Modeling Deformable Surfaces from Single Videos

P. Fua
CVLab, EPFL, Switzerland
http://cvlab.epfl.ch/

Talk Outline

- **2D Deformable Surfaces**
  - Problem Formulation.
  - Fast Matching.
  - Robust Optimization Scheme.
  - Illumination Correction.

- **3D Deformable Surfaces**
  - Linear Formulation.
  - Inextensible surfaces.
  - Sharply folding surfaces.
  - Eliminating the reference image.

ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

CVLab

2D Deformable Surfaces

+

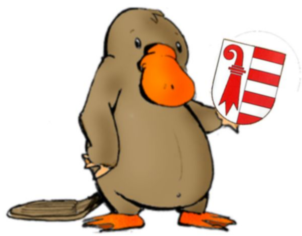

=

Estimating:

- Deformations
- Lighting parameters
- Occlusions

Pilet et al., IJCV 2008

2D Deformable Surfaces

- Problem Formulation
- Fast Matching
- Robust Optimization
- Lighting Correction

ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

CVLab

Problem Formulation

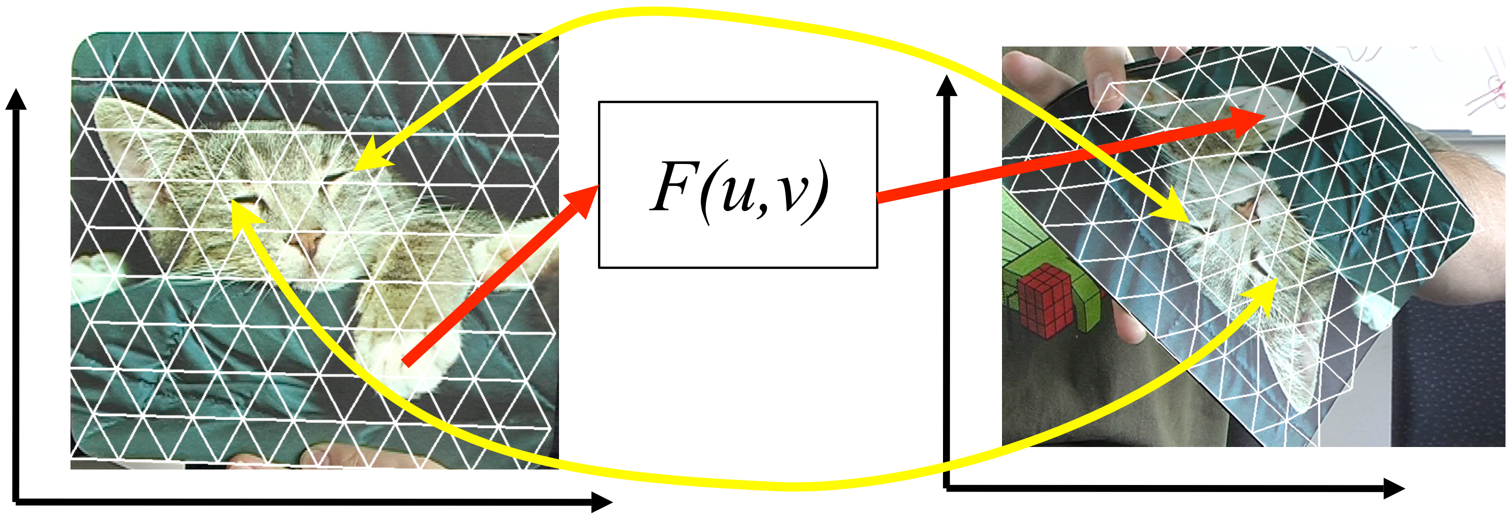

- Input:
  - Correspondences between a **reference** and **input image.**
  - No a priori pose information.
- Output:
  - A mapping **F** from model to input image.

Challenges

Non-rigid deformation without a priori pose:

- High dimensionality (200+ DOF)
- Large search space
- Wide baseline matching

Real-time requirements:

- Fast optimization scheme
- Fast matching

Deformable Model

Wide Baseline Matching

Regularization Term

$$\varepsilon(S) = \varepsilon_C(S) + \lambda_D \varepsilon_D(S)$$
$$S = (X, Y)$$

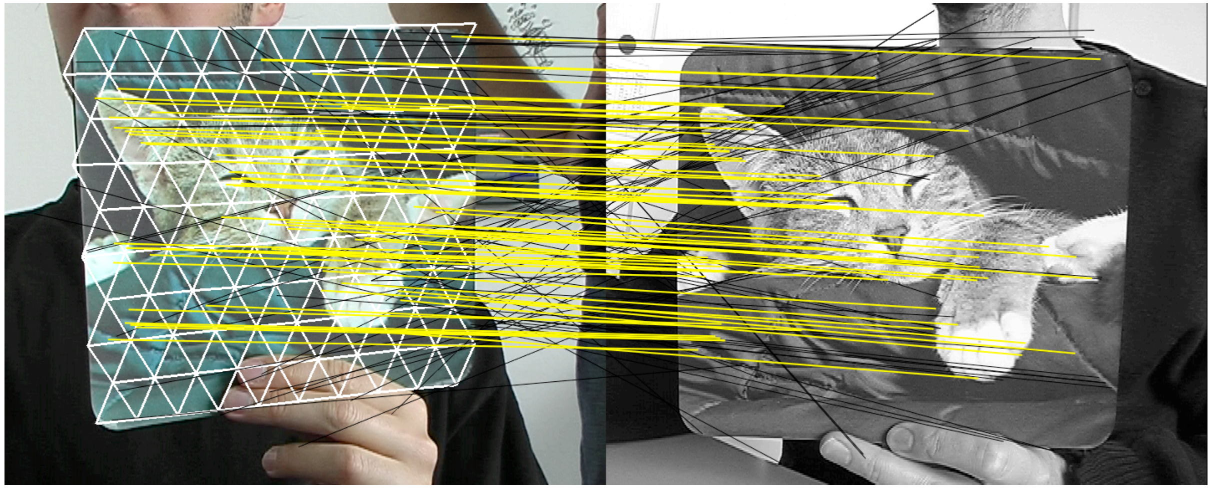

Reference Image Input Image

ε_D Regularization Term

Quadratic function vertex coordinates:

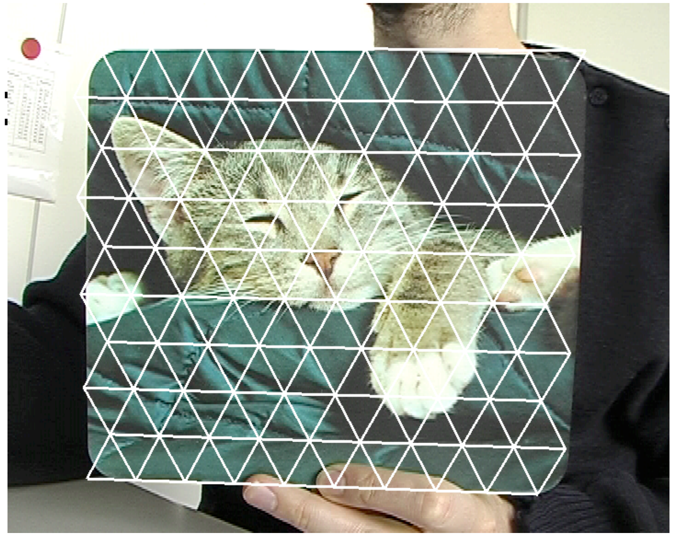

$$\varepsilon_D(S) = \frac{1}{2}\left(X^T K X + Y^T K Y\right)$$

- penalizes non uniform scaling;
- penalizes excessive bending;
- allows perspective distortion;
- allows smooth surface deformation.

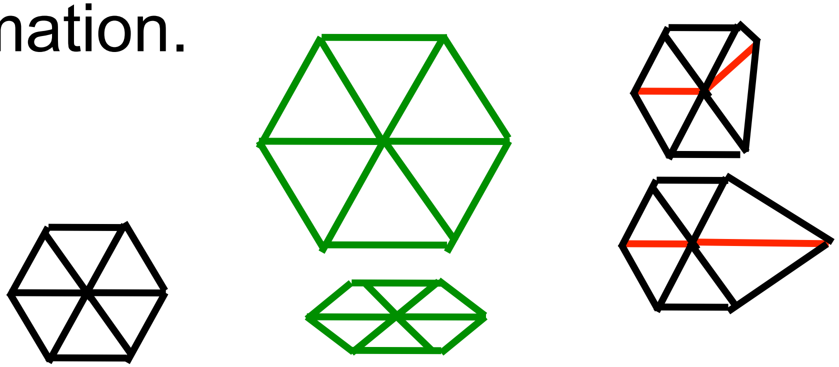

ε_C Correspondence Term

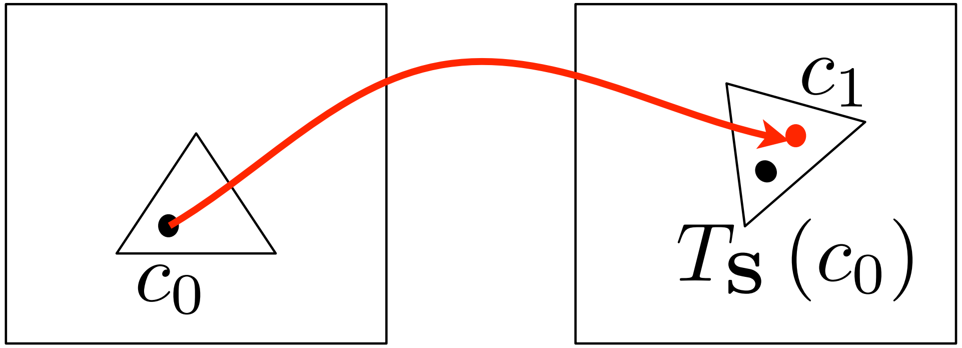

$$\varepsilon_C = -\sum_{c \in C} \left\| c_1 - T_{\mathbf{S}}(c_0) \right\|^2$$

ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

Real-Time Augmentation

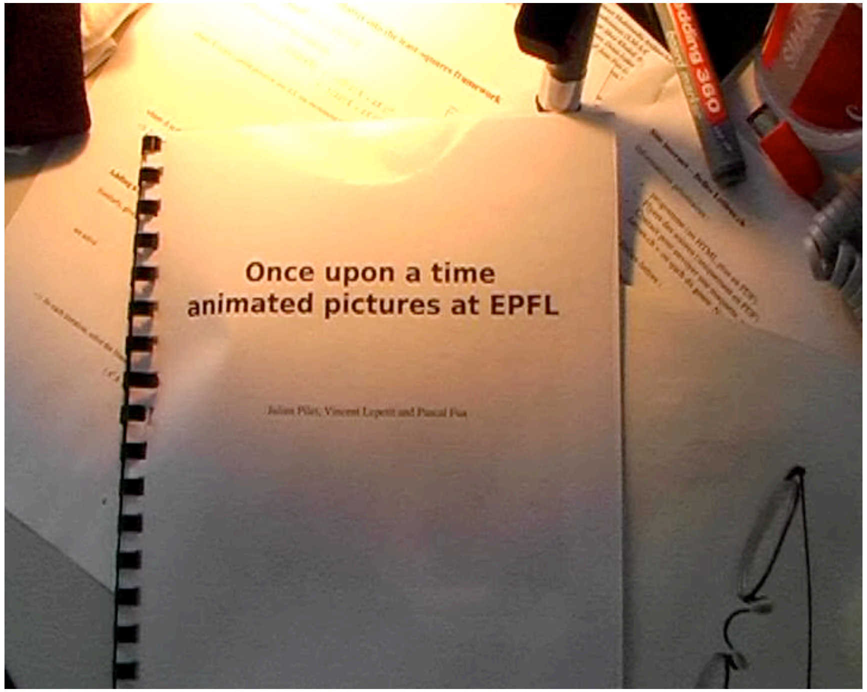

Key Ingredients

- Classification-based approach to matching.
- Robust minimization scheme.
- Intensity ratios for illumination correction.

2D Deformable Surfaces

- Problem Formulation
- Fast Matching
- Robust Optimization
- Lighting Correction

ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

CVLab

ε_C Correspondence Term

Standard approach:

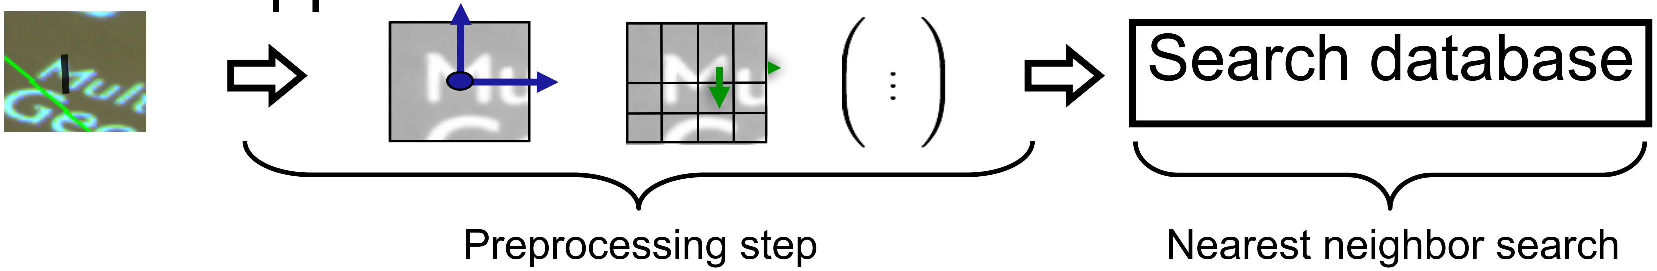

Classification-based approach:

- One class per keypoint.
- Shifts computational burden to offline training.

ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

Binarized Tests for Keypoint Matching

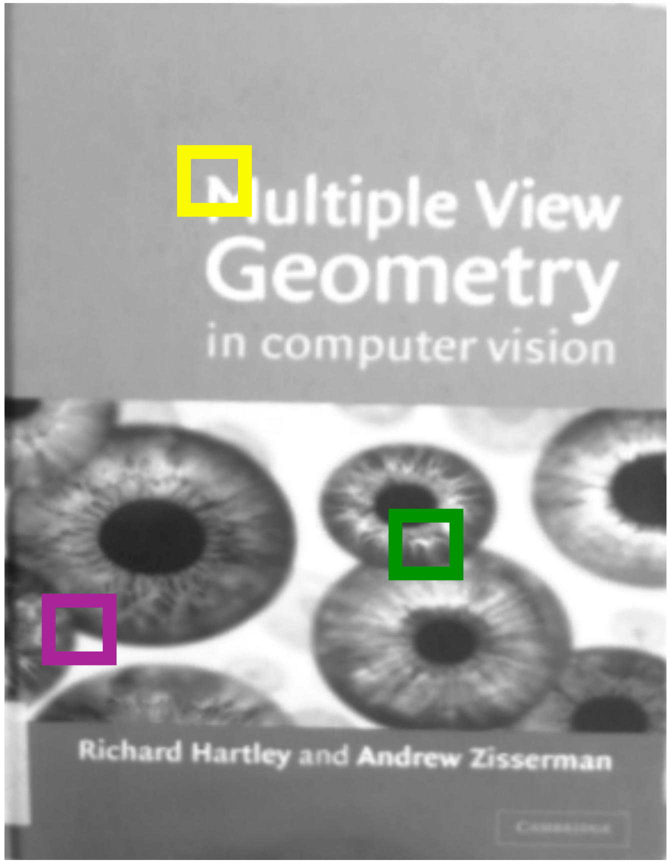

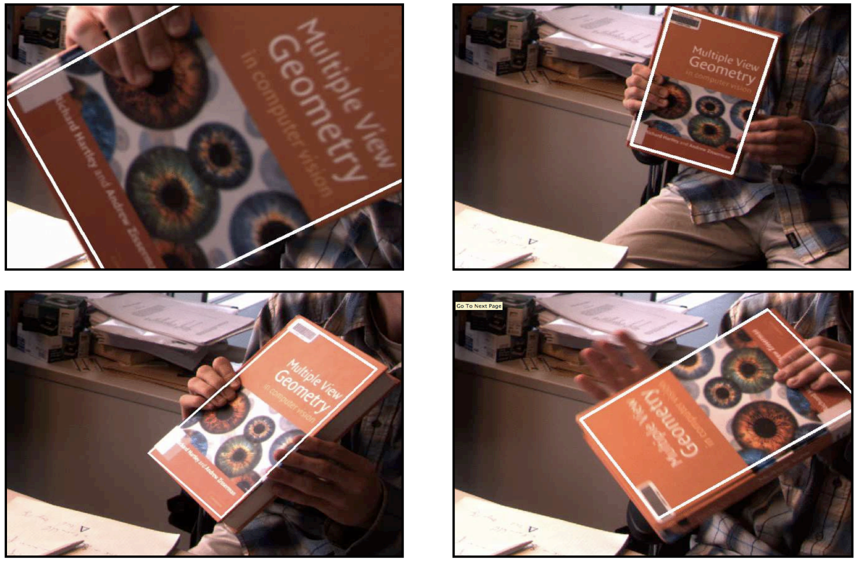

Randomized Tree

Generic tree: The nodes contain simple tests of the form “Is I($\mathbf{m_1}$) > I($\mathbf{m_2}$) ?”

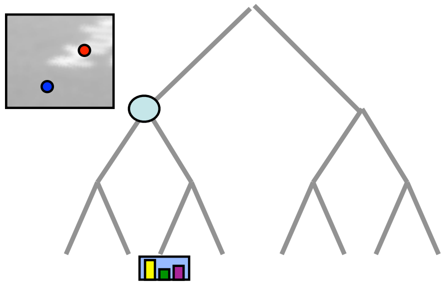

Posterior probabilities can be learnt from:

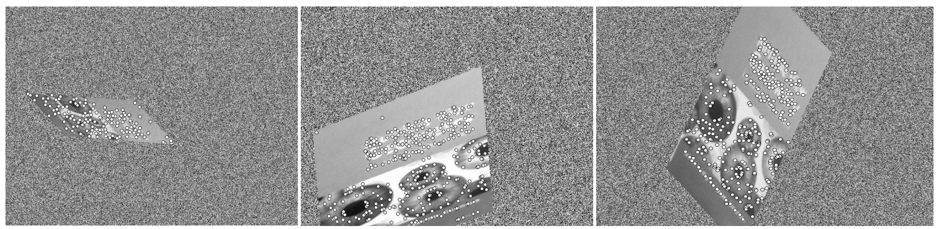

- synthetically warped images

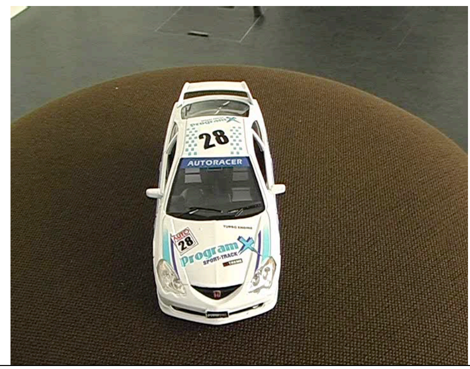

- video sequences

Multiple Trees

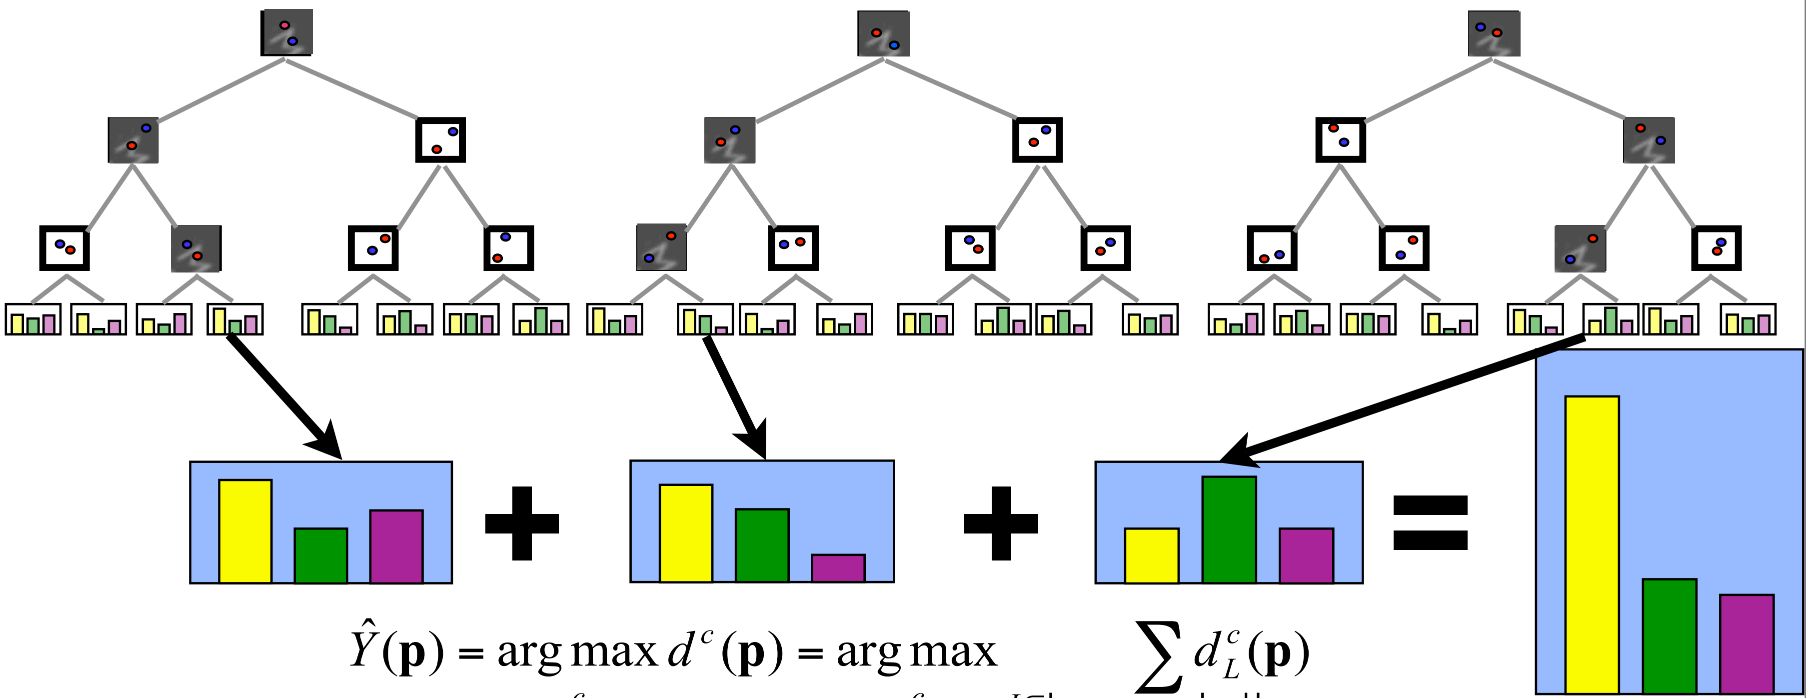

$$\hat{Y}(\mathbf{p}) = \arg\max_c d^c(\mathbf{p}) = \arg\max_c \sum_{L \in \text{leaves reached by } \mathbf{p}} d_L^c(\mathbf{p})$$

Where should the tests be performed?

- Choose locations to maximize information gain.
- Choose locations randomly.

Random vs Optimized Locations

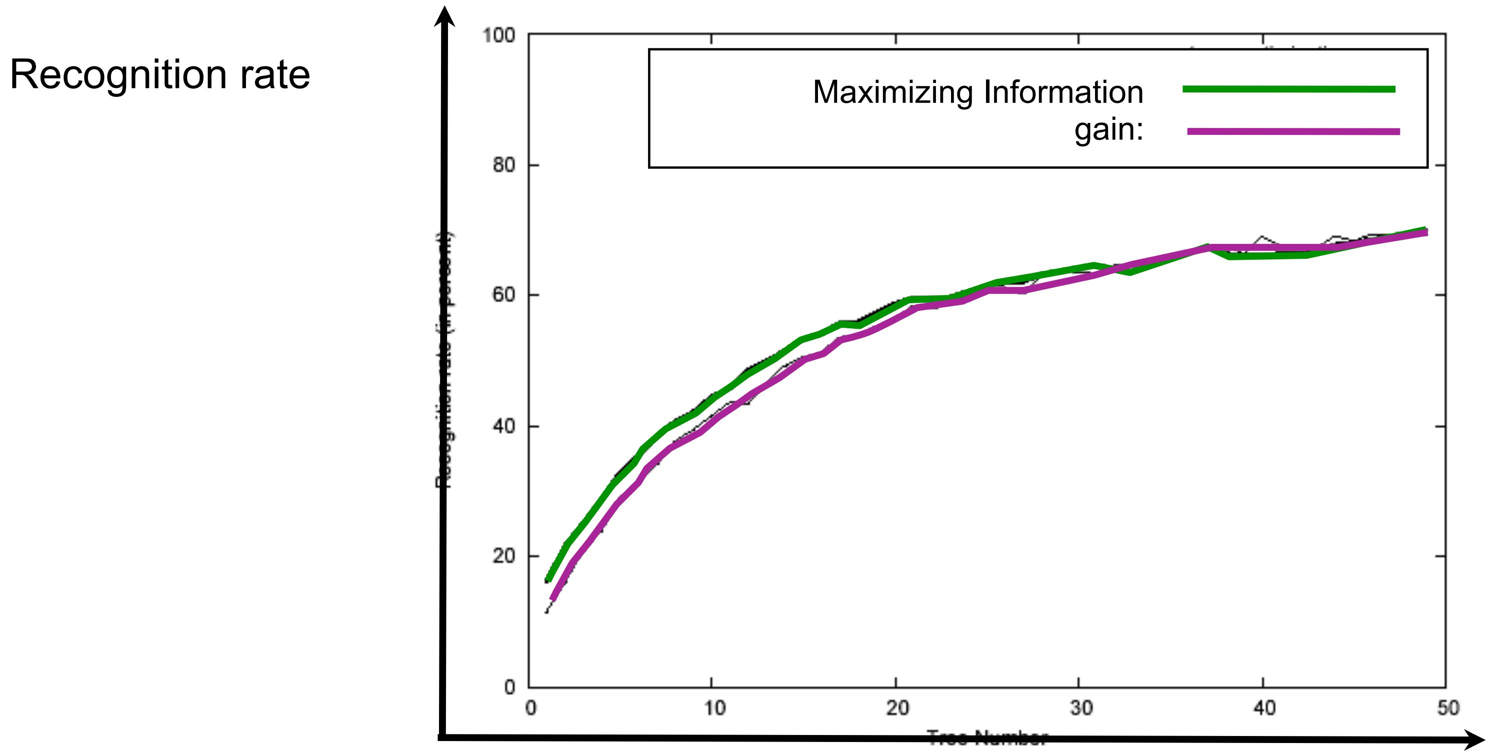

Recognition rates for 200 keypoints.

