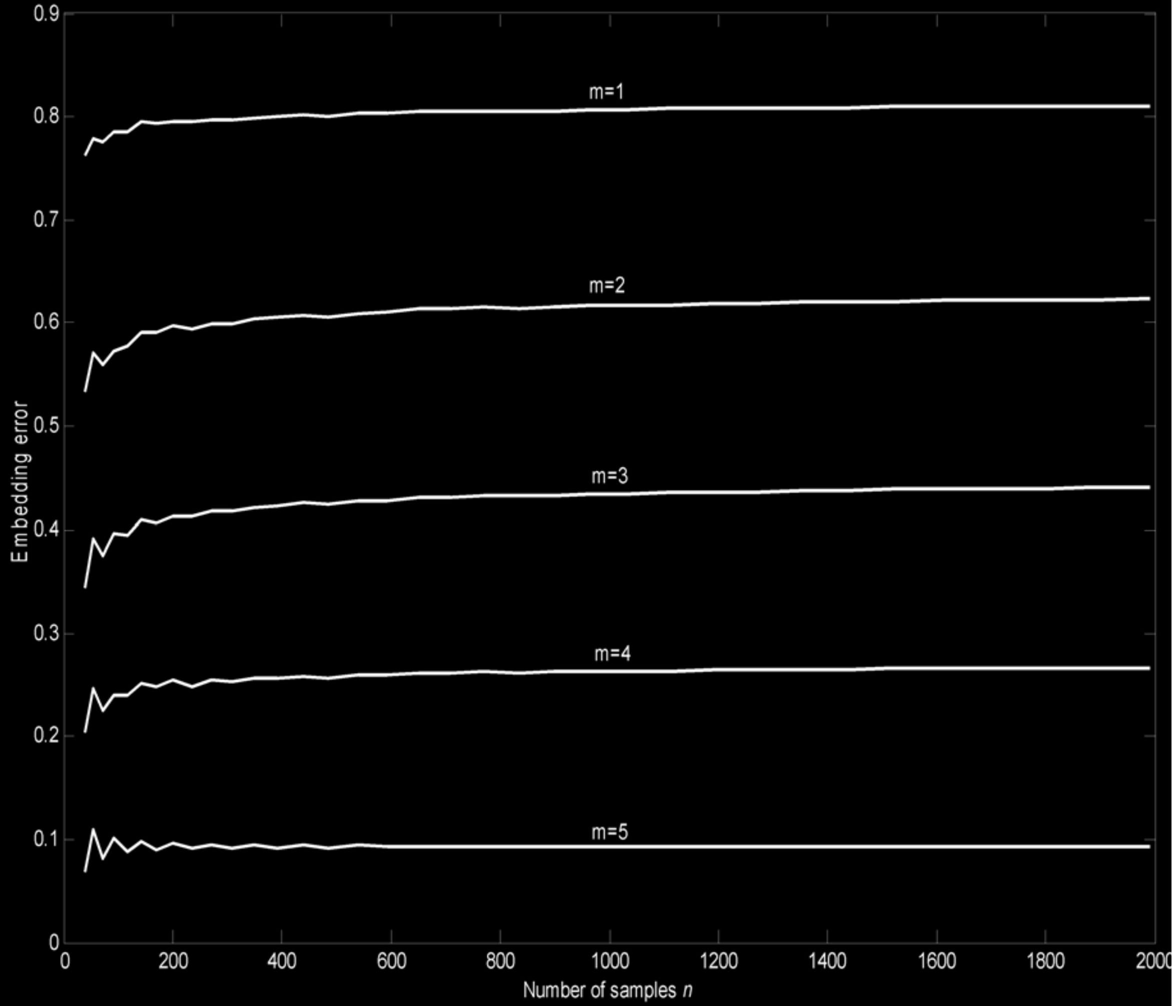

Embedding a sphere in $\mathbf{R}^m$. Asymptotic behavior with number of samples: embedding error decreases as embedding dimension *m* grows.

A. Bronstein, M. Bronstein and R. Kimmel, “A geometric framework for 3D face recognition“, 2002

Flat embedding?