FERNS: Flattening the Tree

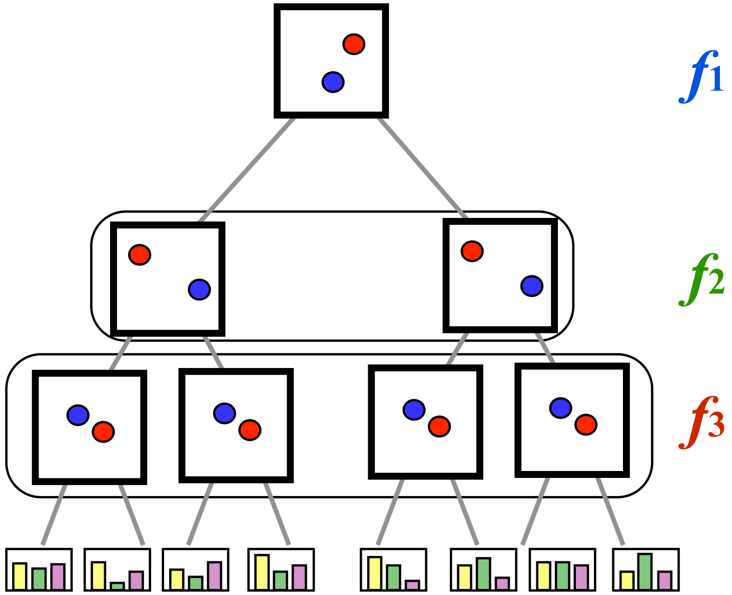

The distributions can be expressed simply, as:

$$P(f_1, f_2, \ldots f_n \mid C = c_i)$$

Results of pixel comparisons (0 or 1)

Class Label

Bayesian Interpretation

We are looking for:

$$P(C = c_i \mid f_1, f_2, \cdots f_n,\ f_{n+1}, \cdots \cdots f_N)$$

proportional to

$$P(f_1, f_2, \cdots f_n,\ f_{n+1}, \cdots \cdots f_N \mid C = c_i)$$

but complete representation of joint distribution infeasible.

Naive Bayesian:

$$\approx \prod_j P(f_j \mid C = c_i)$$

Compromise:

$$\approx \boxed{P(f_1, f_2, \cdots f_n \mid C = c_i)} \otimes \boxed{P(f_{n+1}, \cdots f_{2n} \mid C = c_i)} \otimes \cdots$$

--> probabilities stored in the leaves.

Sum vs Product

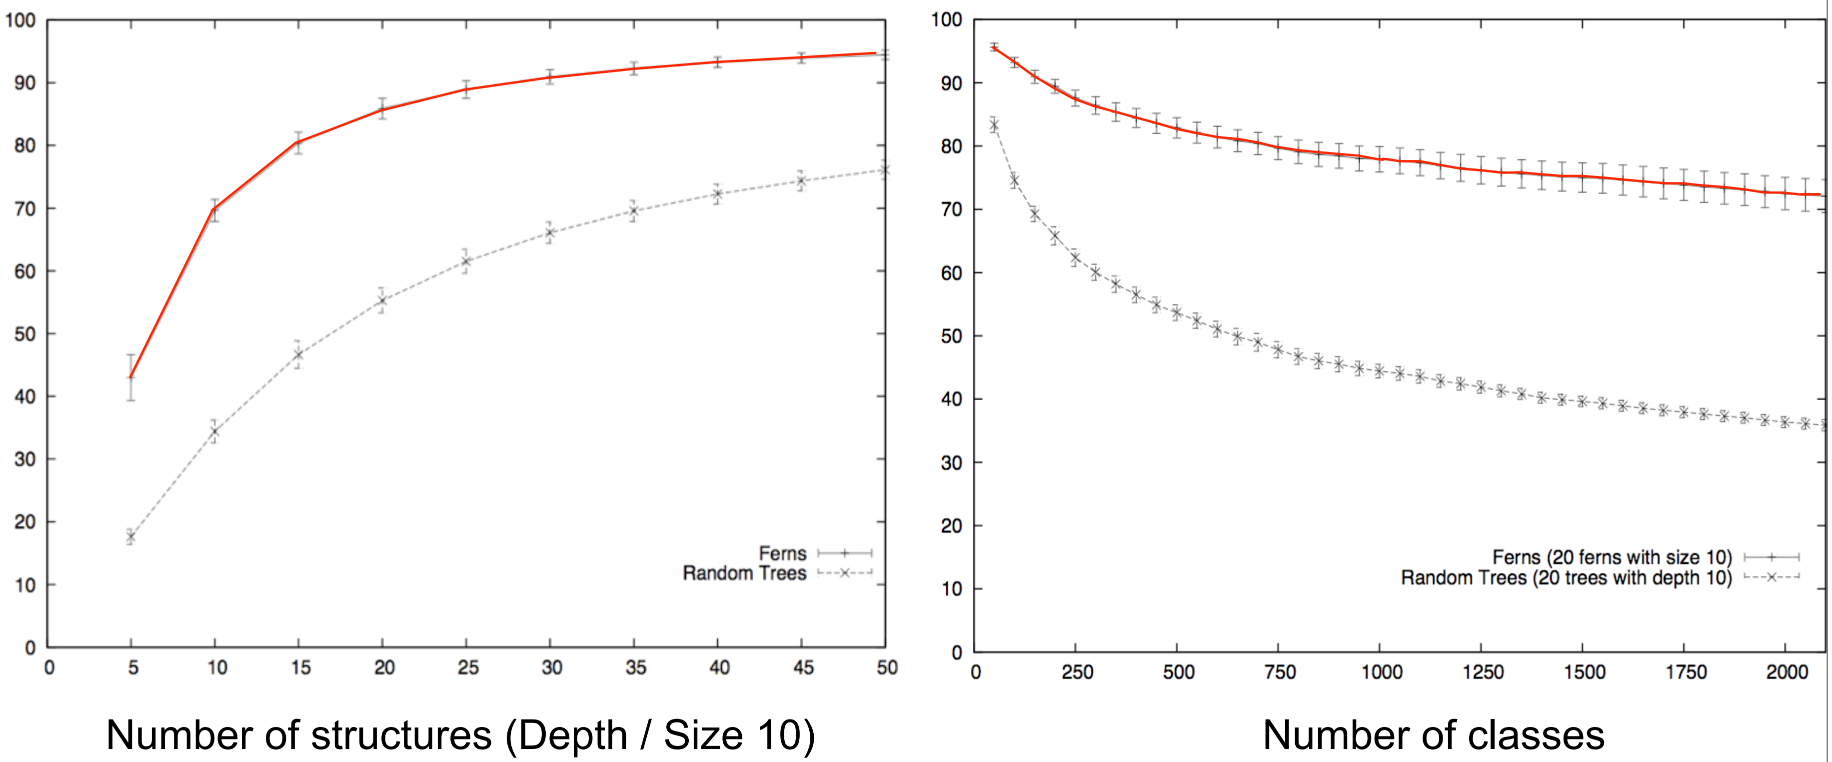

500 classes, no orientation or perspective correction.

Scale and Orientation Invariance

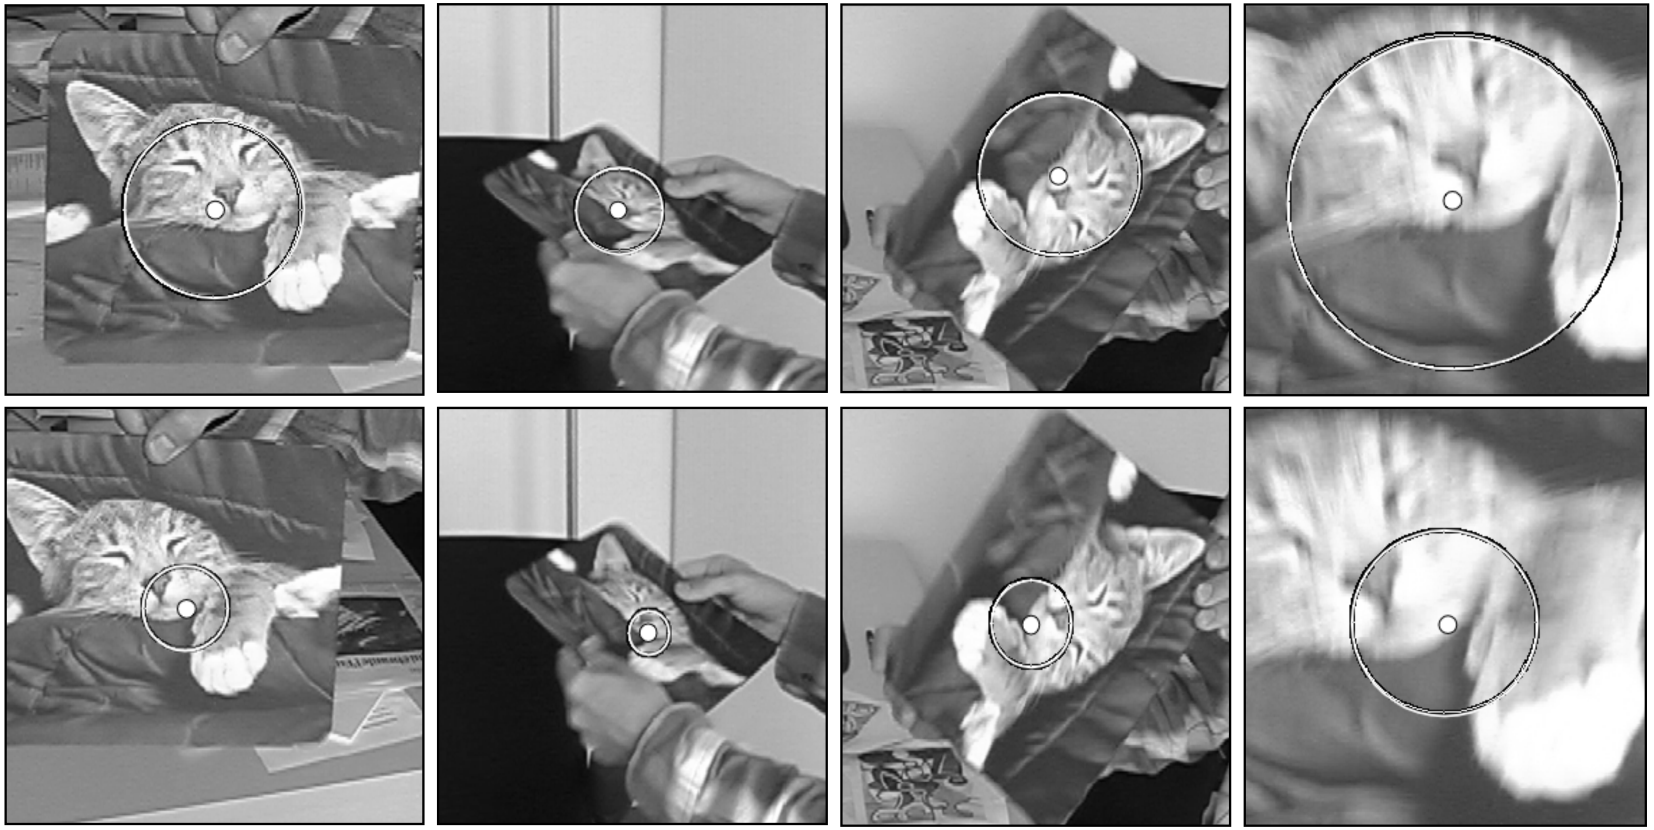

ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

Planar or Not

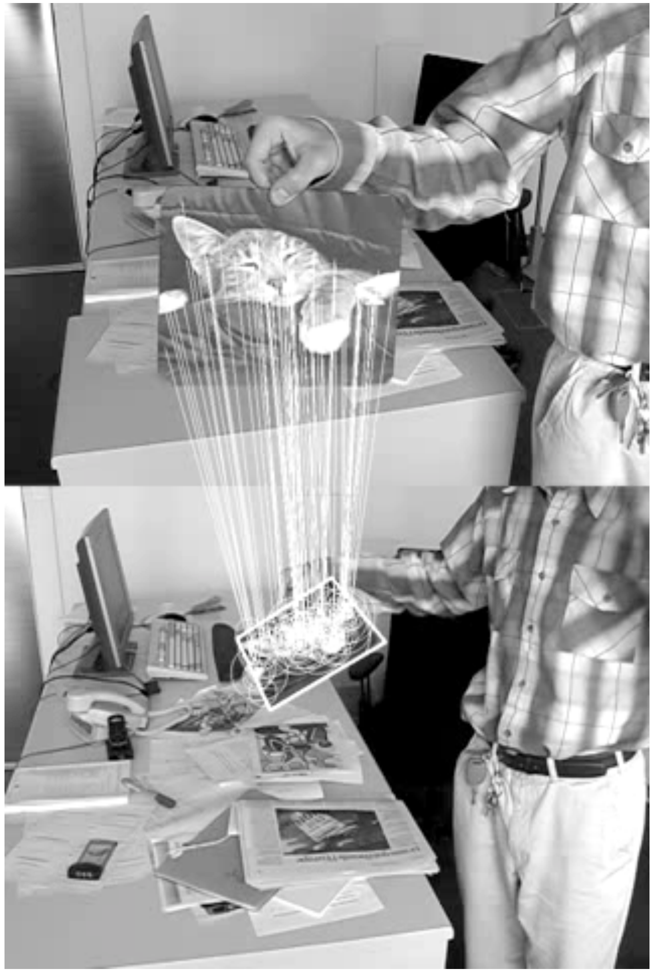

Reference image vs Input Images

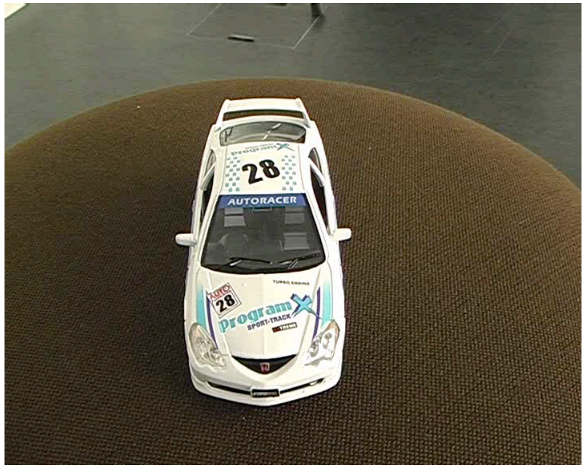

Reference video

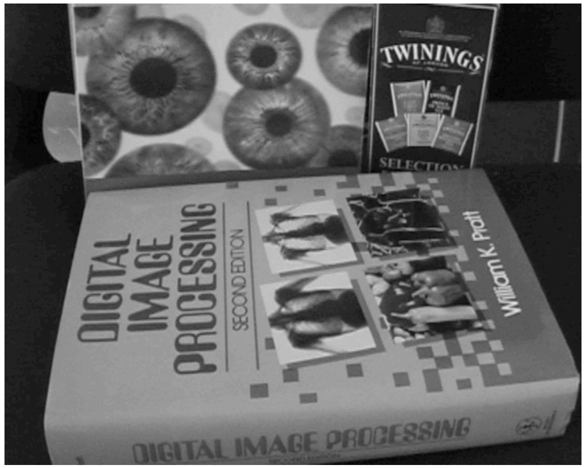

Input Images

BRIEF

Very simple computation that can be seen as computing gradients:

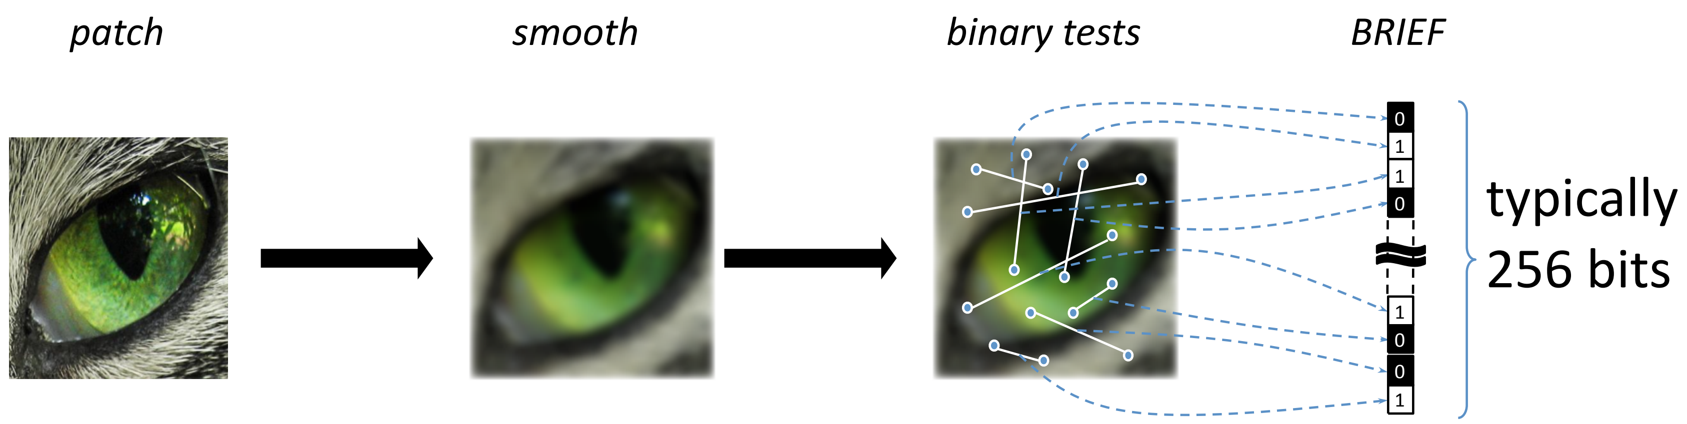

- Most smooth kernels work, even simple box filters.
- 128, 256, or 512 binary tests usually suffice.
- Random arrangment of tests effective iff evenly sampled.
- Not rotationally invariant.