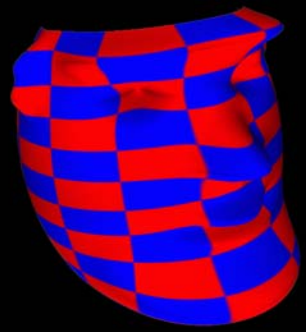

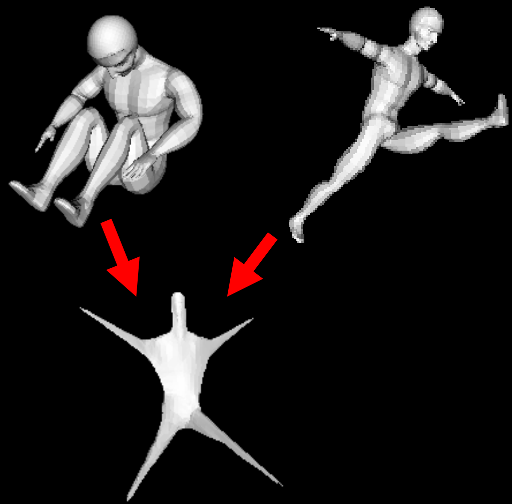

Beyond Bourgain theorem (1985):

For n points metric space there is a flat embedding with distortion O(log n)

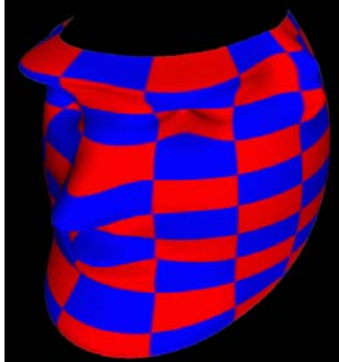

- Does a smooth surface has a flat embedding with bounded distortion error?

Partial answers:

Tight bound for 1D curves (spectral distortion)

Tight numerical bound for spheres (spectral distortion)

And numerical evidence for faces.