Calonder et al. ECCV'10

Benchmarks Datasets

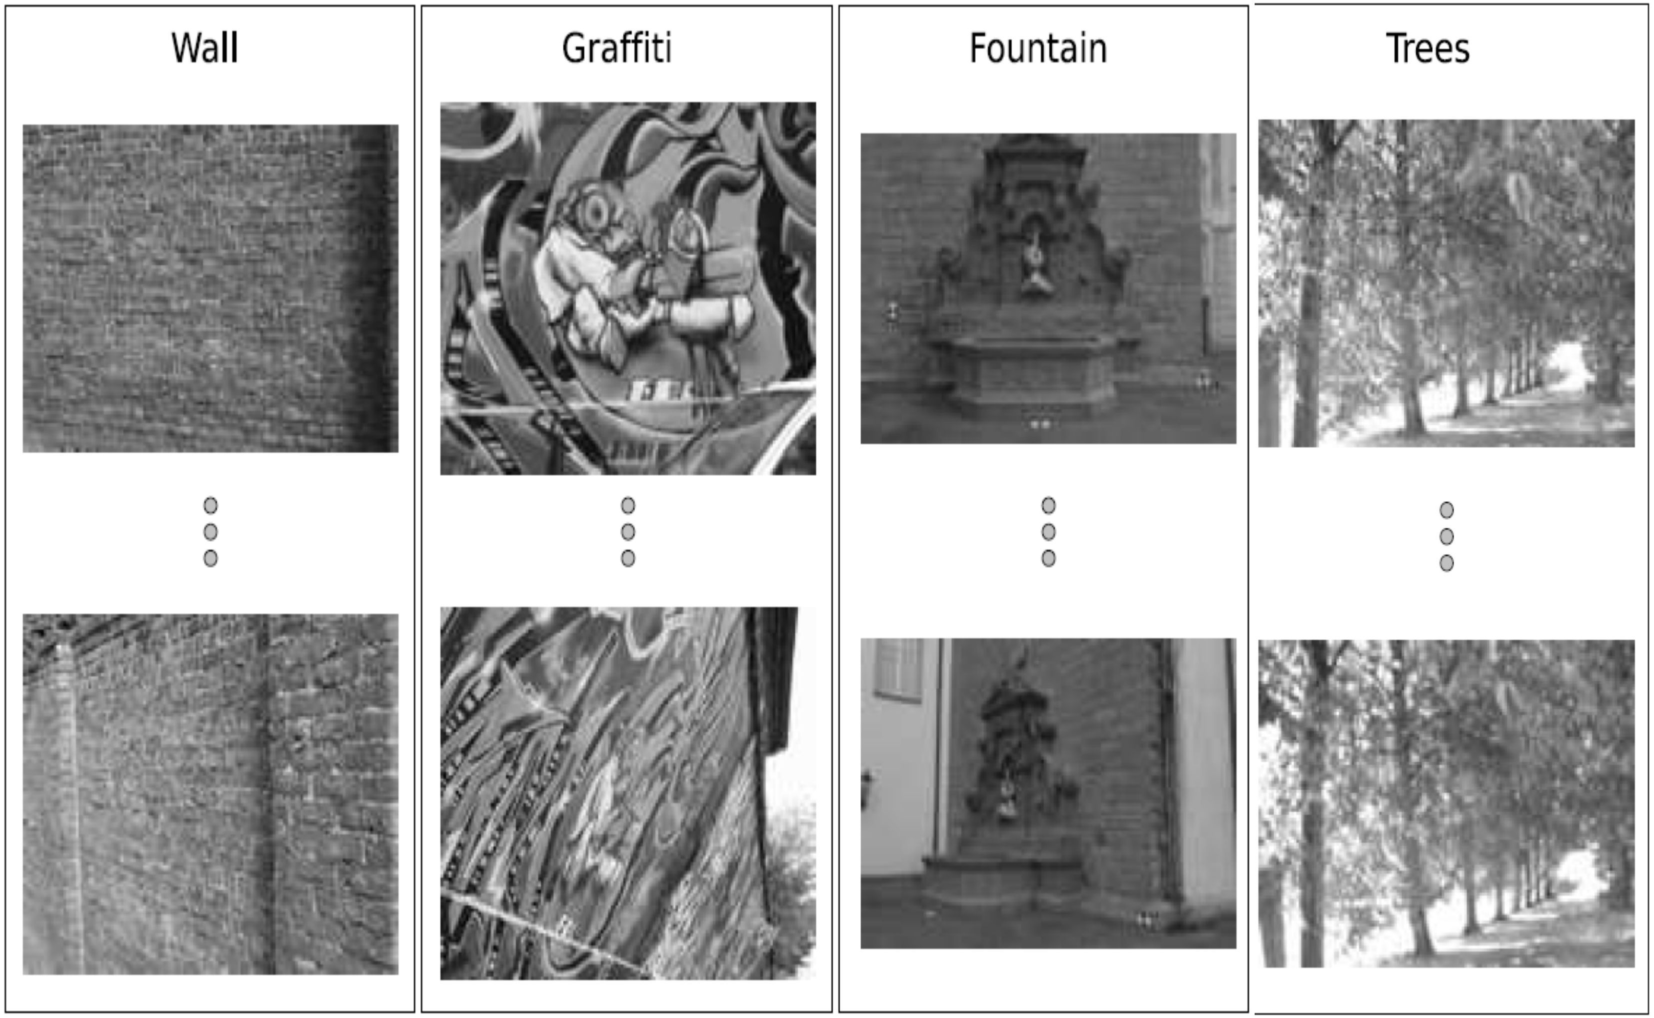

BRIEF vs SURF

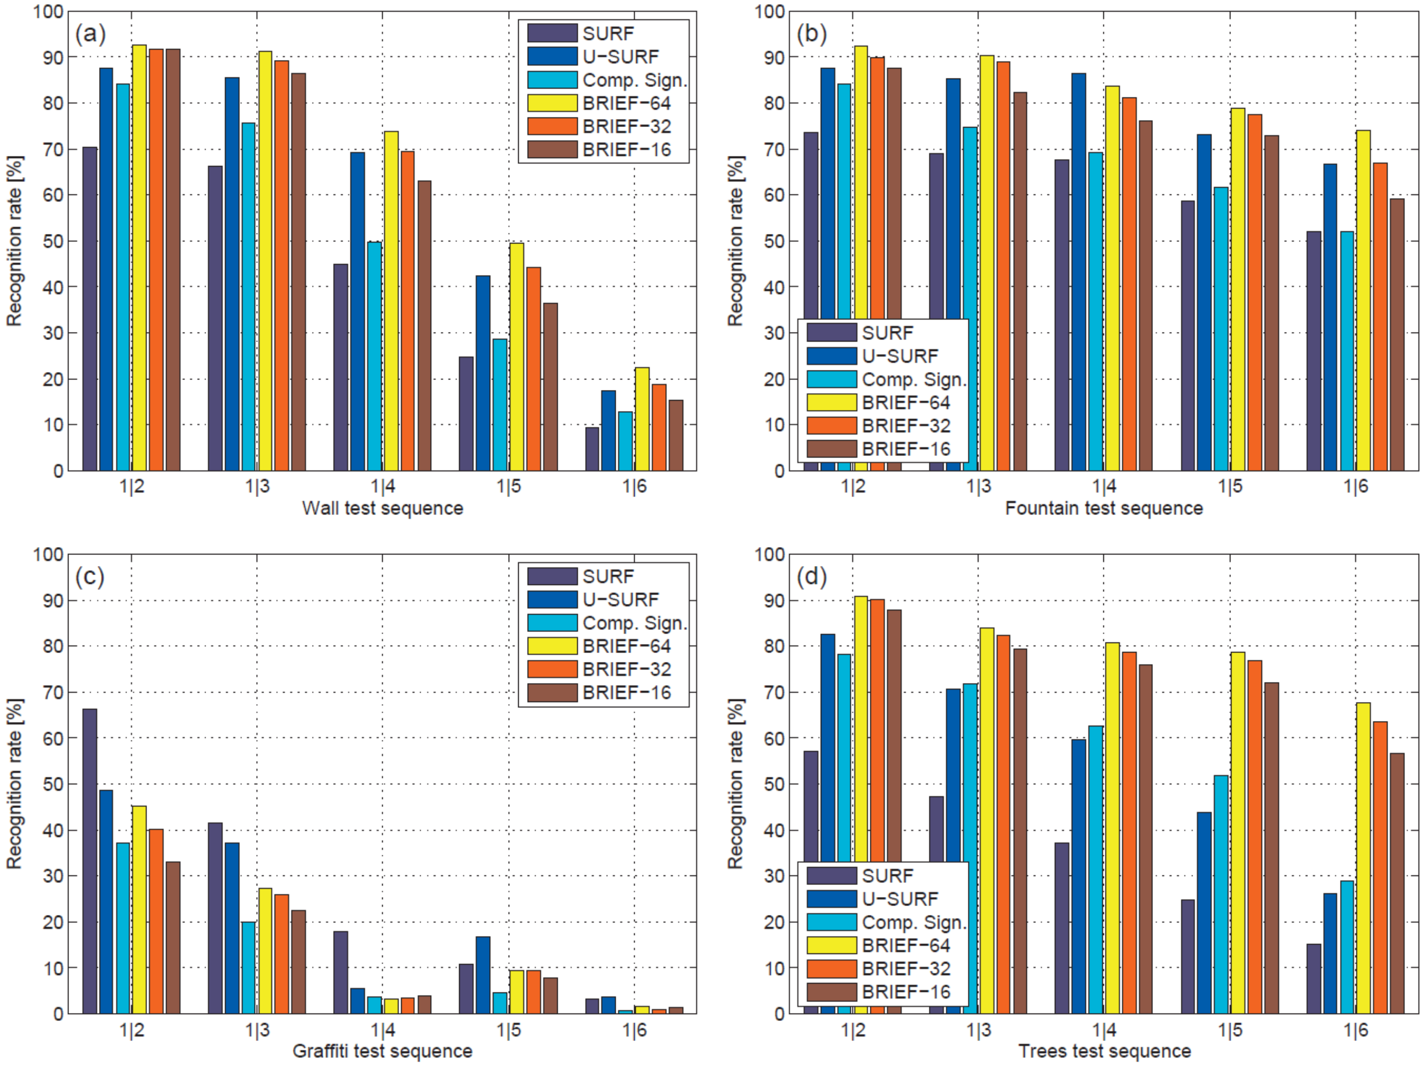

BRIEF vs SIFT

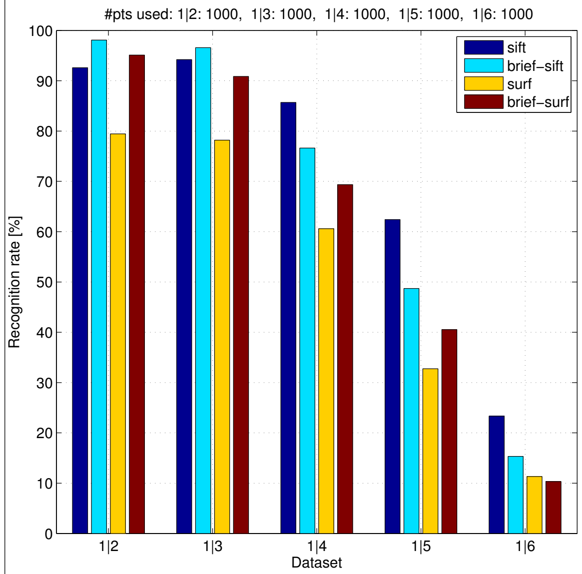

Wall

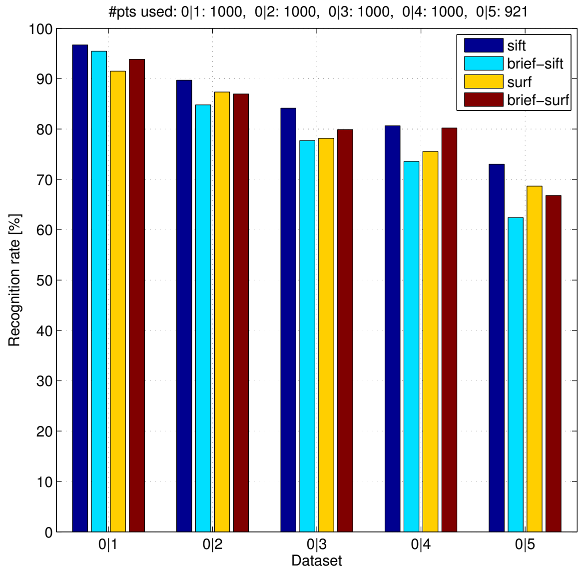

Fountain

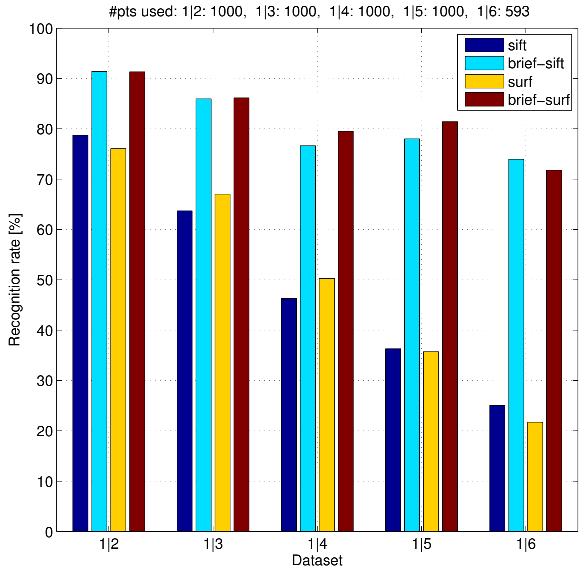

Trees

SIFT > BRIEF > SURF.

Be careful about interpreting benchmarks!

Computational Issues

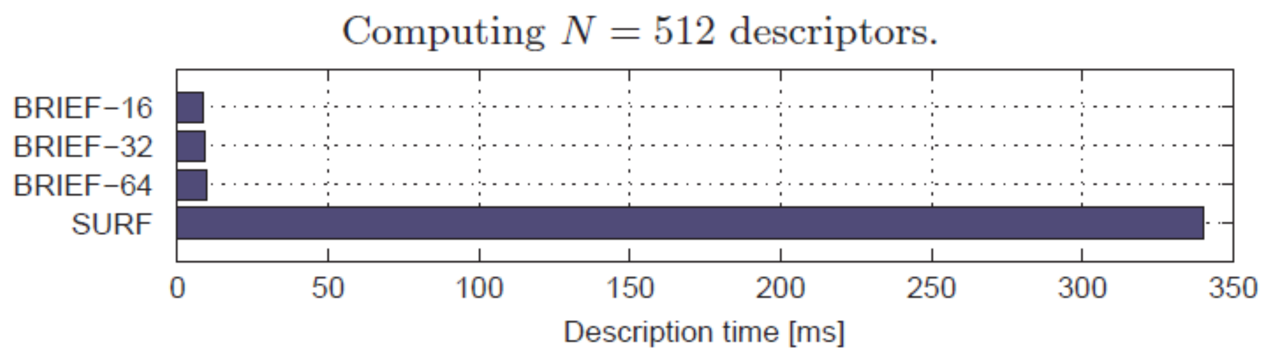

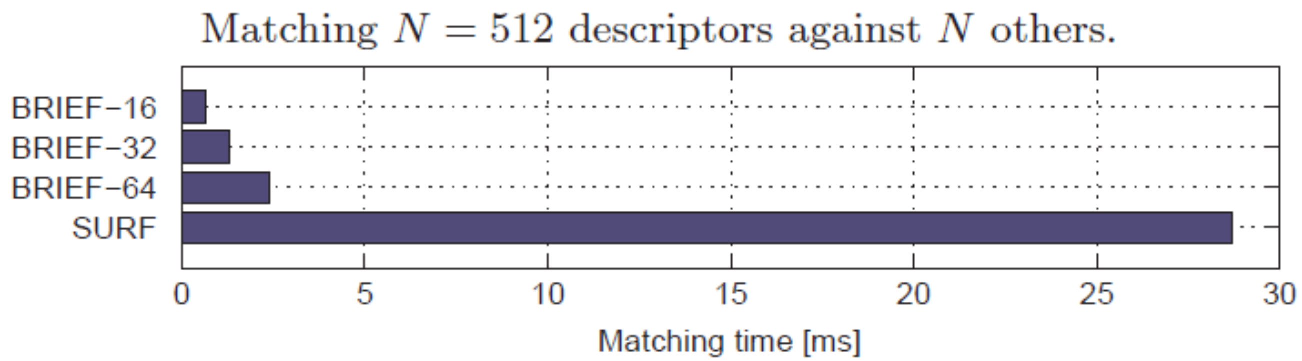

- Integral images can further decrease BRIEF's description time by making smoothing faster.
- Intel Core i7 CPU's POPCOUNT instruction will drastically speed-up the matching of binary vectors.
- Scale and rotational invariance need to be added in some cases.

2D Deformable Surfaces

- Problem Formulation
- Fast Matching
- Robust Optimization
- Lighting Correction

ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

ε_C Correspondence Term

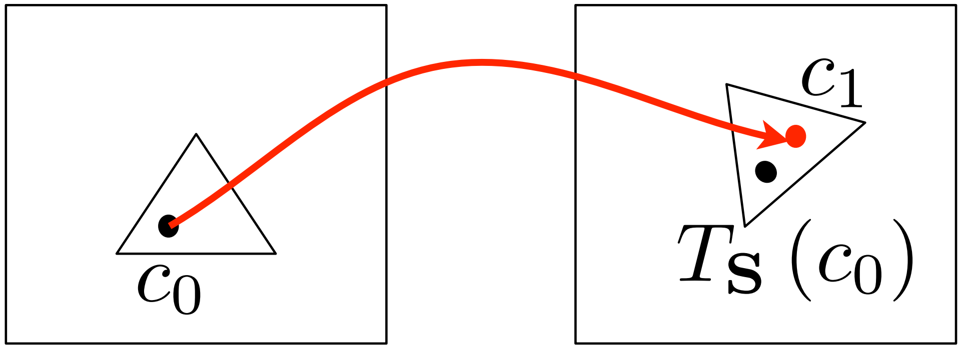

$$\varepsilon_C = -\sum_{c \in C} \left\| c_1 - T_S(c_0) \right\|^2$$

Not robust to outliers!

Robustness to Mismatches

$$\varepsilon_C = -\sum_{c \in C} w_c \rho \left(\left\| c_1 - T_{\mathbf{S}} \left(c_0 \right) \right\|, r \right)$$

where ρ is a robust estimator whose *radius of confidence* is r and $w_c \in [0,1]$ a weight associated to each correspondence.

$$\rho\left(\delta, r\right) = \begin{cases} \frac{3\left(r^2 - \delta^2\right)}{4r^3} & \delta < r \\ 0 & otherwise \end{cases}$$

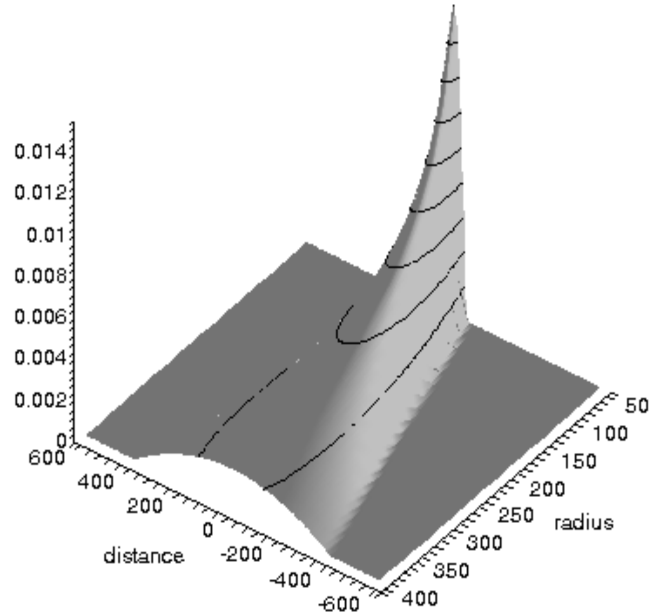

$$\int_{-\infty}^{\infty} \rho(x, r) dx = 1 \;\; \forall r > 0$$

Iterative Reweighting

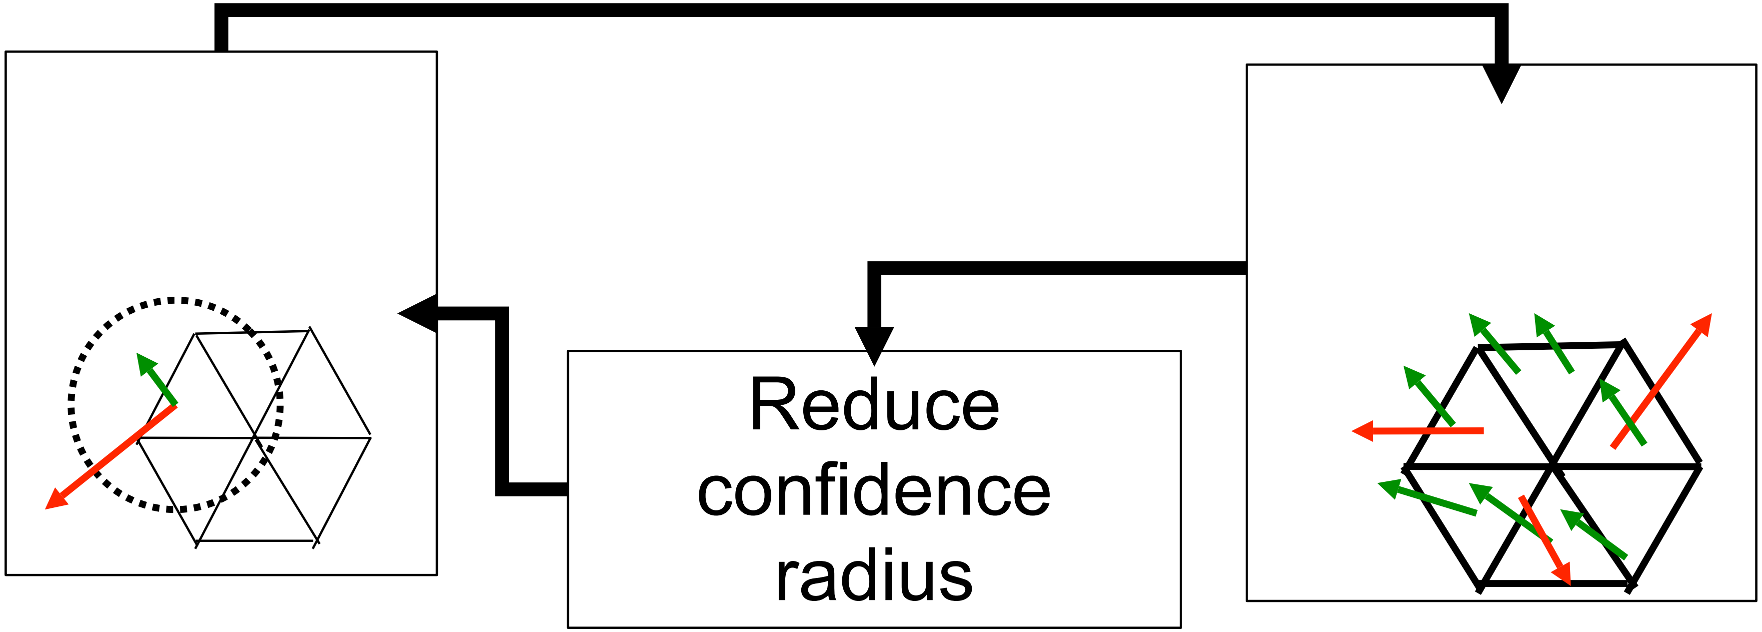

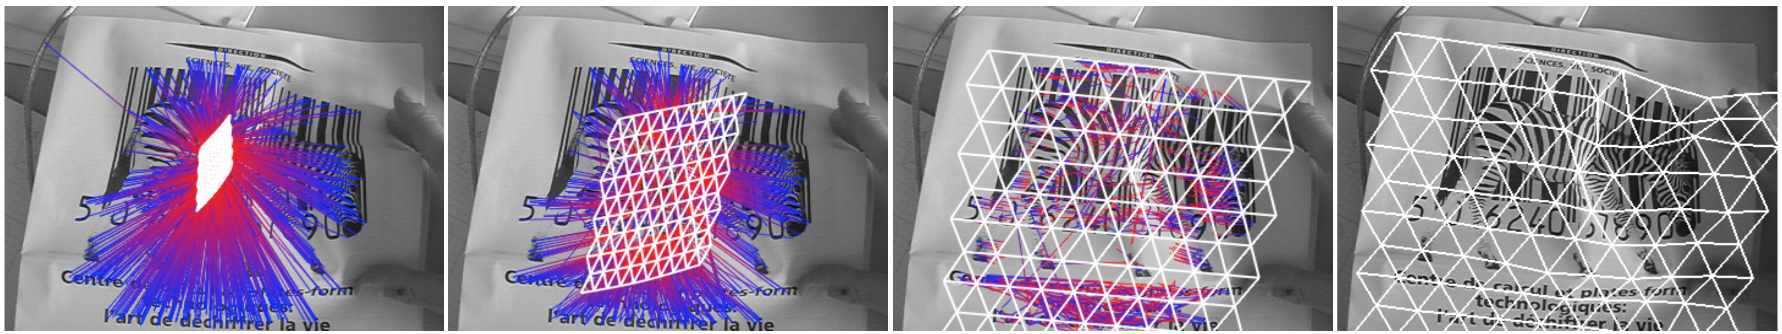

Gauging Robustness

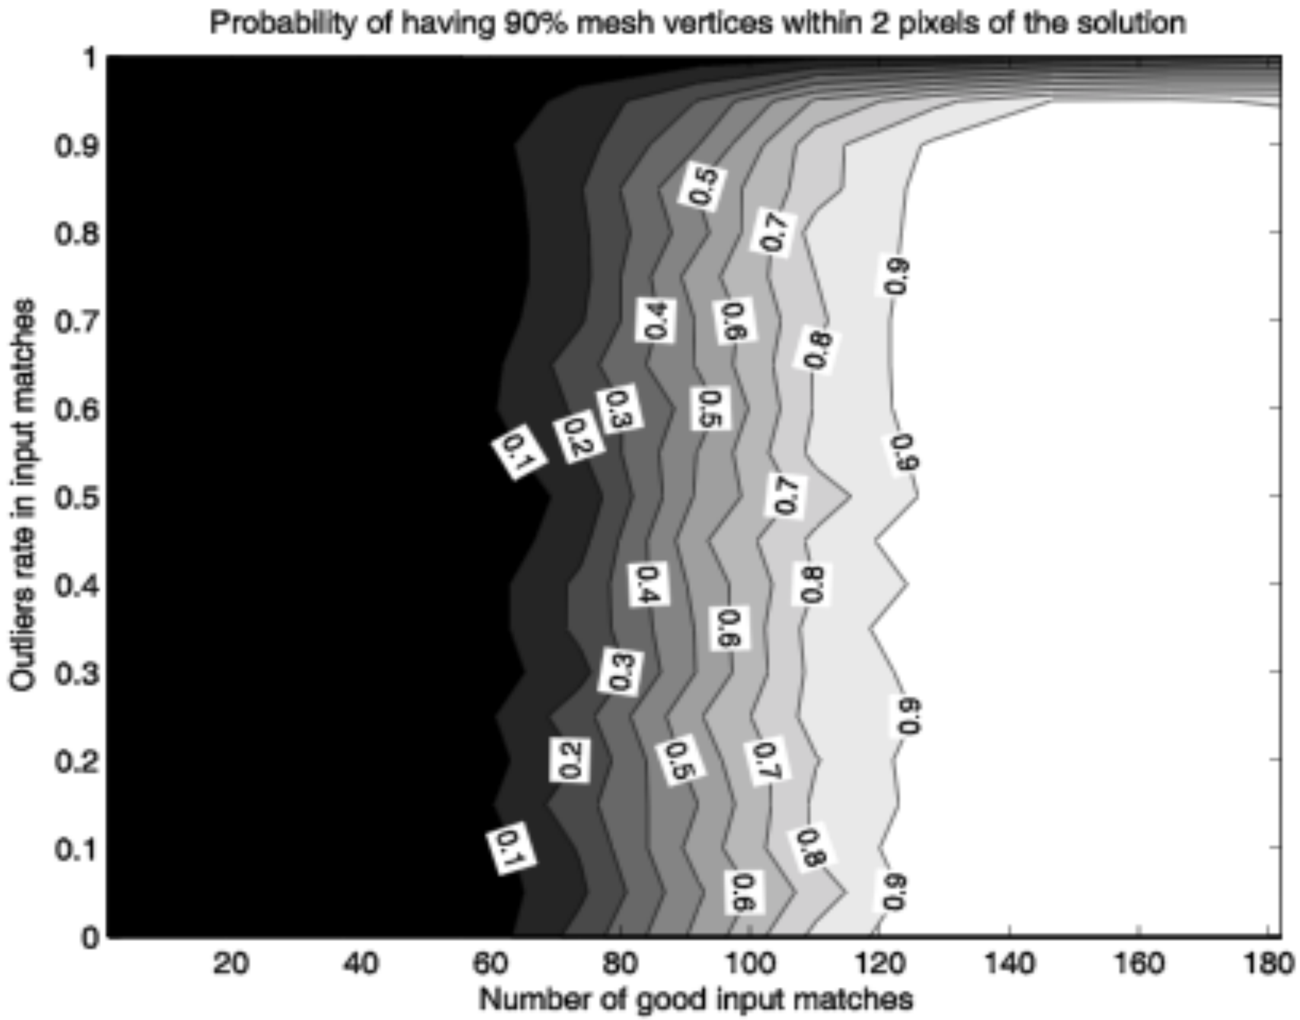

Visualizing the Deformations

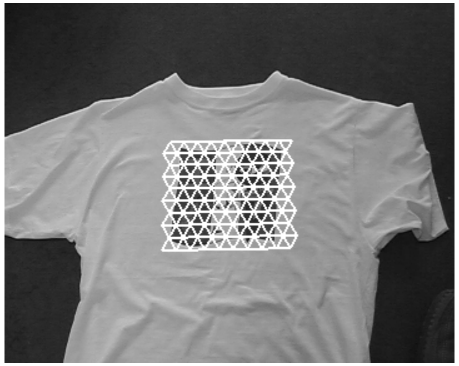

Semi-Implicit Optimization Scheme

Minimize:

$$\varepsilon(S) = \lambda_D \varepsilon_D(S) + \varepsilon_C(S)$$

$$\varepsilon_D(S) = \frac{1}{2}\left(X^T K X + Y^T K Y\right)$$

Satisfied when:

$$0 = \frac{\partial \varepsilon}{\partial X} = \frac{\partial \varepsilon_C}{\partial X} + KX$$

$$0 = \frac{\partial \varepsilon}{\partial Y} = \frac{\partial \varepsilon_C}{\partial Y} + KY$$

M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active Contour Models. IJCV, 1988.

Semi-Implicit Optimization Scheme

Introduce viscosity term:

$$
\begin{aligned}
KX + \alpha \dot{X} &= -\frac{\partial \varepsilon_C}{\partial X} \\
KY + \alpha \dot{Y} &= -\frac{\partial \varepsilon_C}{\partial Y}
\end{aligned}
$$

Time discretization:

$$
\begin{aligned}
0 &= KX_t + \alpha(X_t - X_{t-1}) + \left.\frac{\partial \varepsilon_C}{\partial X}\right|_{X=X_{t-1}Y=Y_{t-1}} \\
0 &= KY_t + \alpha(Y_t - Y_{t-1}) + \left.\frac{\partial \varepsilon_C}{\partial Y}\right|_{X=X_{t-1},Y=Y_{t-1}}
\end{aligned}
$$

Semi-Implicit Optimization Scheme

Solve at each iteration:

$$
\begin{aligned}
(K + \alpha I) X_t &= \alpha X_{t-1} - \left.\frac{\partial \varepsilon_C}{\partial X}\right|_{X = X_{t-1}, Y = Y_{t-1}} \\
(K + \alpha I) Y_t &= \alpha Y_{t-1} - \left.\frac{\partial \varepsilon_C}{\partial Y}\right|_{X = X_{t-1}, Y = Y_{t-1}}
\end{aligned}
$$

--> Fast because K has only a few non zero diagonals.

ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

Newton Optimization Scheme

Taylor expansion of data term:

$$
\begin{aligned}
\varepsilon(X,Y) &= \lambda_D \varepsilon_D(X,Y) + \varepsilon_C(X,Y) \\
\varepsilon_D(X,Y) &= \frac{1}{2}\left(X^T K X + Y^T K Y\right) \\
\varepsilon_C(X+dX, Y+dY) &= A + BdX + CdY + \frac{1}{2} dX^t D dX + \frac{1}{2} dY^t E dY
\end{aligned}
$$

Zhu and Lyu, ECCV'07

Newton Optimization Scheme

At the minimum:

$$
\begin{aligned}
0 &= \frac{\partial \varepsilon}{\partial X} = B + DdX + K(X + dX) \\
0 &= \frac{\partial \varepsilon}{\partial Y} = C + EdY + K(Y + dY)
\end{aligned}
$$

Solve at each iteration:

$$(K + D)dX = -B - KX$$

$$(K + E)dY = -C - KY$$

Semi-Implicit vs Newton

Residuals as a function of the number of iterations: Semi-Implicit in green and Newton in blue.