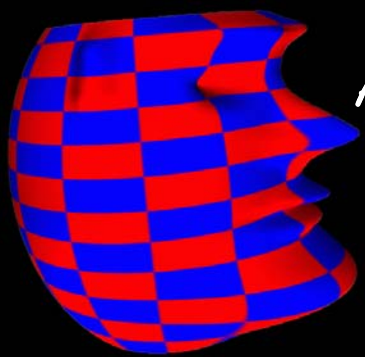

Next goal: A sampling theorem for isometric surfaces

Embedding arbitrary metric spaces in flat domains

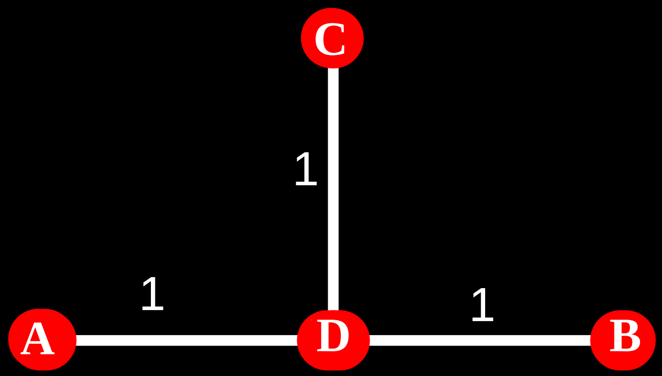

	A	B	C	D
A	0	2	2	1
B	2	0	2	1
C	2	2	0	1
D	1	1	1	1

Personal communication N. Linial

Embedding arbitrary metric spaces in flat domains

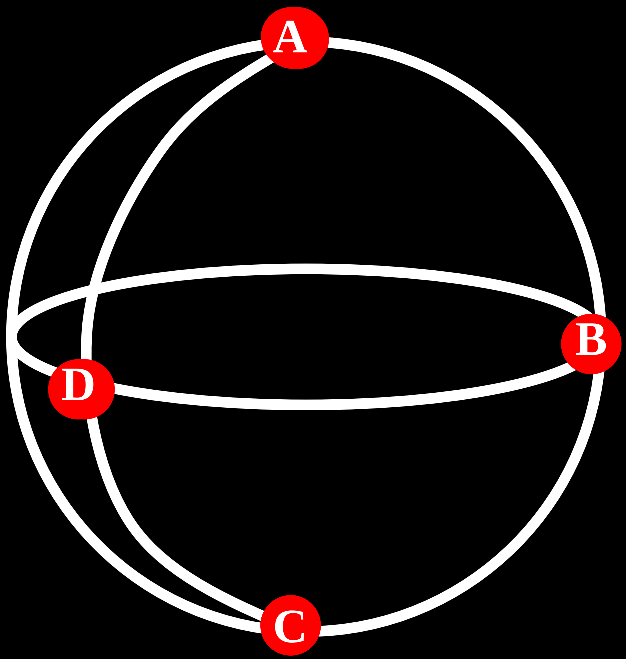

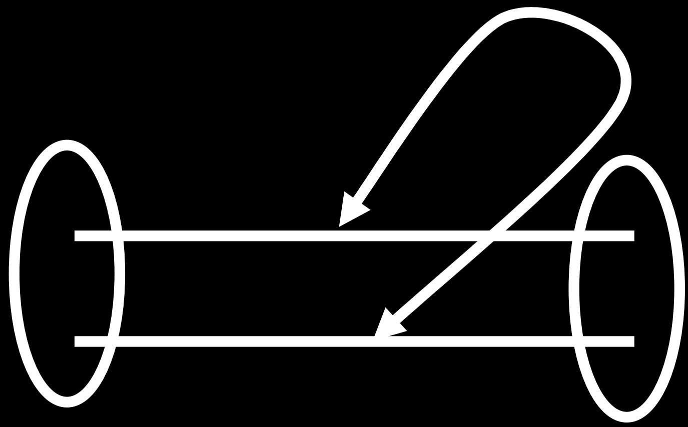

Personal communication N. Linial

Released: Oct. 2003
Publisher: Springer

ISBN: 0387955623

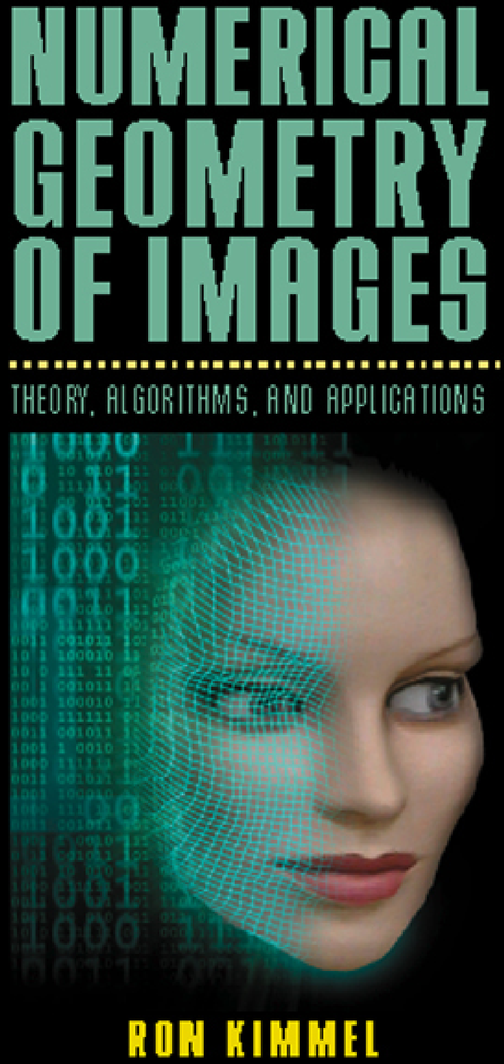

Future directions:

- Better models
- Better numerical methods
- Better coding
- Better hardware
- Better analysis of the large scale problem

Front cover rendered by A&M Bronstein


3DFACE