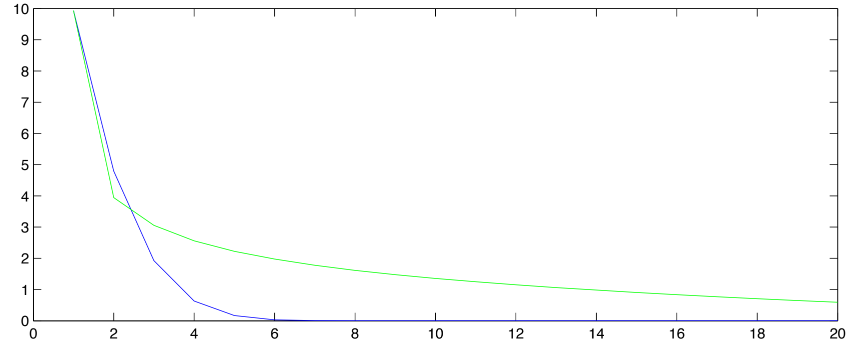

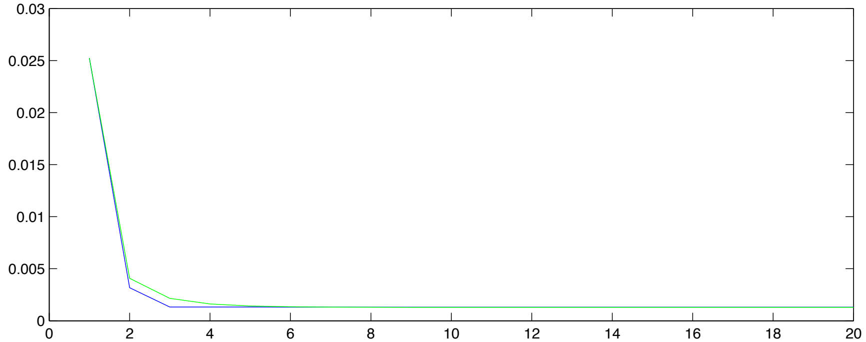

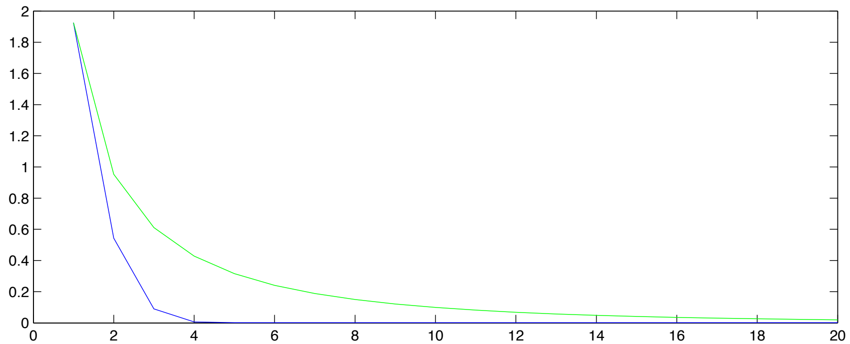

$$\sum_i (log(x_i) - y_i)^2 + \frac{1}{2} X^t K X$$

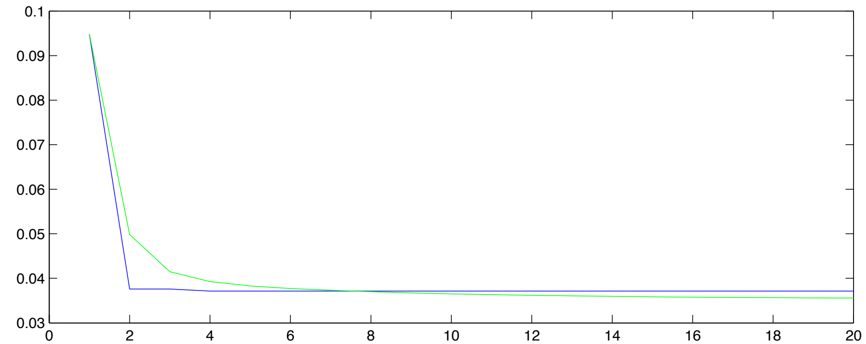

$$\sum_i (sin(x_i) - y_i)^2 + \frac{1}{2} X^t K X$$

2D Deformable Surfaces

- Problem Formulation
- Fast Matching
- Robust Optimization
- Lighting Correction

Lighting

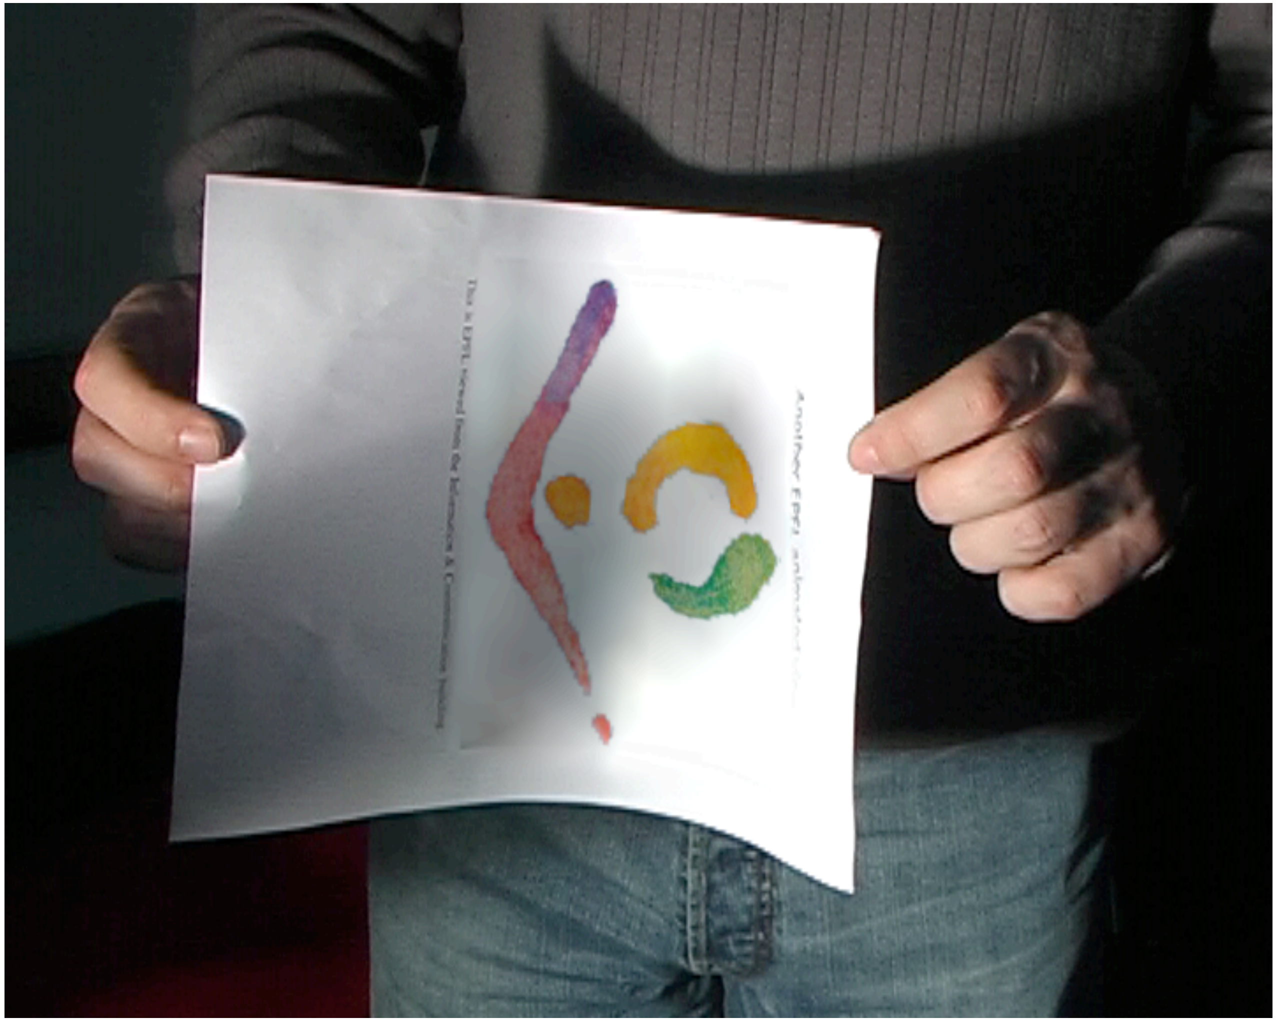

Intensity Ratios

$$
\begin{aligned}
\text{Reference image: } I_{r,p} &= L_r A_p \\
\text{Input image: } I_{i,p} &= L_{i,p} A_p \\
\\
\text{White image: } I_{r,w} &= L_r A_w \\
\text{Synthetic image: } I_{x,p} &= L_{i,p} A_w \\
&= A_w L_r \frac{I_{i,p}}{I_{r,p}} \\
&= I_{r,w} \frac{I_{i,p}}{I_{r,p}}
\end{aligned}
$$

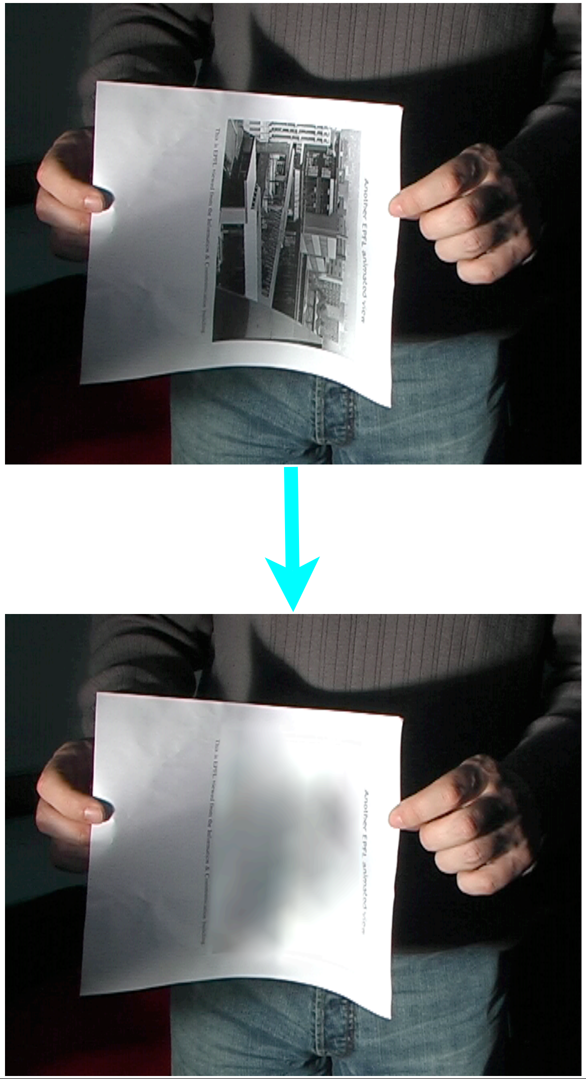

Background Subtraction

Standard approach:

- Pixel-wise statistical background model.

Modified approach:

- Account for the fact that illuminations changes tend to be correlated.
- Model variations of intensity ratios as GMMs.

--> Effective for occlustion detection.

ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

CVLab

Realistic Augmentation

3D Deformable Surfaces

Reference

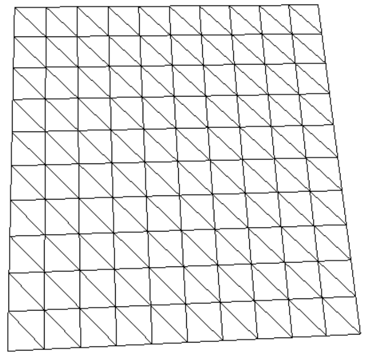

Input

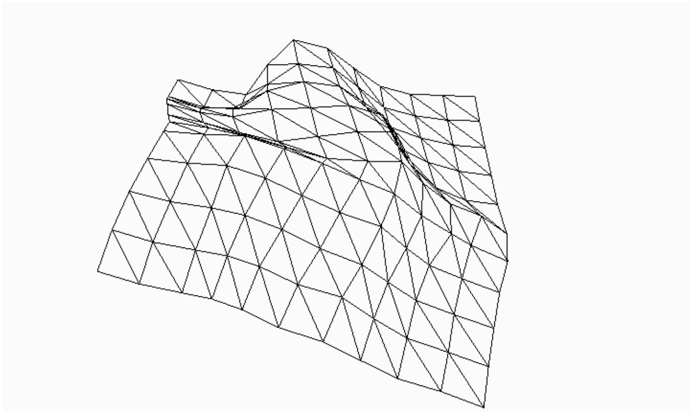

Problem Formulation

Input:

- Reference image.
- Corresponding 3D surface.
- Projection matrix P.
- 3D-to-2D correspondences between reference configuration and input image.

Unknowns:

- Mesh vertex coordinates corresponding to input image

$$\mathbf{X} = [x_1, y_1, z_1, \cdots, x_{n_v}, y_{n_v}, z_{n_v}]^T$$

Ambiguity

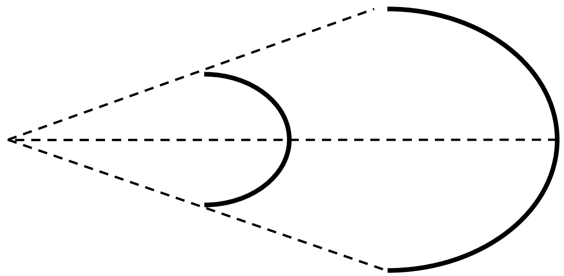

Scale ambiguity

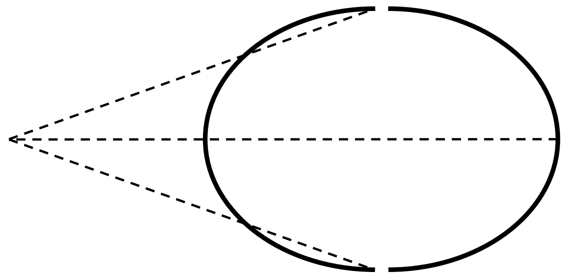

Bas-Relief Ambiguity

- 3D Shape or deformation models are needed.

➢ How can we design models that do not make unwarranted assumptions?

ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

3D Deformable Surfaces

- Linear Formulation.
- Inextensible surfaces.
- Sharply folding surfaces.
- Eliminating the reference image.
- Applications.

Linear Formulation

- Calibrated camera, **A** intrinsic parameters matrix.
- Coordinates expressed in the camera referential.
- Unknown mesh vertex coordinates: $\mathbf{X} = (\mathbf{v}_1^{\mathrm{T}}, \ldots, \mathbf{v}_{n_V}^{\mathrm{T}}), \ \mathbf{v}_i = (x_i, y_i, z_i)^{\mathrm{T}}$
- Correspondences
  - Barycentric coordinates from reference configuration: (a_i, b_i, c_i)
  - Current image location: $(u_i, v_i)^T$

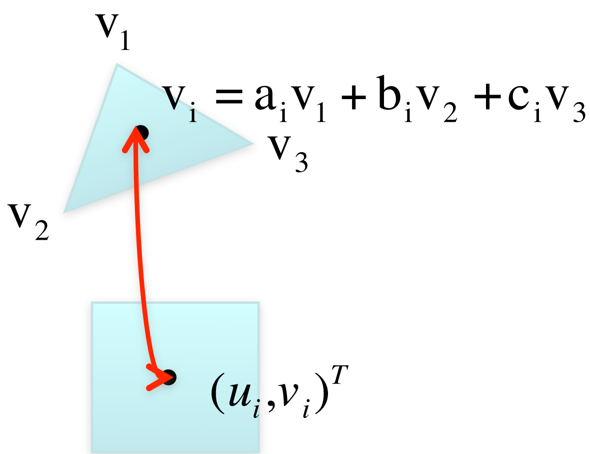

$$\begin{pmatrix} u_i \\ v_i \\ 1 \end{pmatrix} = \frac{1}{k_i} \mathbf{A} (a_i \mathrm{v}_1 + b_i \mathrm{v}_2 + c_i \mathrm{v}_3)$$

Salzmann et al., CVPR'07

ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

Linear Formulation

$$\mathbf{A}\left(b_1\mathbf{v}_1 + b_2\mathbf{v}_2 + b_3\mathbf{v}_3\right) = k\begin{bmatrix} u \\ v \\ 1 \end{bmatrix}$$

$\rightarrow$ k can be expressed in terms of the vertex coordinates using the last row.

$$\begin{bmatrix} b_1\mathbf{H} & b_2\mathbf{H} & b_3\mathbf{H} \end{bmatrix}\begin{bmatrix} \mathbf{v}_1 \\ \mathbf{v}_2 \\ \mathbf{v}_3 \end{bmatrix} = \mathbf{0}\ ,$$

with

$$\mathbf{H} = \mathbf{A_{2\times 3}} - \begin{bmatrix} u \\ v \end{bmatrix}\mathbf{A_3}\ ,$$

where $\mathbf{A_{2\times 3}}$ contains the first two rows of $\mathbf{A}$, and $\mathbf{A_3}$ is the third one.

--> Each correspondence gives rise to two linear equations.

Linear System and Singular Values

X must be solution of **MX** =0

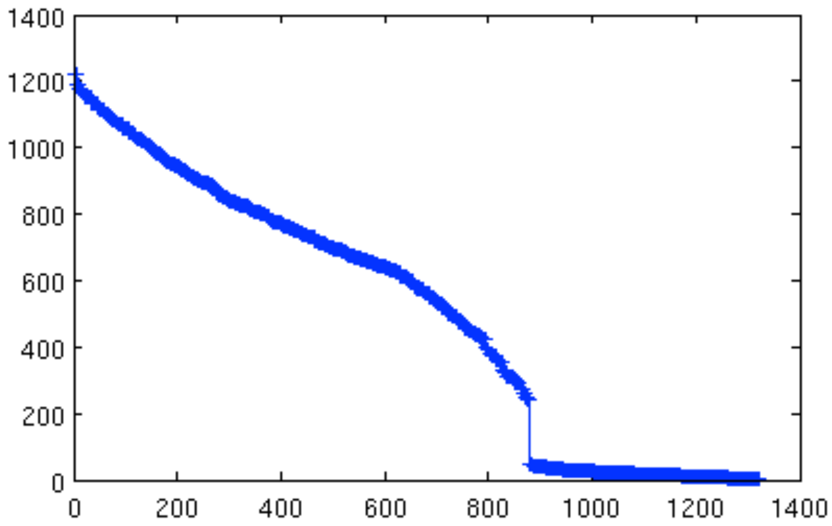

Singular values of **M**

ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

CVLab

3D Deformable Surfaces

- Linear Formulation.
- Inextensible surfaces.
- Sharply folding surfaces.
- Eliminating the reference image.
- Applications.

Inextensible Meshes

A solution of the linear system belongs to the kernel of $\mathbf{M}$:

$$\mathbf{MX} = 0 \Rightarrow \mathbf{X} = \sum_i \beta_i \mathbf{p}_i ,$$

where the $\mathbf{p}_i$ are the eigenvectors corresponding to small eigenvalues.

Inextensible mesh:

$$\left\| \sum_i \beta_i \mathbf{p}_i^{\mathrm{j}} - \sum_i \beta_i \mathbf{p}_i^{\mathrm{k}} \right\|^2 = \mathrm{cte}$$

for all neighboring vertices j and k.

--> A system a quadratic equations that could be solved in closed form using extended linearization, but with too many variables for existing solvers.

Salzmann et al., ECCV'08

Dimensionality Reduction

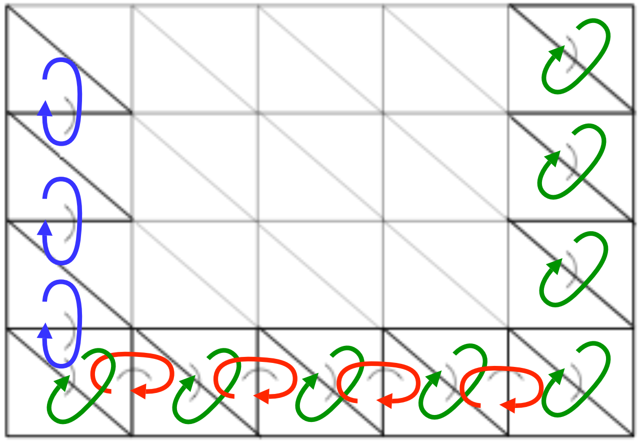

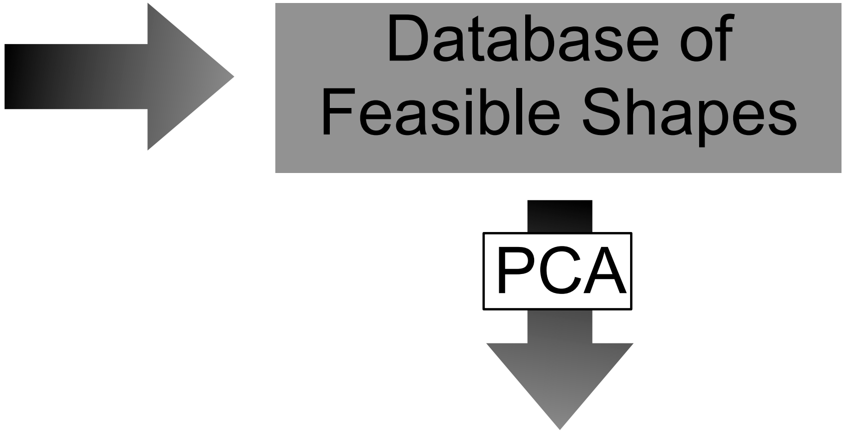

$$\mathbf{X} = \mathbf{X}_0 + \sum_i \alpha_i \mathbf{S}_i$$

$$= \mathbf{X}_0 + \mathbf{S}\mathbf{A}$$

$$\text{with } \mathbf{A} = \begin{bmatrix} \alpha_1 & \dots & \alpha_N \end{bmatrix}^T$$

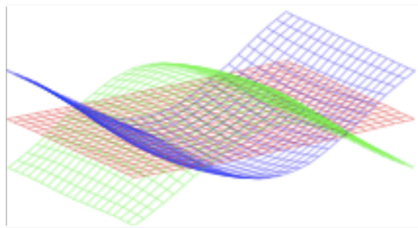

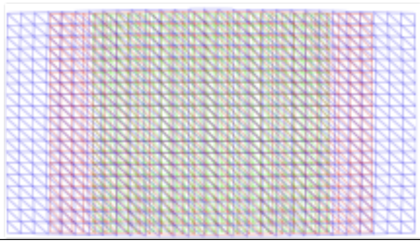

Salzmann et al., PAMI’07

Degrees of Freedom

For an inextensible triangulation with V vertices, E=Ei+Eb edges, and F facets with no holes:

- Euler formula

  V+F-E=1 .

- Interior edges shared by two facets

  3F=2Ei+Eb.

- Degrees of freedom

  3V-E=6+Eb

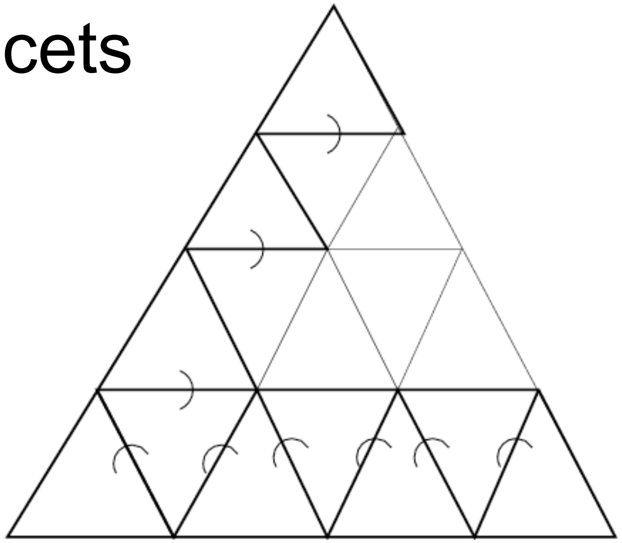

Spinnaker Modes

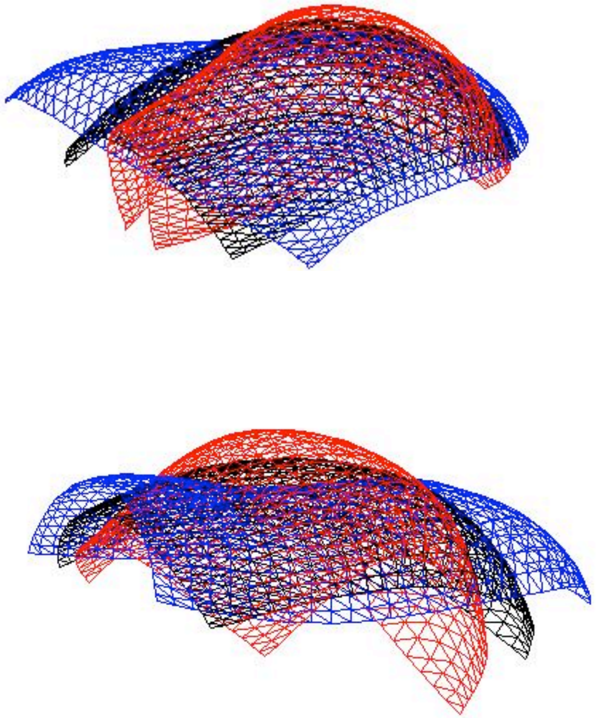

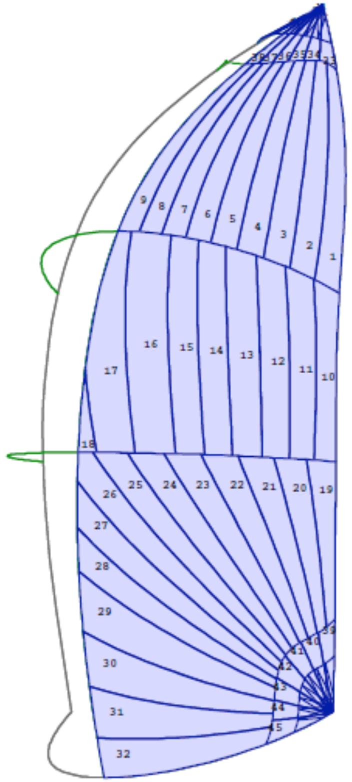

Reduced System

$$\mathbf{MX} = 0$$

$$\mathbf{X} = \mathbf{X}_0 + \mathbf{SA}$$

$$\Rightarrow \begin{bmatrix} \mathbf{MS} & \mathbf{MX}_0 \\ \mathbf{W} & 0 \end{bmatrix} \begin{bmatrix} \mathbf{A} \\ 1 \end{bmatrix} = 0$$

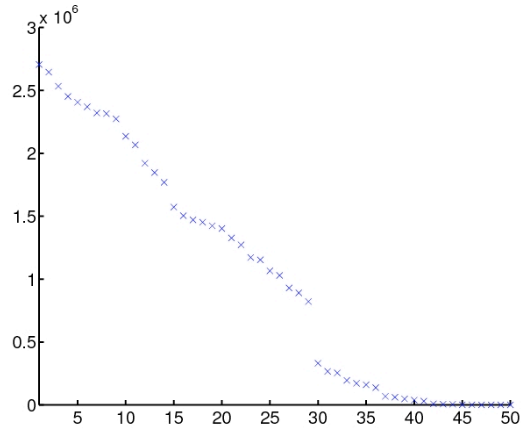

where the $\mathbf{W}$ is a diagonal matrix of modal penalty terms that depends on the eigenvalues of the training data covariance matrix.

- **A** can also be written as a weighted sum of eigenvectors of the extended matrix.
- The inextensibility constraints give rise to a smaller set of quadratic equations than can now be solved.

Independent Detection

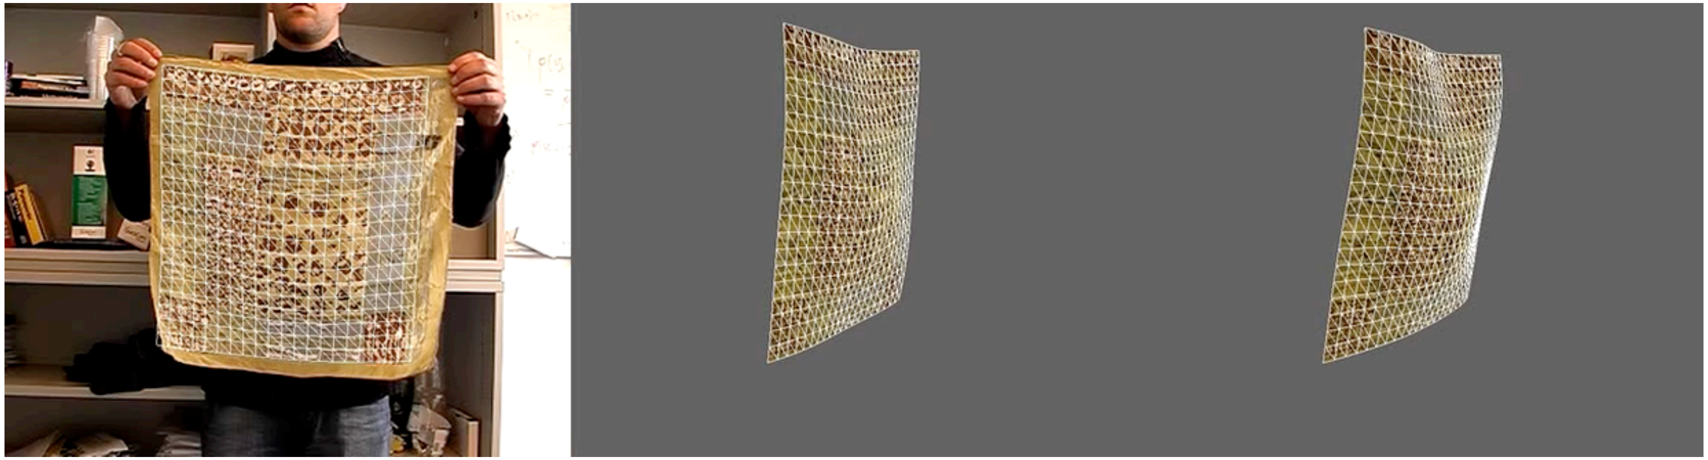

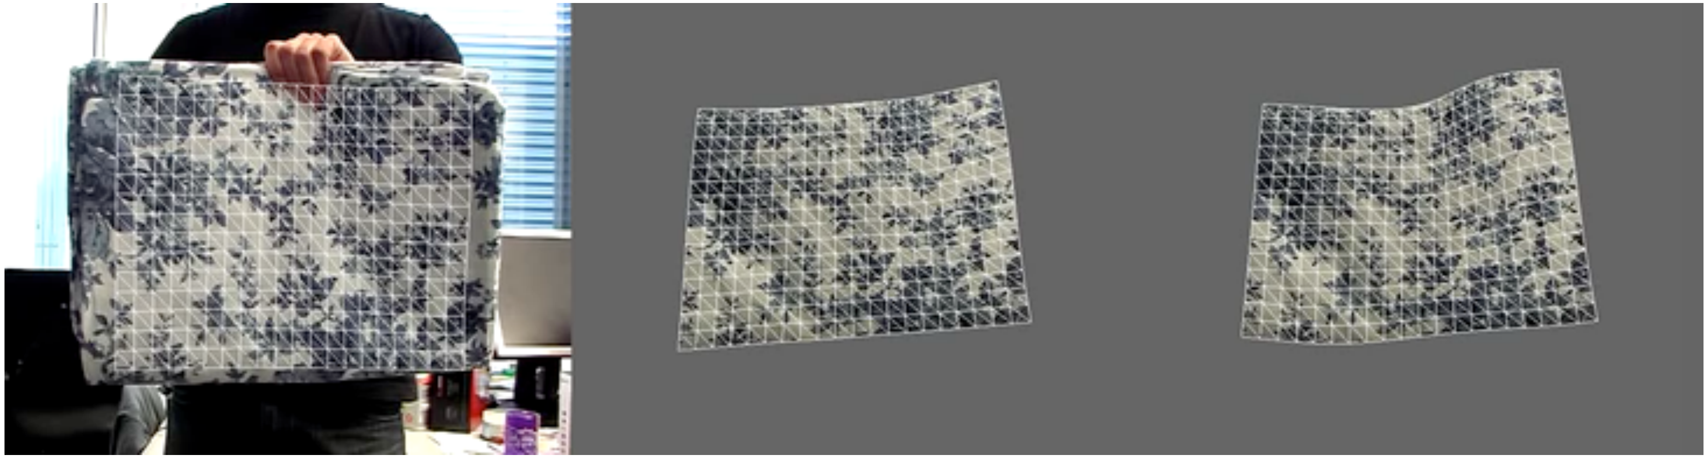

Limitation

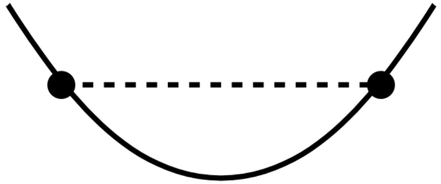

In the presence of sharp folds:

- The Euclidean distance between discrete points decreases.
- Inextensibility constraints are not appropriate anymore.

3D Deformable Surfaces

- Linear Formulation.
- Inextensible surfaces.
- Sharply folding surfaces.
- Eliminating the reference image.
- Applications.

Handling Creases

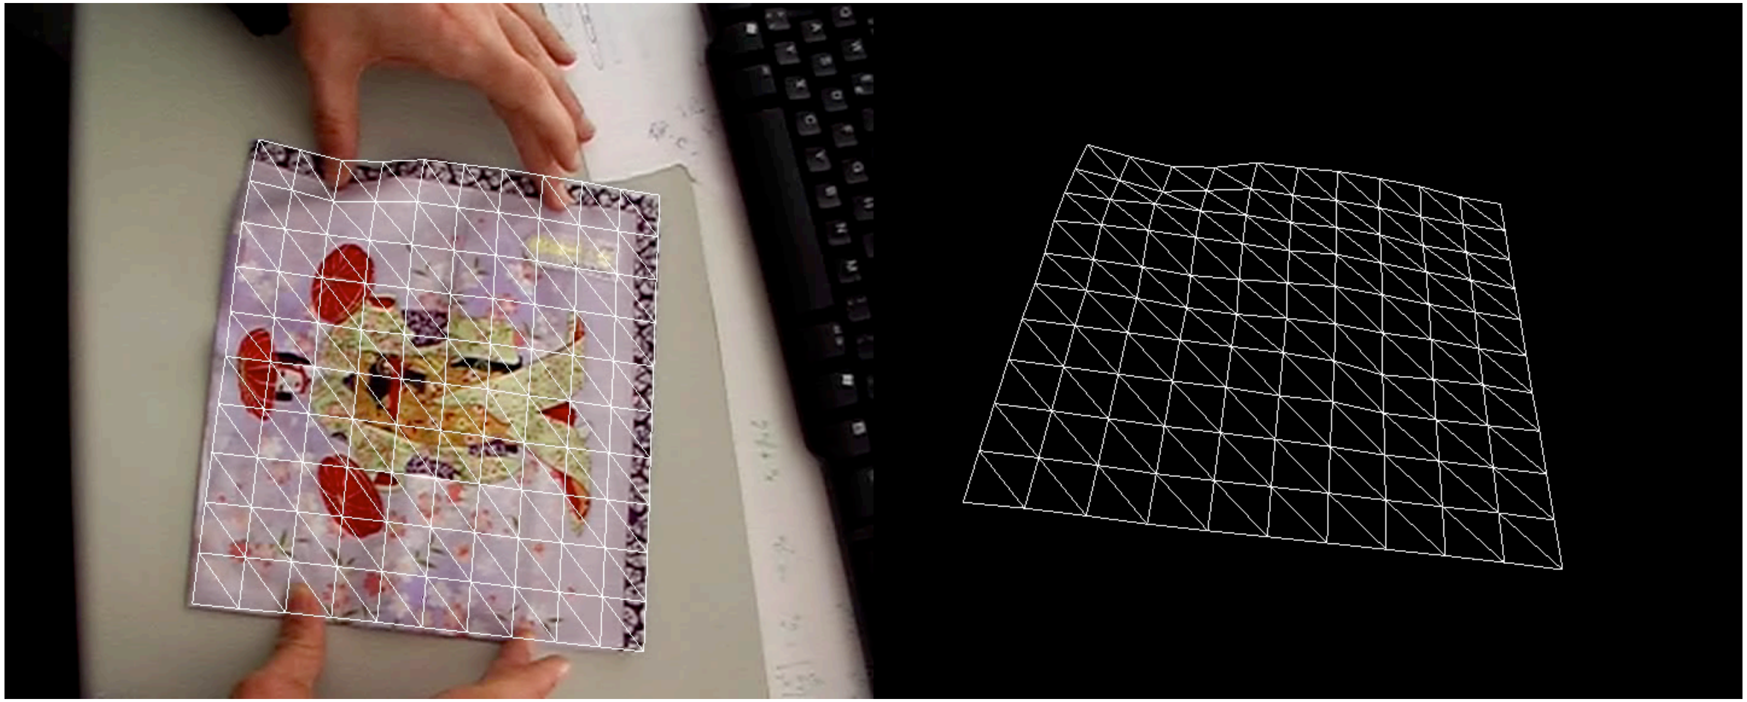

Replace inextensibility constraints by distance inequalities that:

- Let us reconstruct surfaces with sharp folds.
- Yield a convex formulation of the reconstruction problem.

Salzmann et al., CVPR'09

Inequality Constraints

In the presence of sharp folds, geodesic distances remain constant, but Euclidean ones may decrease.

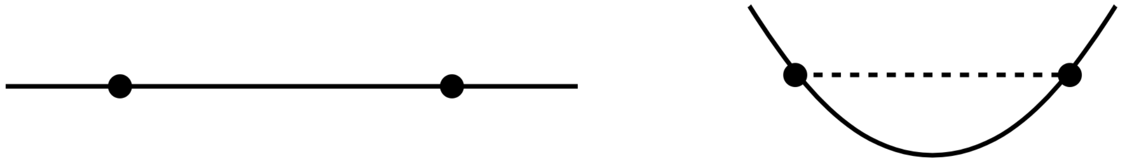

~~$\left\| \mathrm{v}_j - \mathrm{v}_k \right\| = l_{j,k}$~~ ➔ $\left\| \mathrm{v}_j - \mathrm{v}_k \right\| \leq l_{j,k}$

Naive formulation :

$$\mathbf{X}_{opt} = \arg\min \left\| \mathbf{MX} \right\| \, ,$$

subject to

$$\left\| \mathbf{v}_j - \mathbf{v}_k \right\| \leq d_{jk}$$

for all neighboring vertices j and k.

Pushing the Mesh Away

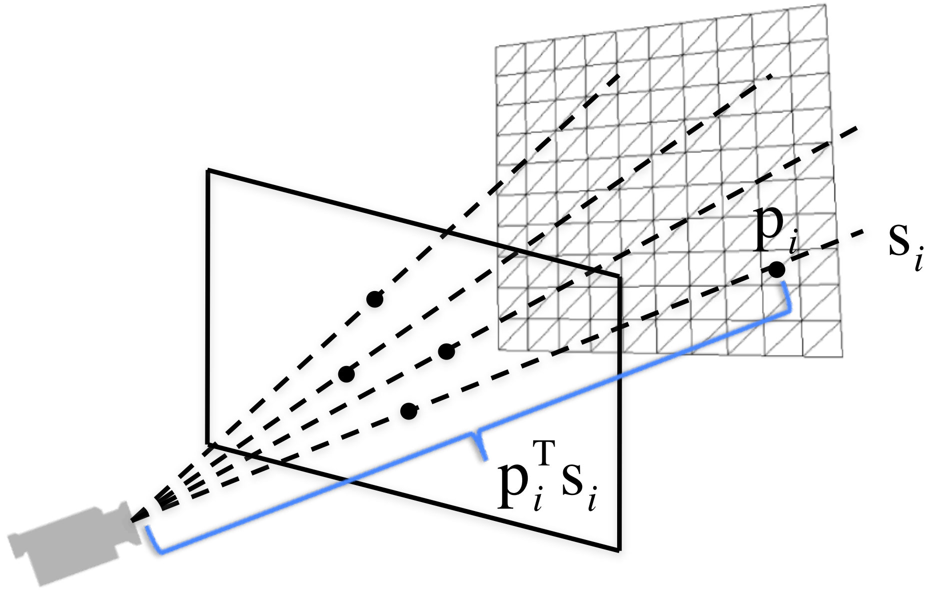

- Inequality constraints do not prevent the mesh from shrinking.

- To this end, we push the points along their lines-of-sight as far as the constraints allow.

Convex Formulation

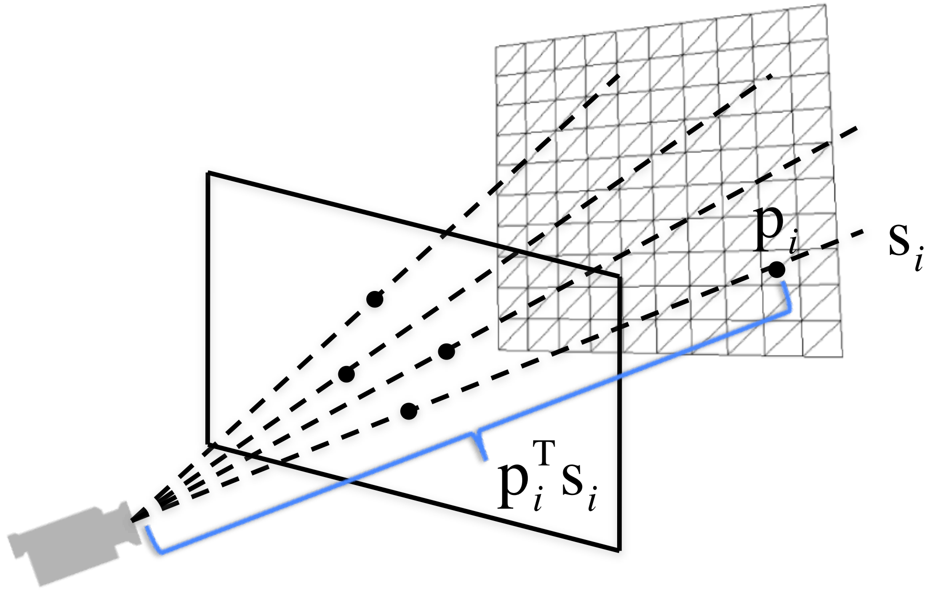

$$\underset{X}{\text{maximize}} \quad w_d \sum_{i \le N_c} p_i^T s_i - \|MX\|$$

$$\text{subject to} \quad \|v_j - v_k\| \le l_{j,k} \quad , \quad \forall (j,k) \in \text{Edges}$$

This is an SOCP problem, which can be solved using standard numerical routines.

Shape Regularization

- Regularization is needed to enforce smoothness on poorly textured parts.
- To handle sharp folds, the global models must be replaced by local ones.

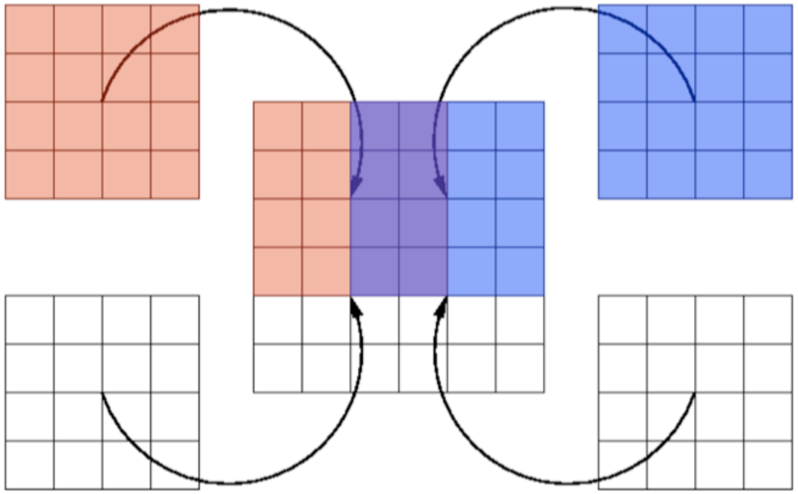

→ Introduce a linear model for individual surfaces patches

$$X^i = X_0^i + \Lambda c^i$$

Local Deformation Model

- To avoid having to explicitly force the coefficients of overlapping patches to be consistent, we express them as

$$c^i = \Lambda^T (X^i - X_0^i) \ ,$$

which arises from the orthonormality of the modes.

- Regularization is achieved by penalizing the coefficients associated to high energy modes, which is done by minimizing

$$\sum_{i \leq N_p} w_i \left\| \Sigma^{-1/2} \Lambda^T (X^i - X_0^i) \right\| \ ,$$

where Σ contains the eigenvalues of the training data covariance matrix.

Temporal Regularization

For short video-sequences, we can enforce temporal consistency by introducing a second order---constant speed---motion model:

$$\Delta_{t-1,t} = \Delta_{t,t+1}$$

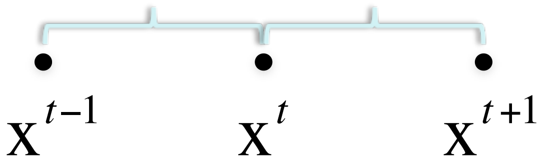

→ We solve our optimization problem for 3 frames simultaneously, and regularize the motion between frames by minimizing

$$\left\| X^{t-1} - 2X^{t} + X^{t+1} \right\| .$$

Synthetic Data

- Optical motion capture:

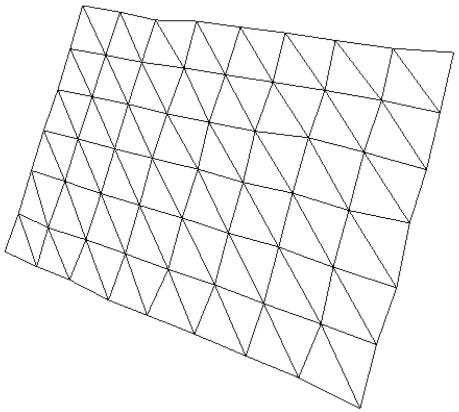
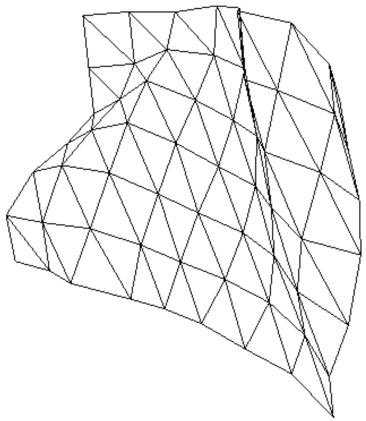
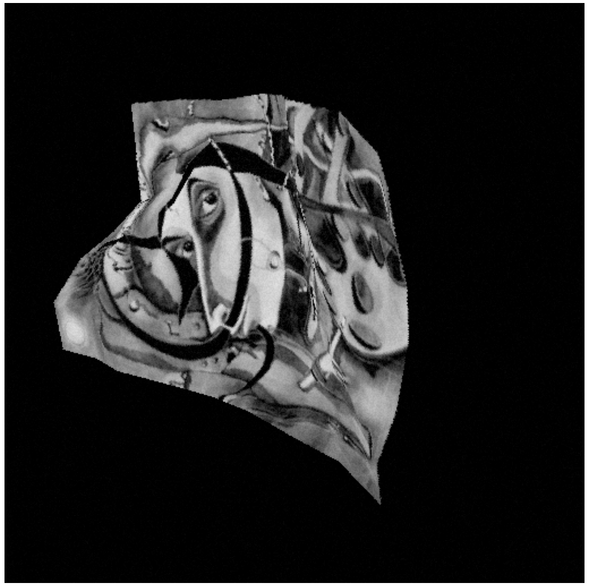

- Correspondences:
  - Sample the barycentric coordinates, project the 3D points, add Gaussian noise with variance 5 to the image locations.
  - Compute SIFT matches between the input images and the reference.

Synthetic Correspondences

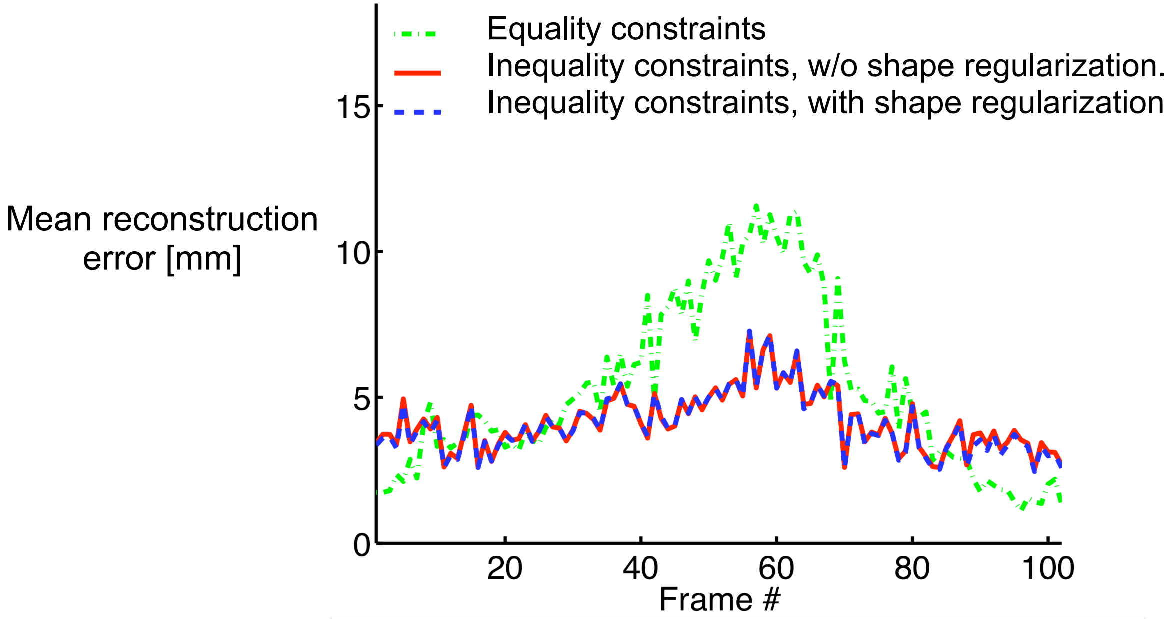

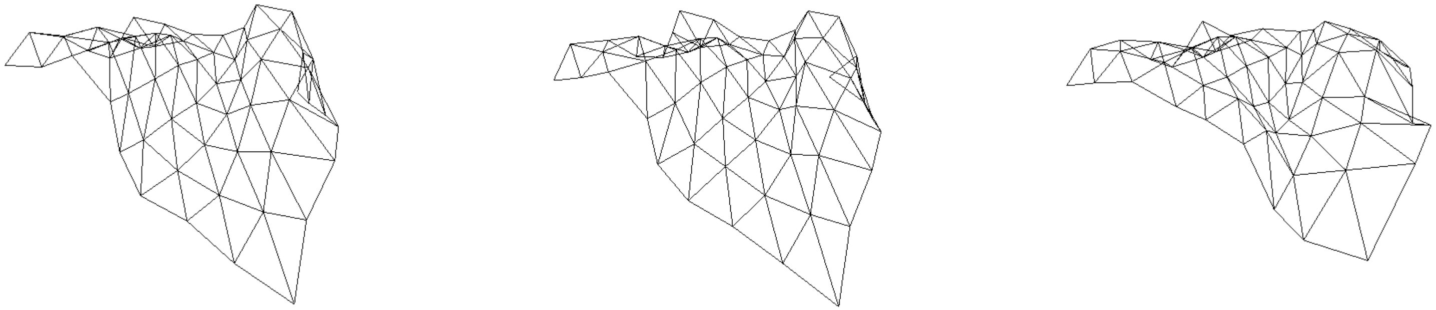

SIFT Correspondences

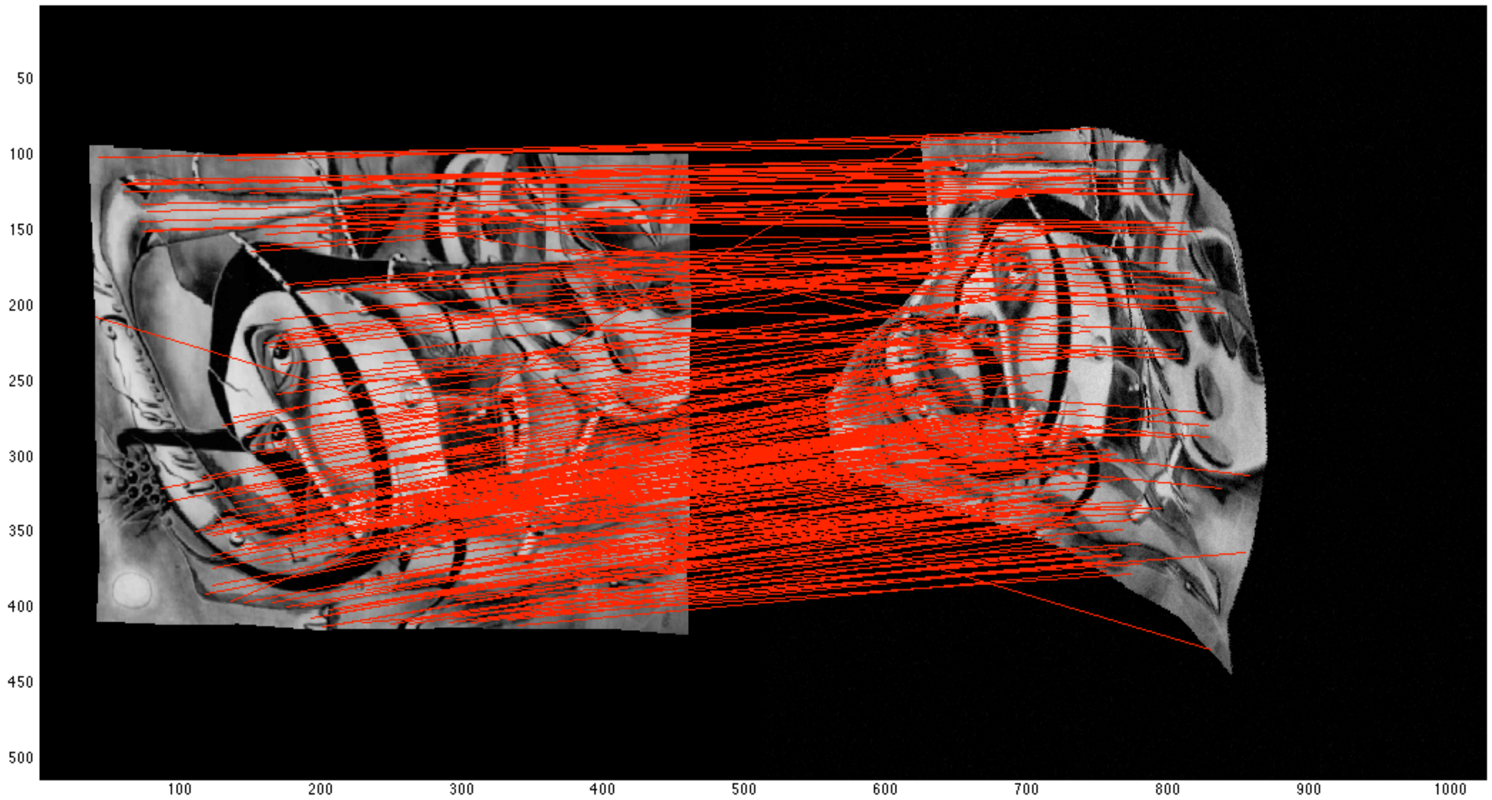

SIFT Correspondences

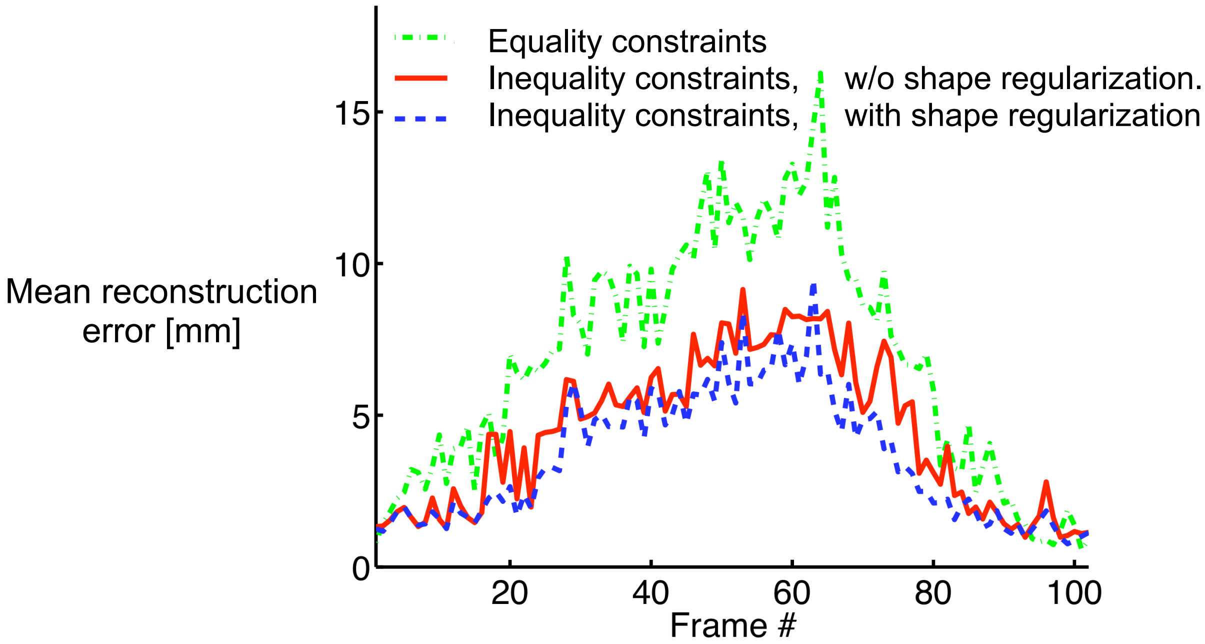

Introducing Outliers

- Synthetic correspondences.
- Varying outlier rate.

Single frames

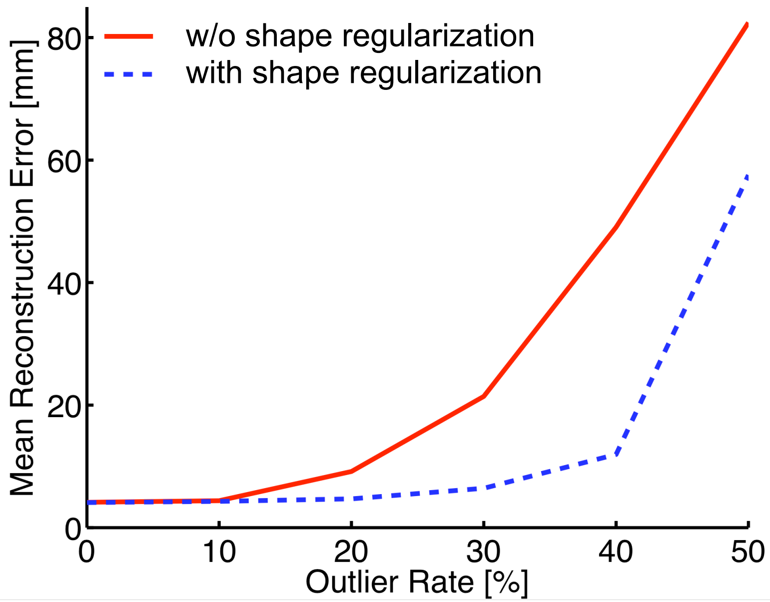

Multiple frames

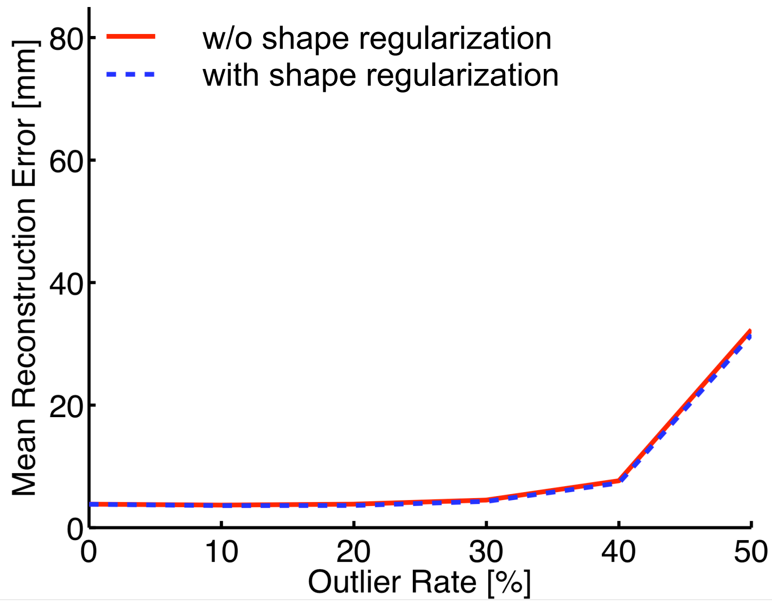

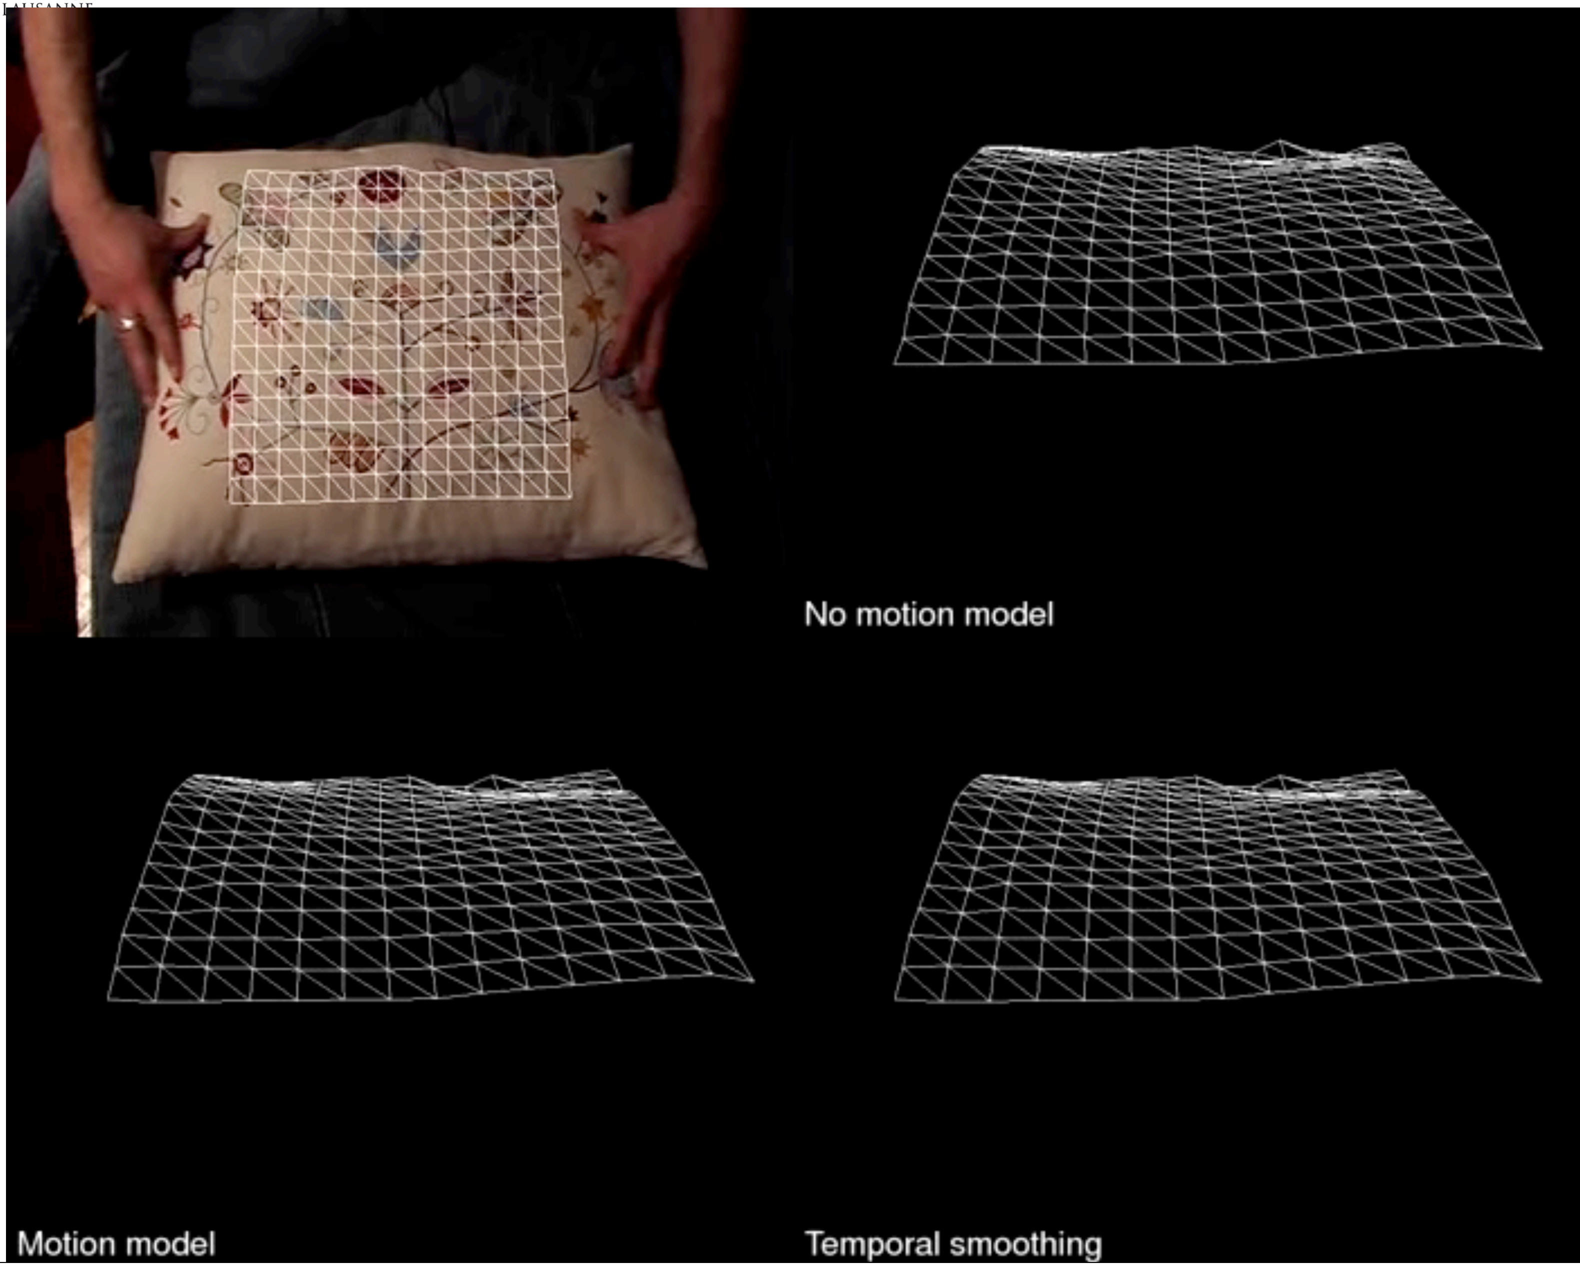
No motion model
Motion model
Temporal smoothing

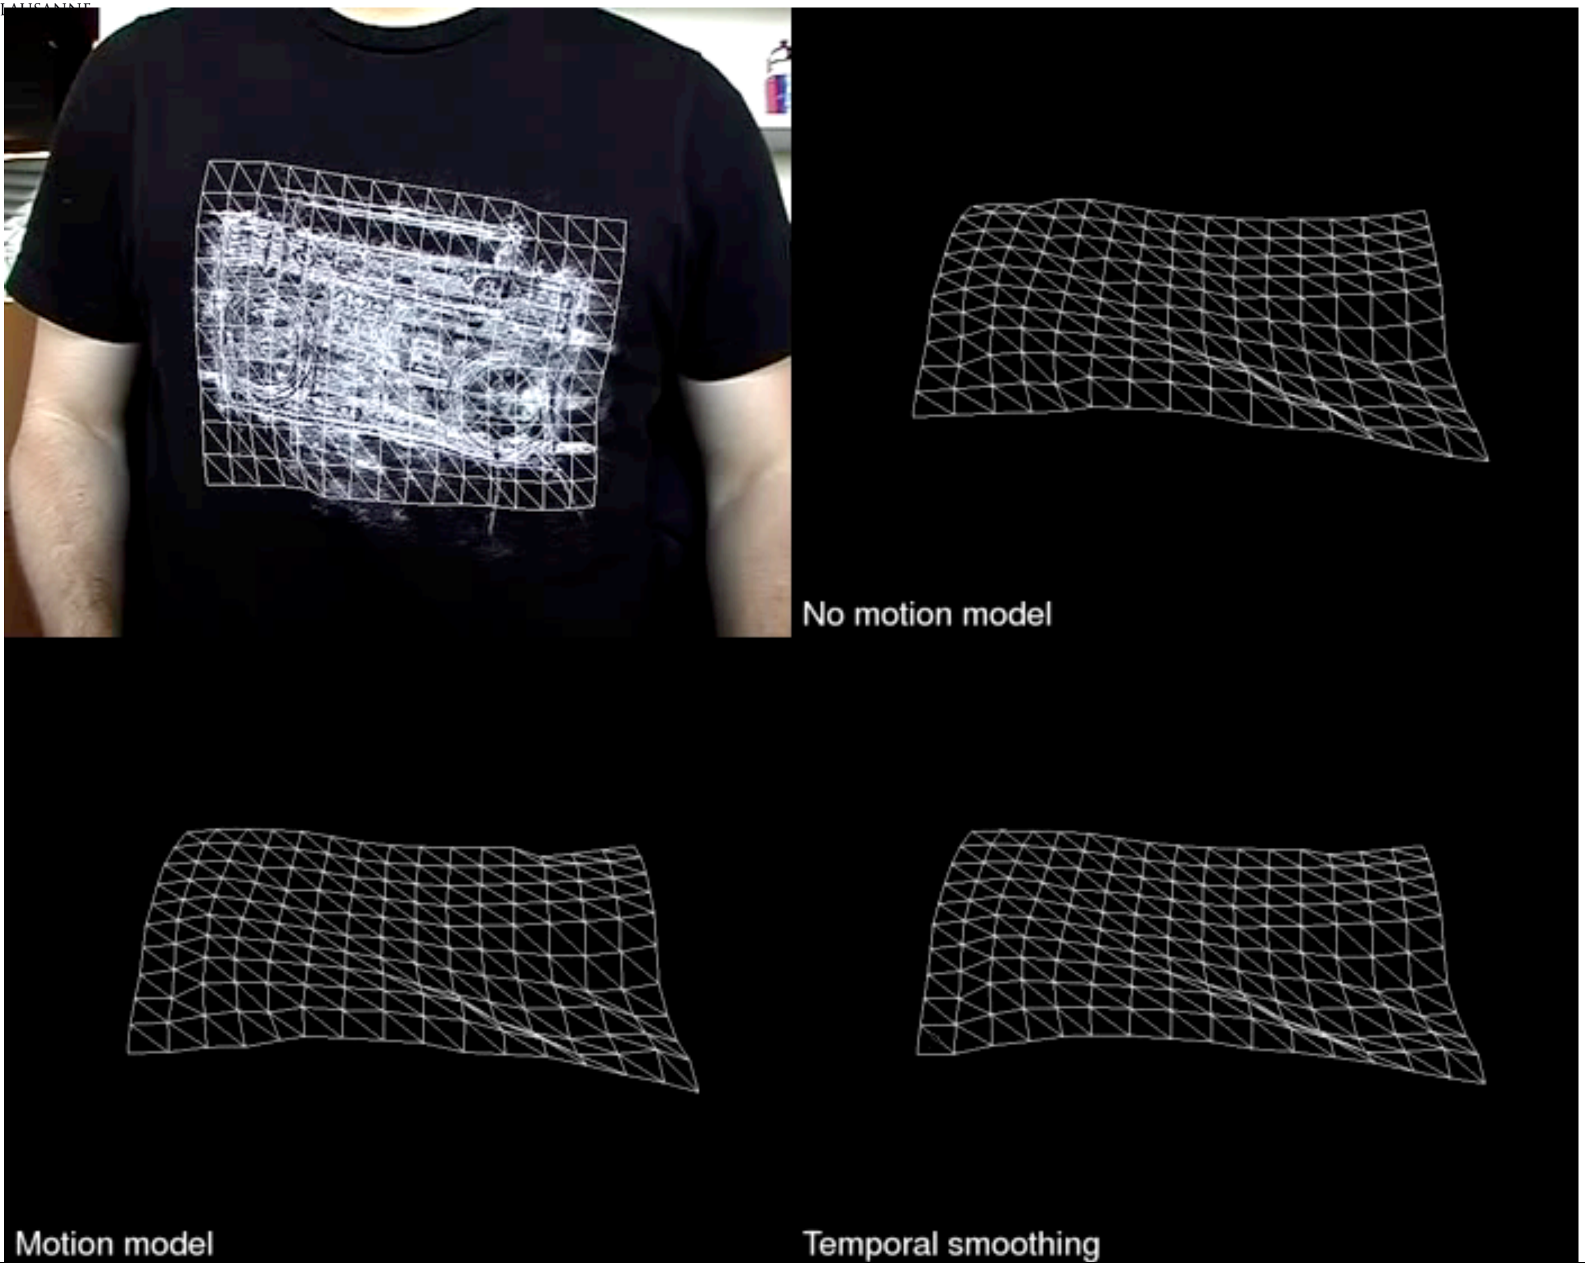
No motion model
Motion model
Temporal smoothing

Repetitive Patterns

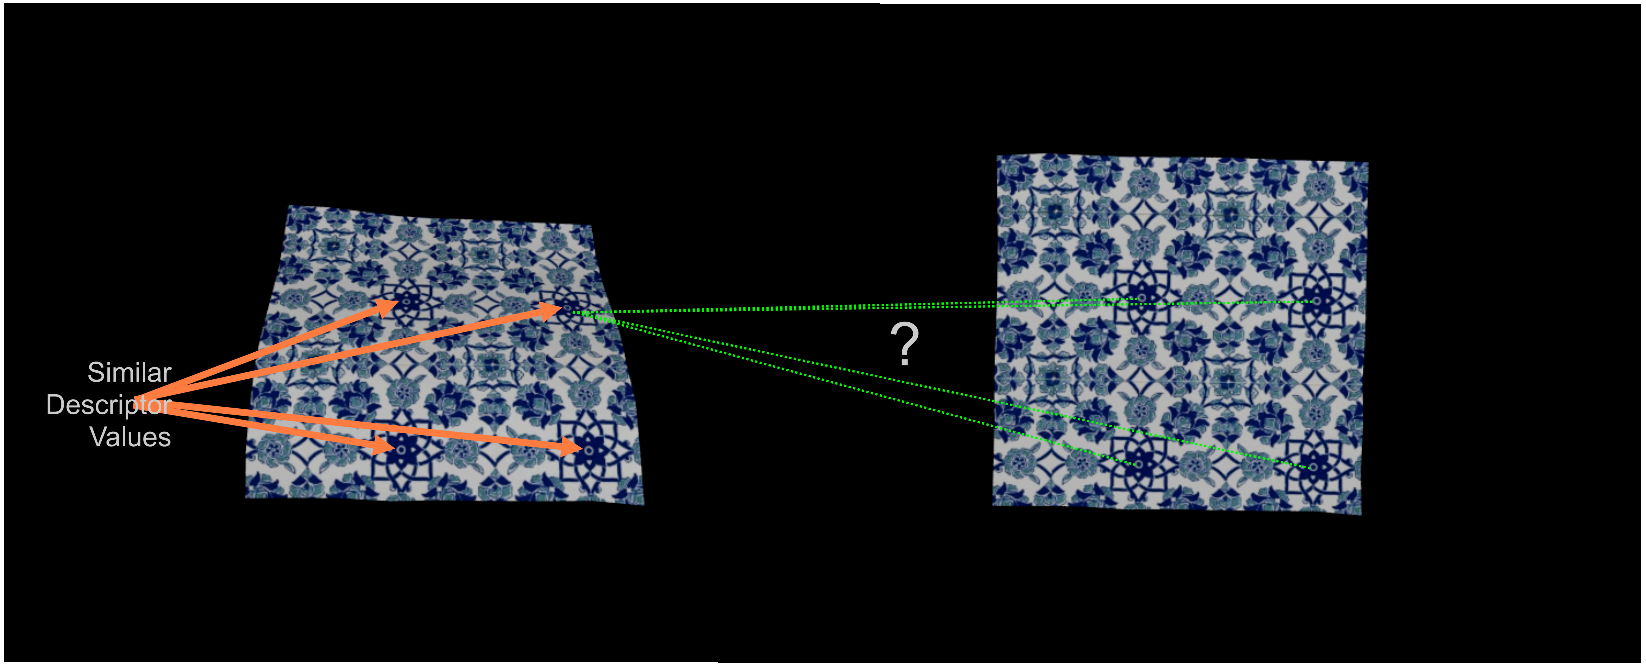

Problem:

• Correspondences are difficult to establish.

Solution:

• Simultaneously solve for correspondences and 3D shape.

Shaji et al., CVPR'10

Mixed Integer Quadratic Problem

$$\min_{\text{matches,reconstruction}} \left(\underbrace{\mathrm{E}(\text{matching score})}_{binary\ terms} + \underbrace{\mathrm{E}(\text{match}(\text{Reconstruction Cost}))}_{mixed\ terms} \right)$$

- Instance of a NP Hard Problem.
- Branch-and-bound methods that works well for this particular problem.

Iterations

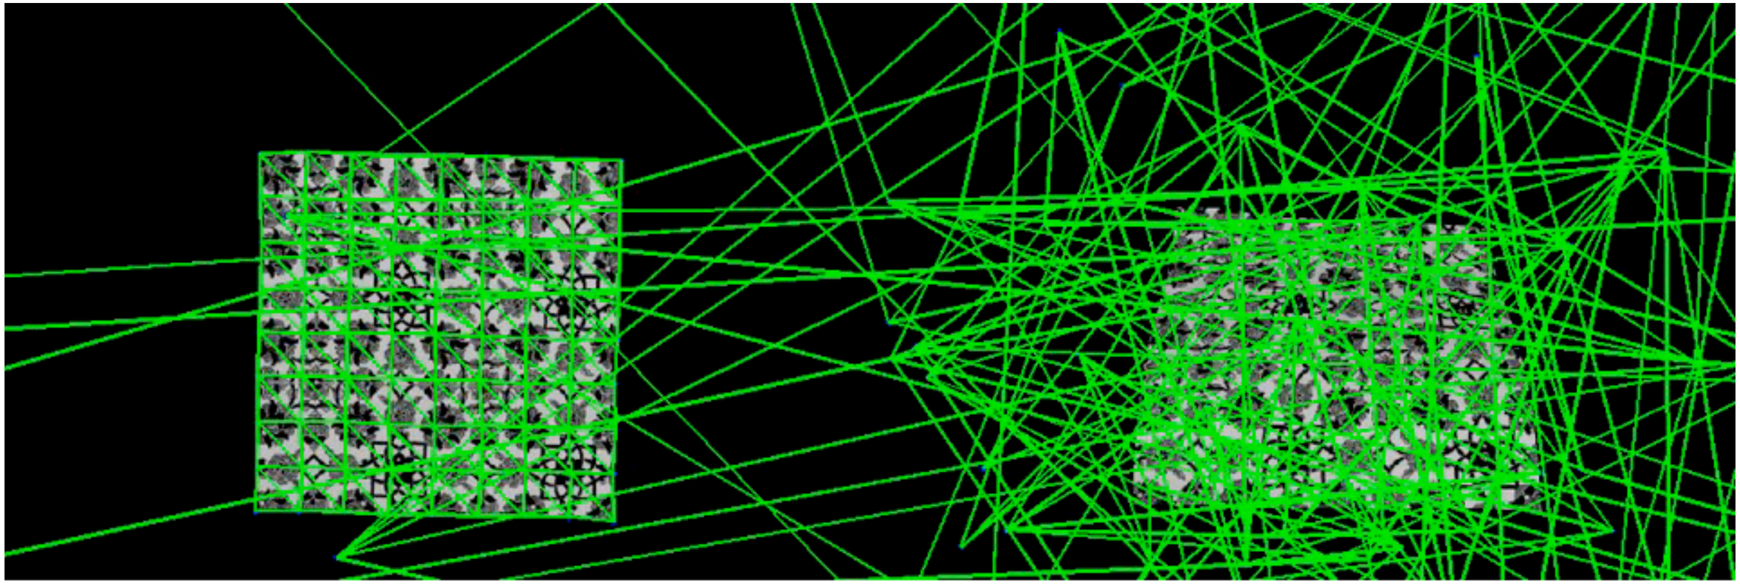

Comparison

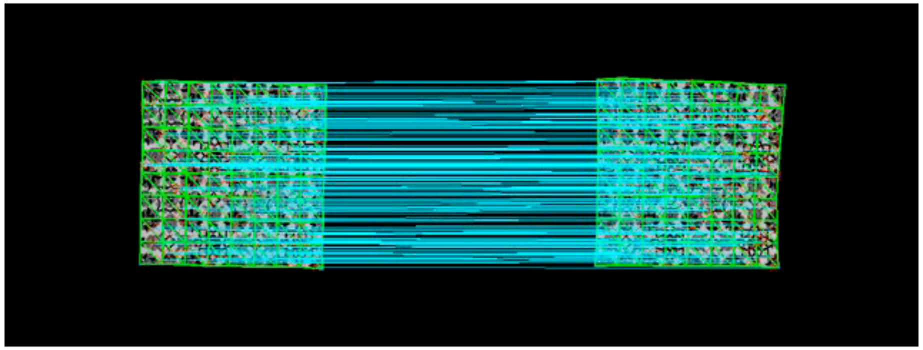

Ground Truth Mesh
Reconstruction by our method
Reconstruction by method of Salzmann et.al

Cushion

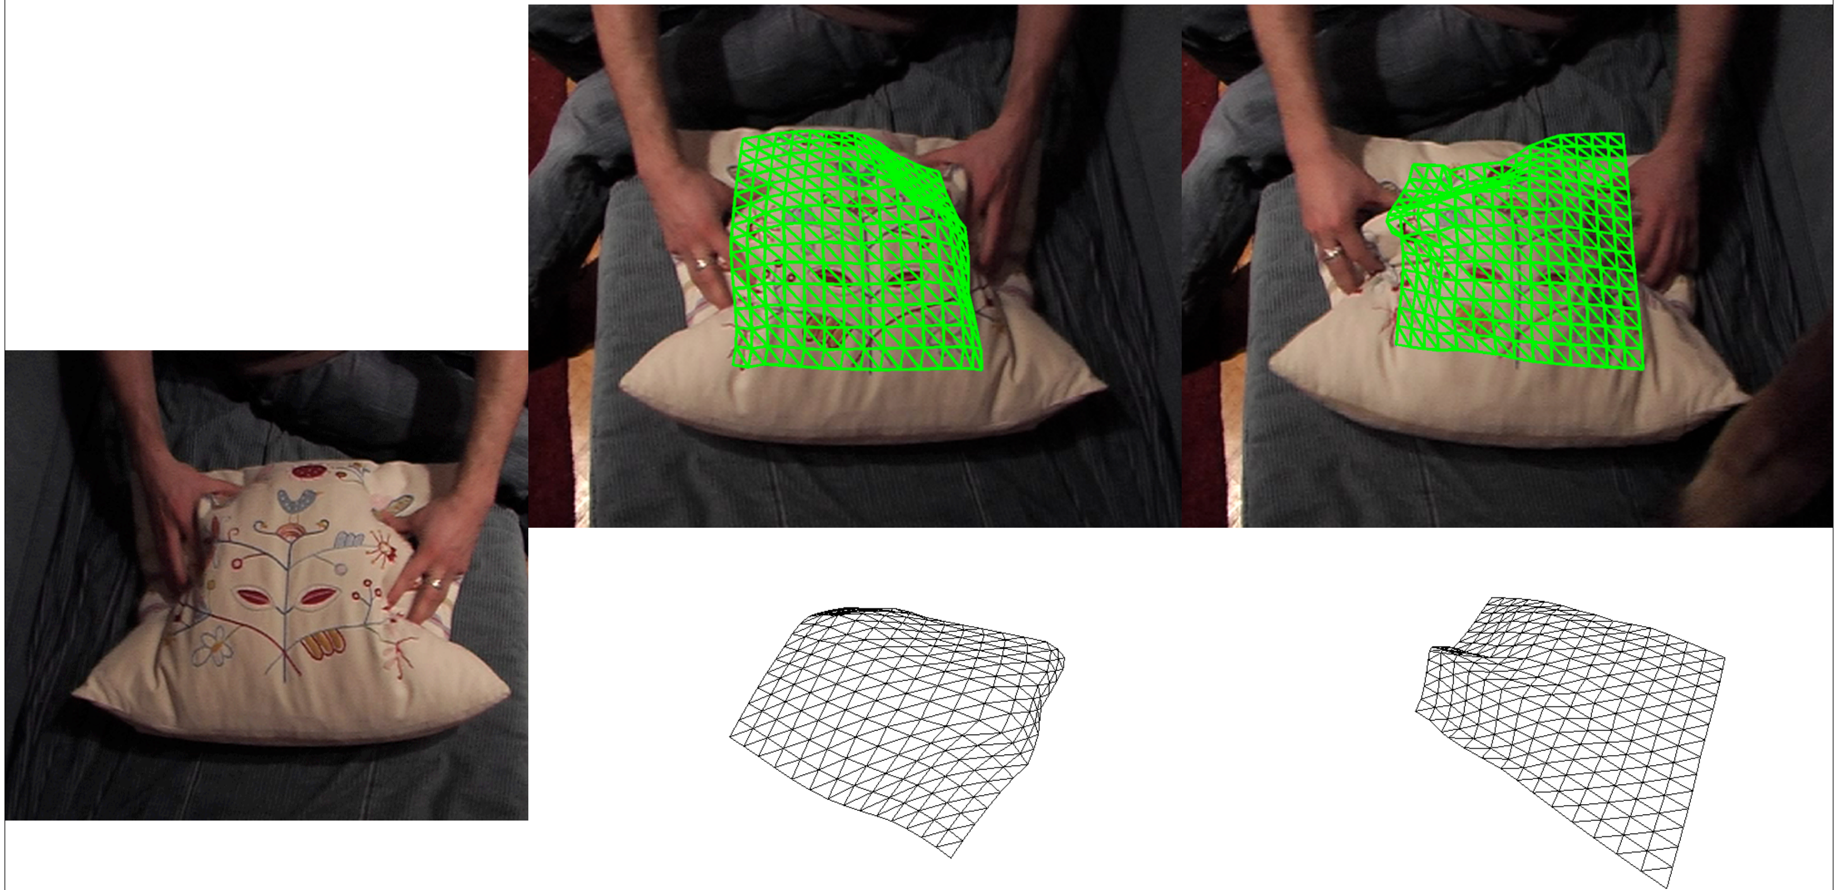

Talk Outline

- Linear Formulation.
- Inextensible surfaces.
- Sharply folding surfaces.
- Eliminating the reference image.
- Applications.

Problem Formulation

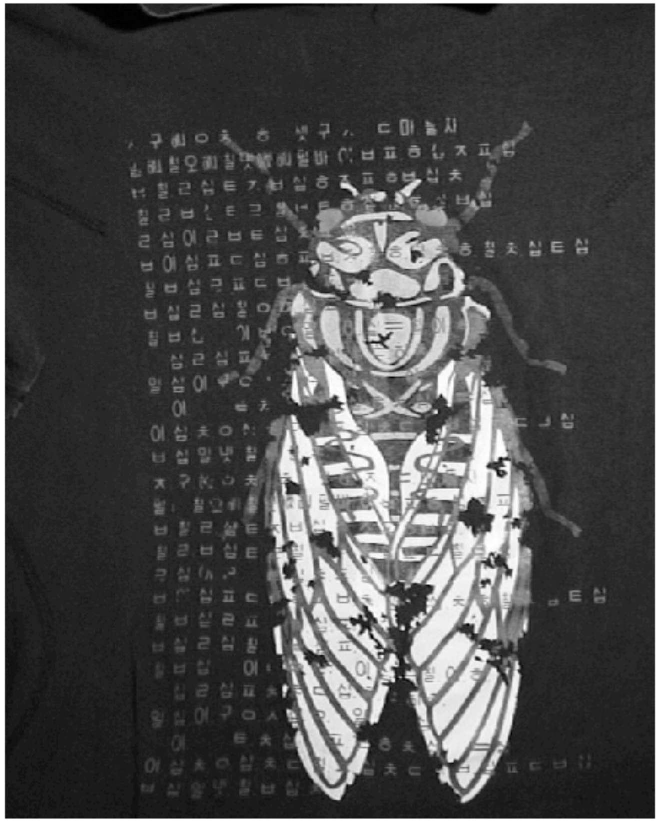

Input Frame

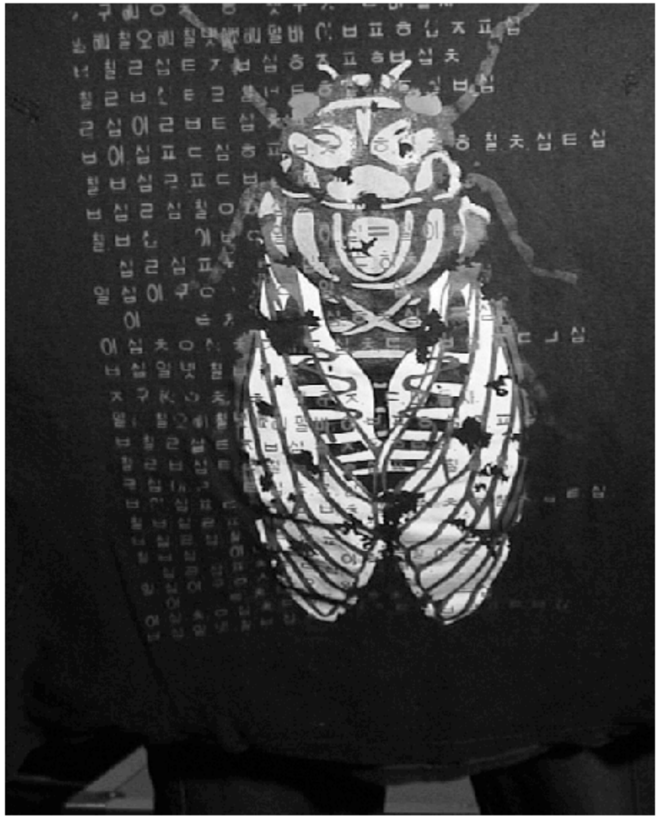

Support Frame

Varol et al., ICCV'09

From Local to Global

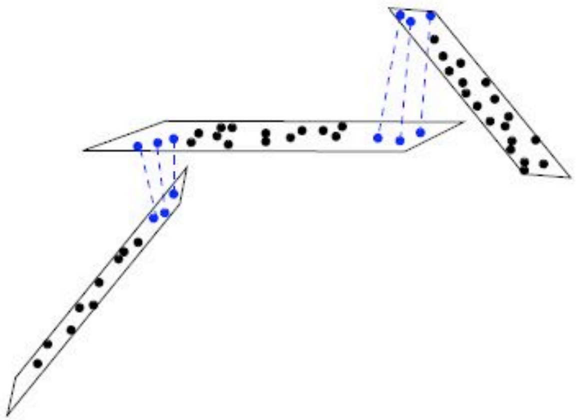

Compute local homographies

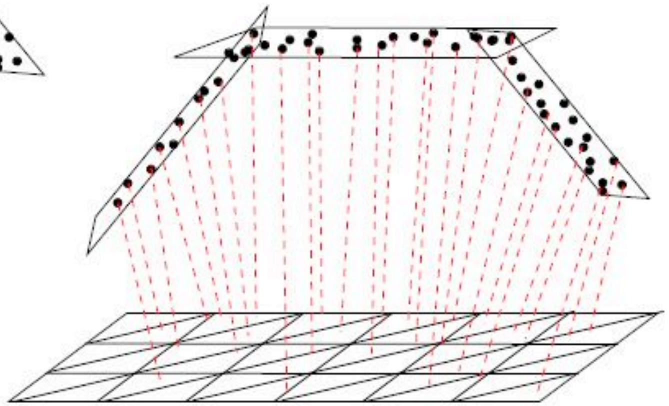

Enforce consistency

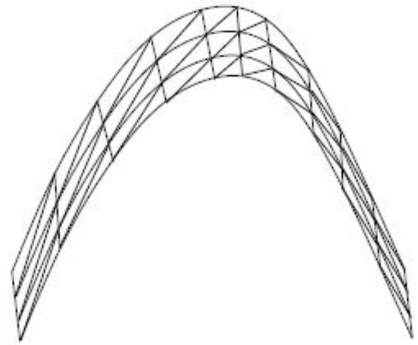

Fit a global surface

Local Homographies

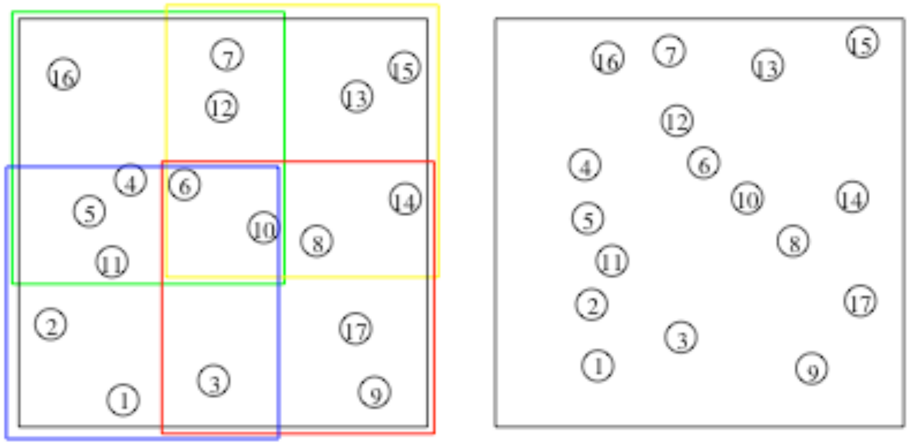

Input Image Support Image

Assuming that the patch is fixed and that the support camera moves

$$\mathbf{P}_i = \mathbf{K}[\mathbf{R}_i|\mathbf{t}_i]$$

$$\mathbf{H}_i = \mathbf{R}_i - \frac{\mathbf{t}_i\mathbf{n}_i^T}{d^i} = \mathbf{R}_i - \mathbf{t}'_i\mathbf{n}_i^T$$

$\rightarrow \mathbf{R}_i$, t_i, and n_i can be recovered up to a scale factor.

Enforcing Consistency

- Turn pairs of matching points in each patch into 3D points that lie on a plane by solving a linear system per patch. The reconstruction is performed up to a local scale factor.

- Compute the scale factors so that points shared by several patches have the same 3D coordinates by solving a global linear system.

--> 3D cloud up to a global scale factor.

Consistent Point Cloud

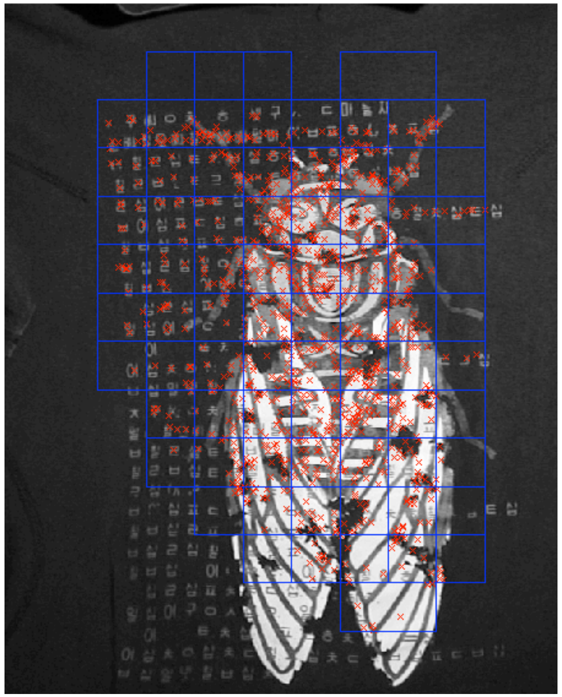

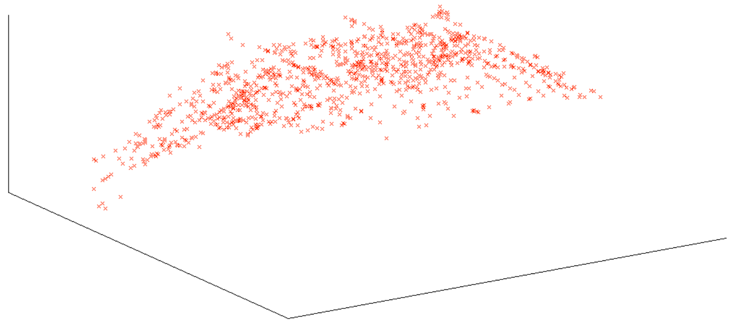

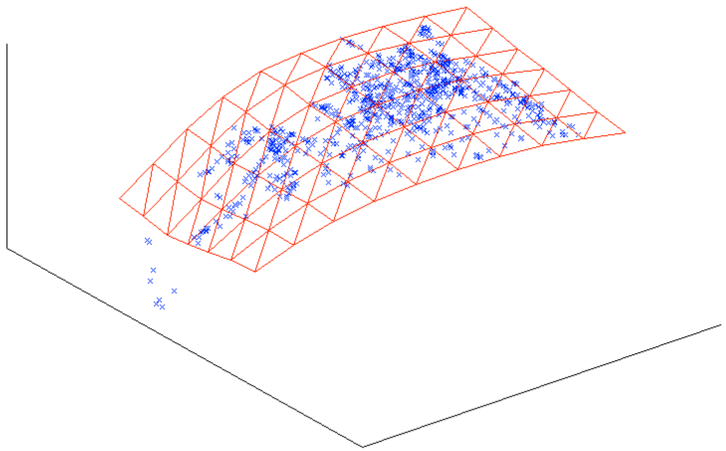

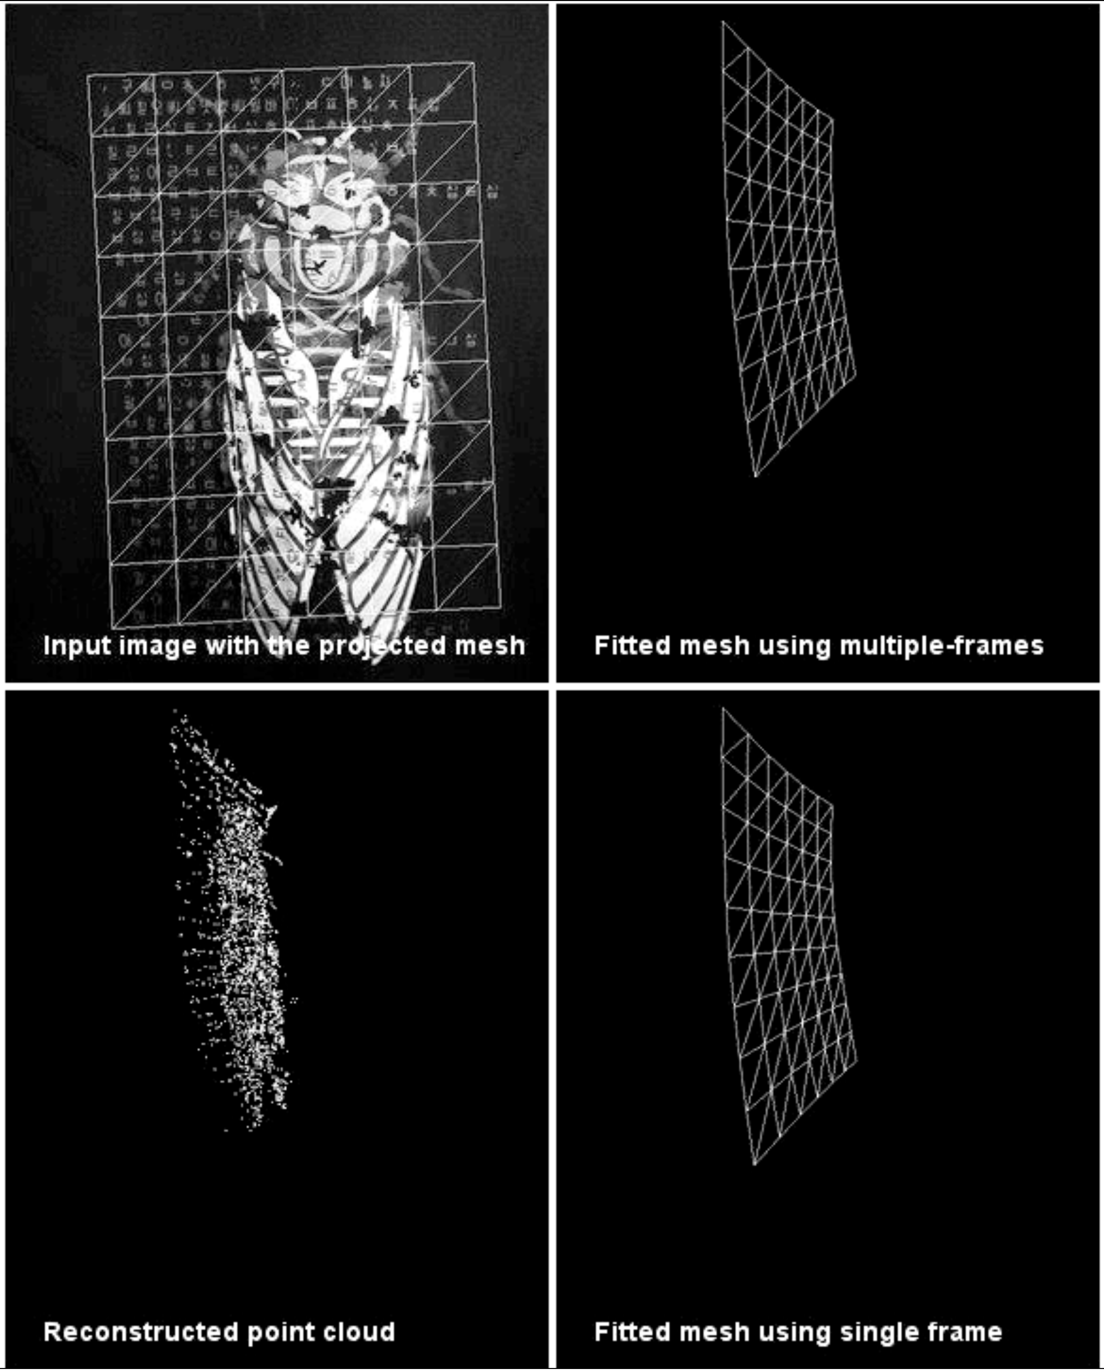
Input image with the projected mesh
Fitted mesh using multiple-frames
Reconstructed point cloud
Fitted mesh using single frame

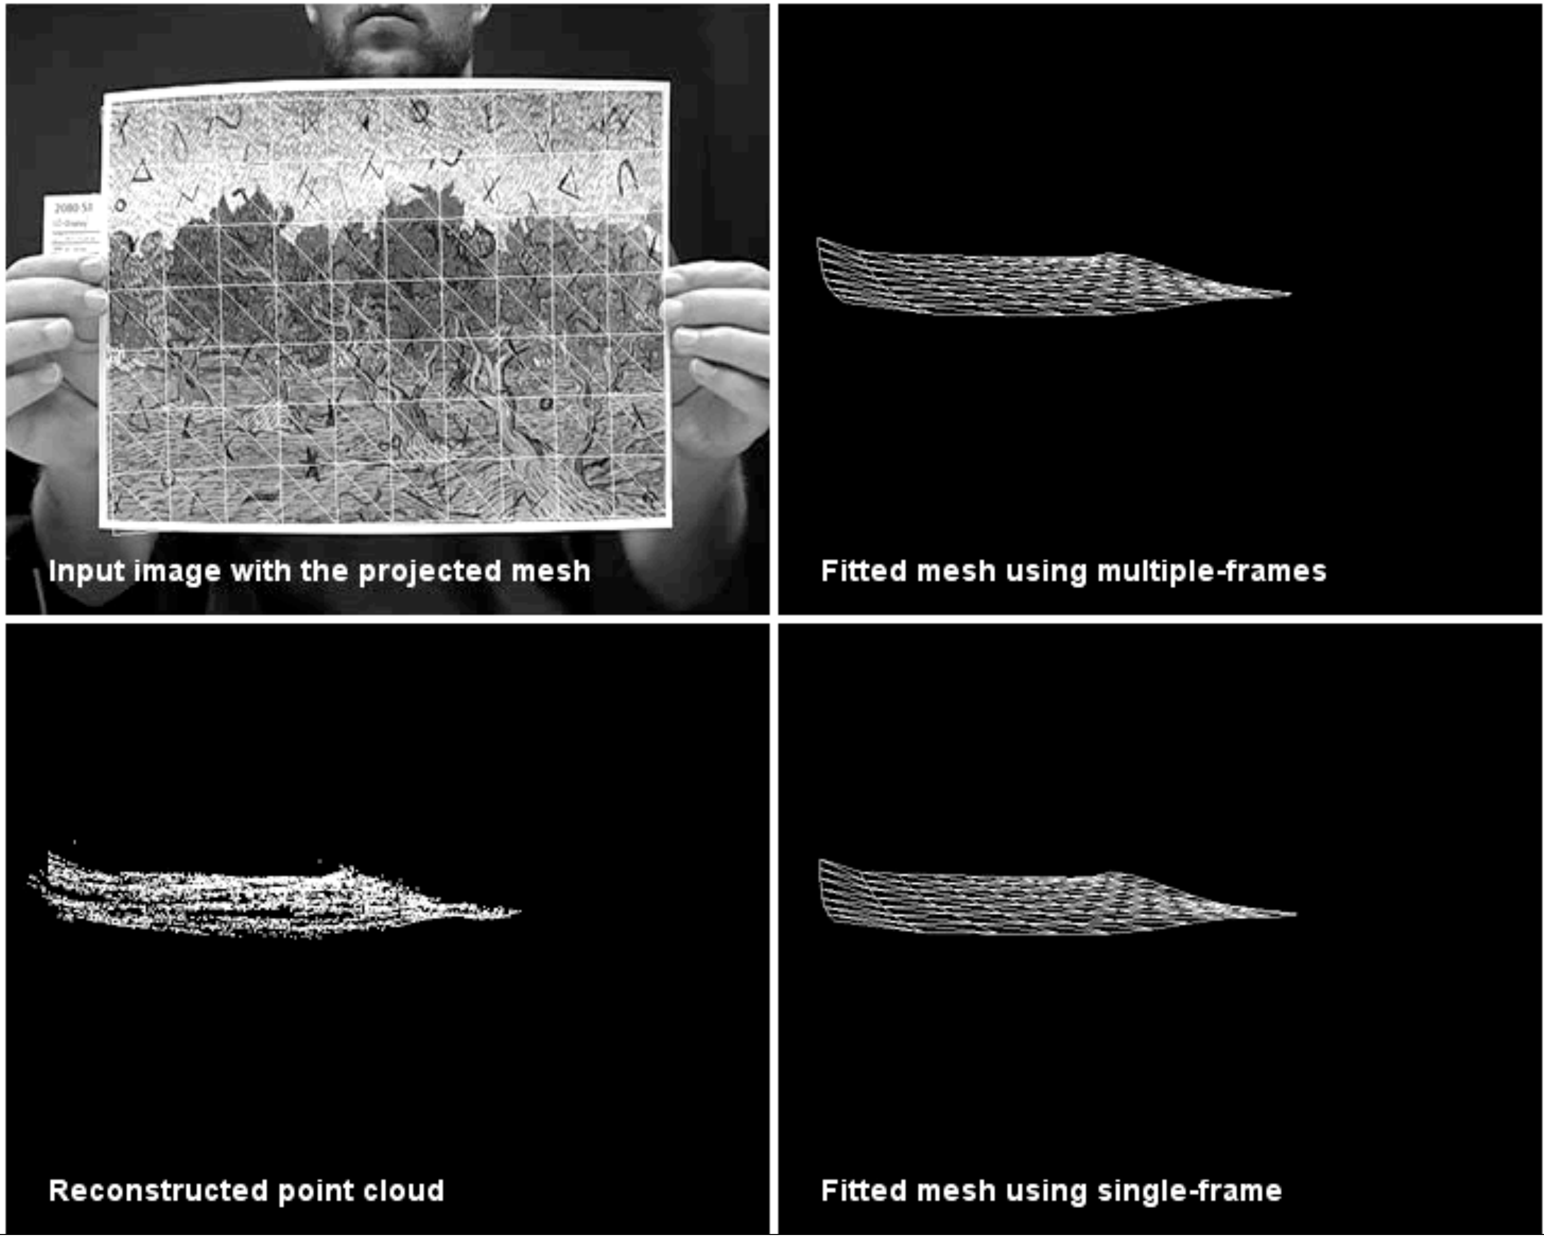
Input image with the projected mesh
Fitted mesh using multiple-frames
Reconstructed point cloud
Fitted mesh using single-frame

Triangle Soup

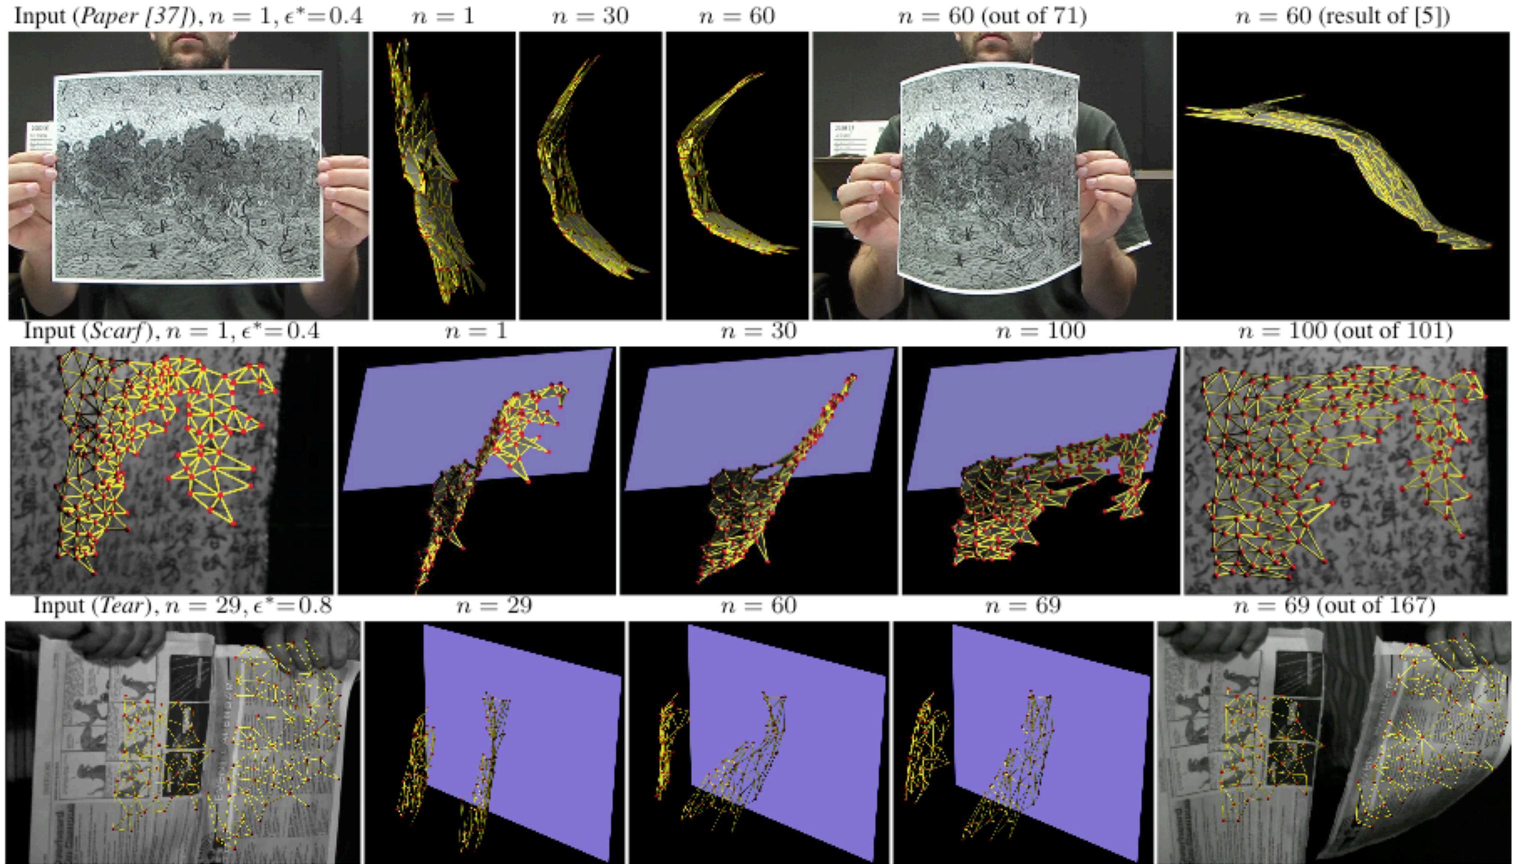

1. Track triangle vertices over 4 or more frames.
2. Assume edge-length is preserved and reconstruct in 3D.
3. Enforce consistency of the resulting *triangle soup*.

Taylor, Jepson, and Kutulakos, CVPR'10

3D Deformable Surfaces

- Linear Formulation.
- Inextensible surfaces.
- Sharply folding surfaces.
- Eliminating the reference image.
- Applications.

ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

Augmented Reality

Magic Book

Magic Cushion

Scherrer et al. , Leonardo, special issue SIGGRAPH'09

Hydroptère

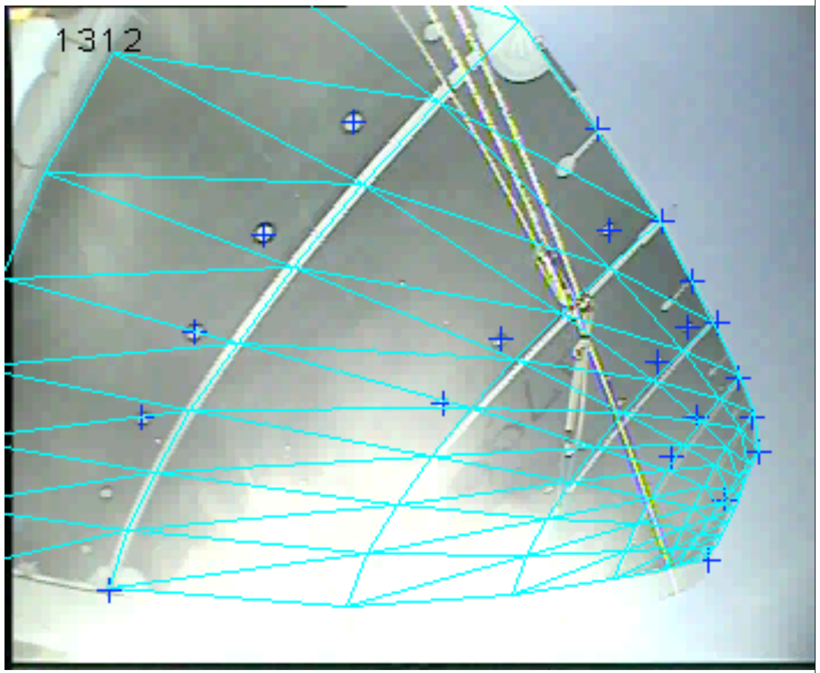

Runs at 8 Hz on an ordinary PC

Alinghi

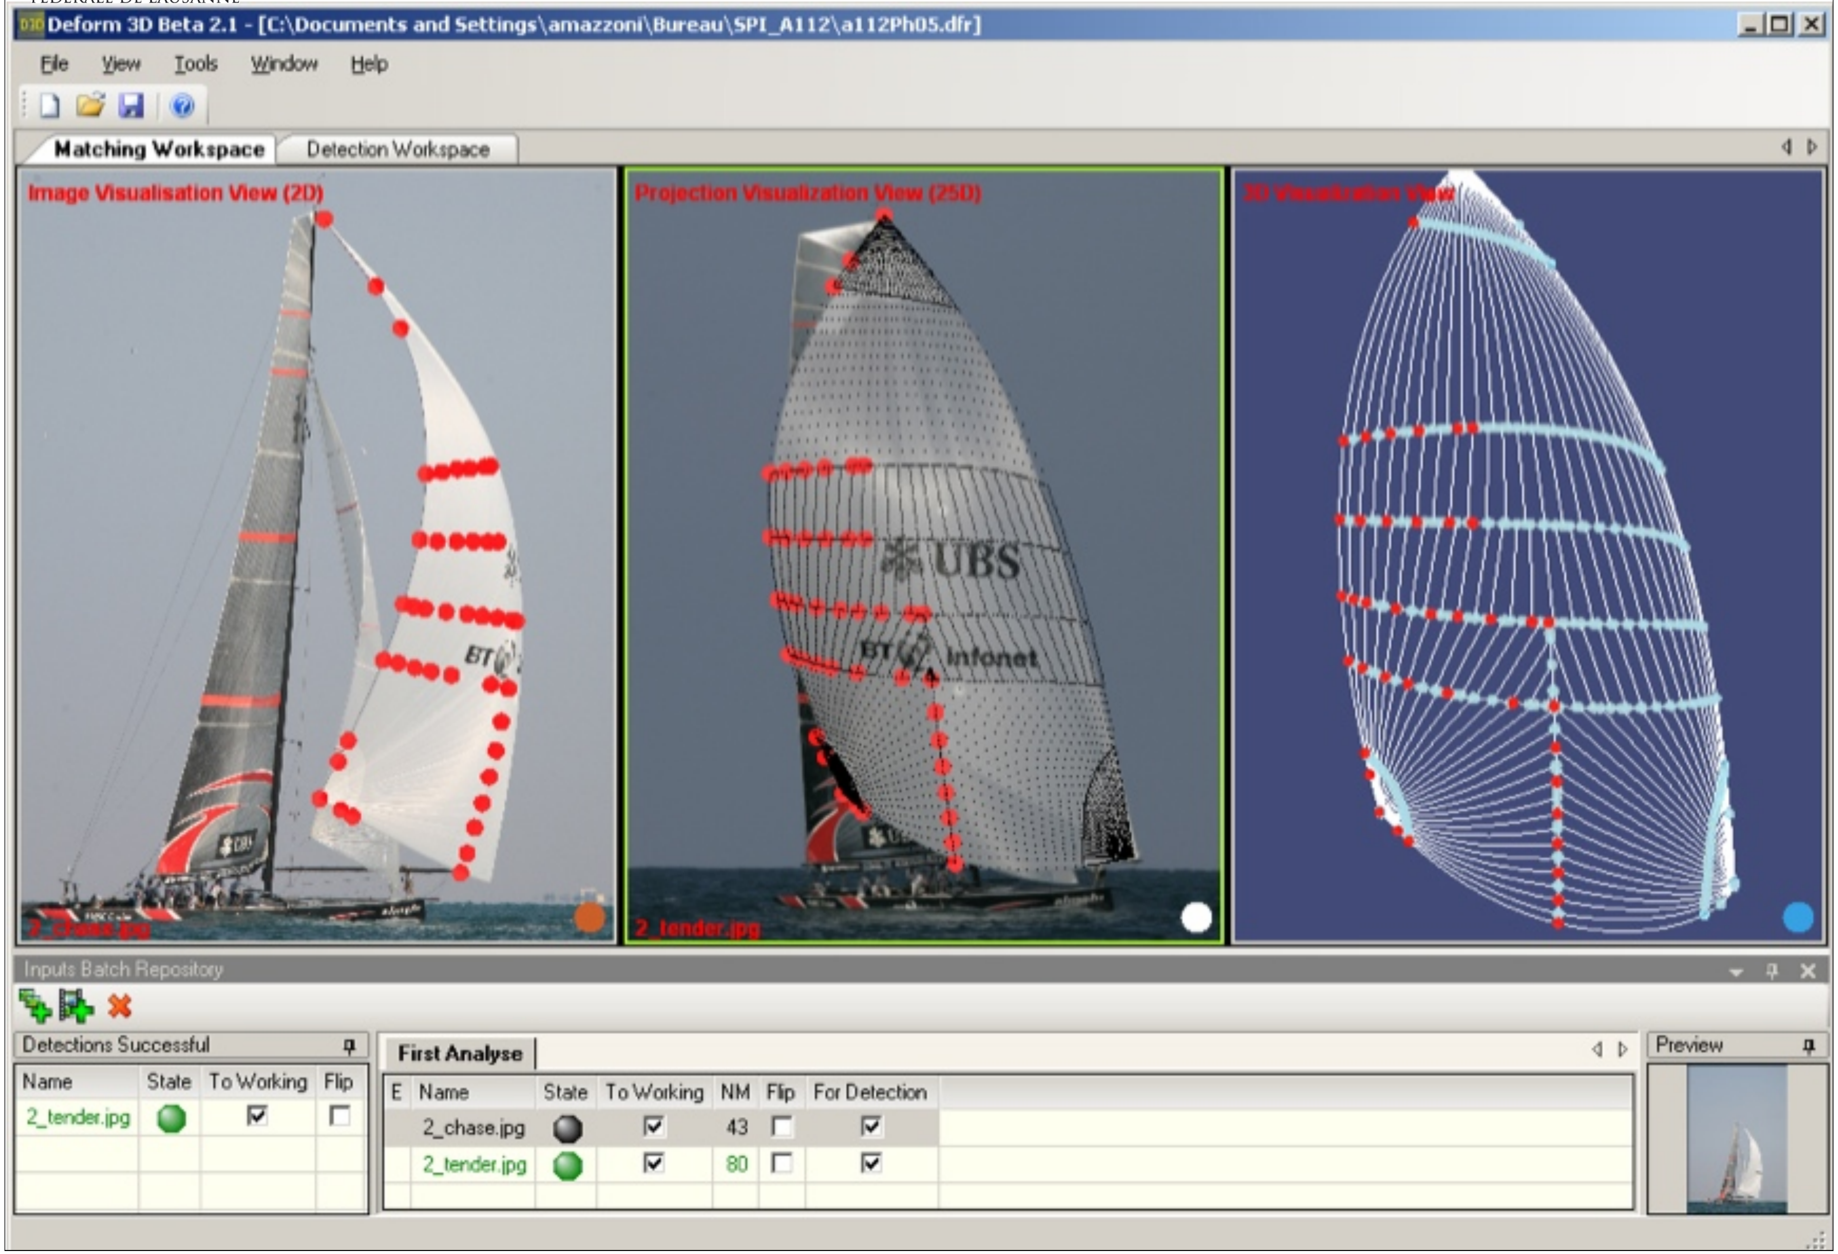
Deform 3D Beta 2.1 - [C:\Documents and Settings\amazzoni\Bureau\SPI_A112\a112Ph05.dfr]
File
View
Tools
Window
Help
Matching Workspace
Detection Workspace
Image Visualisation View (2D)
Projection Visualization View (25D)
UBS
Infonet
2_tender.jpg
Inputs Batch Repository
Detections Successful
Name
State
To Working
Flip
2_tender.jpg
First Analyse
E
Name
State
To Working
NM
Flip
For Detection
2_chase.jpg
43
2_tender.jpg
80
Preview

Wing Deformation

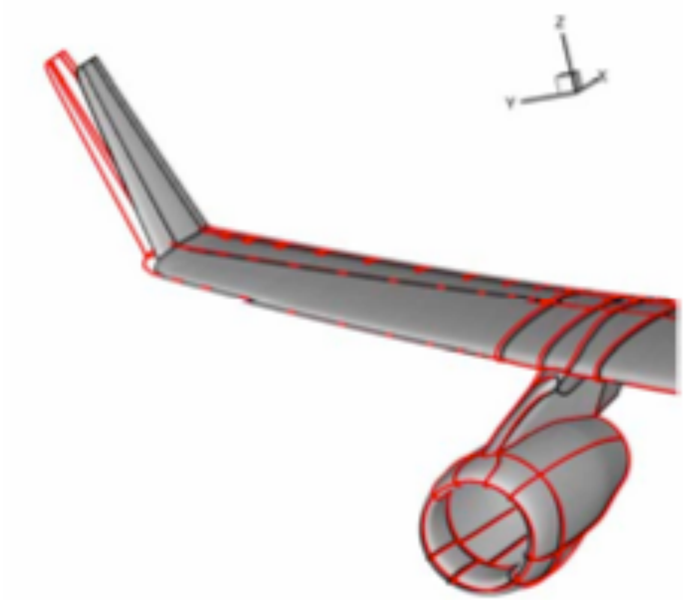

- Compare predicted and observed values.
- Improve simulation software until the two match.

--> Virtual wind tunnel.

Wing Deformation

Intelligence Gathering

- Automated reading of those banners requires unwarping the surfaces.

Laparoscopic Surgery

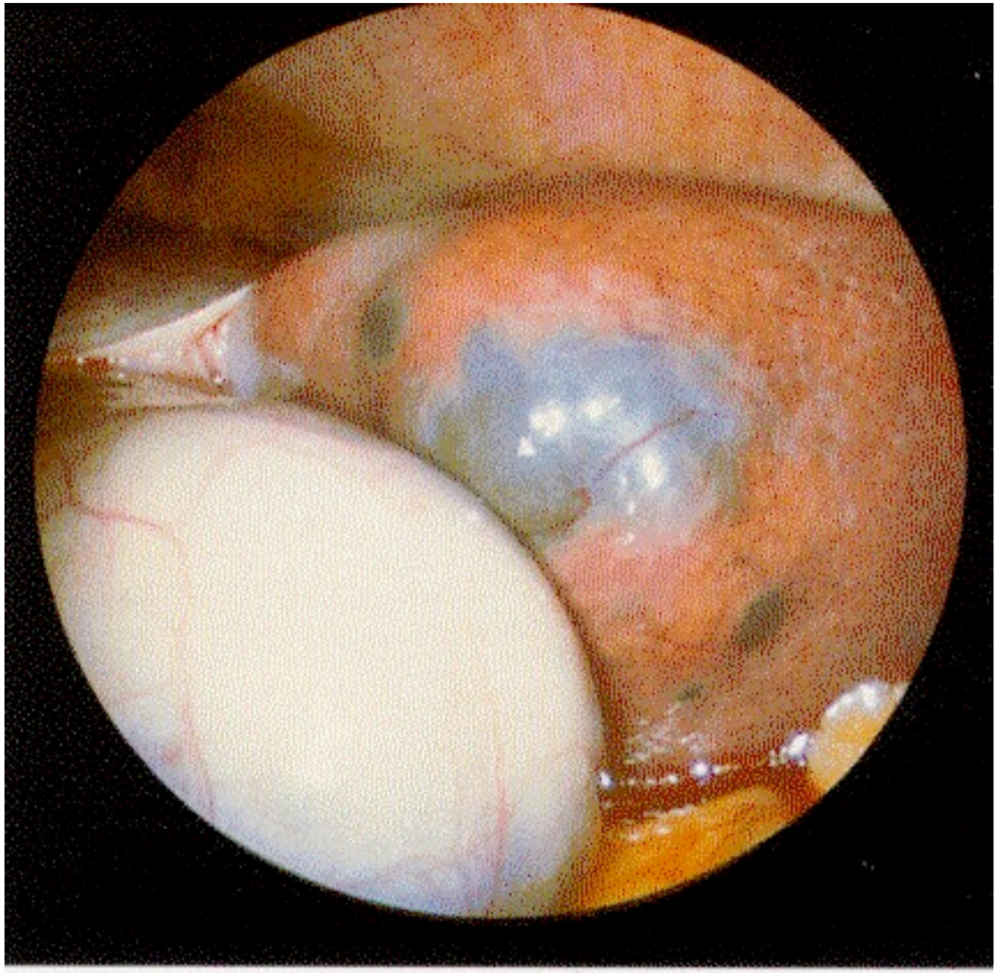

A Generic Paradigm

Automated 3D deformable surface detection:

- Reconstruct textured parts of a surface.
- Learn a deformation model from those.
- Apply it to reconstruct the rest of the surface.

→ A robust method that is easy to deploy.

ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

CVLab

Thanks to

- M. Calonder
- R. Hartley
- S. Ilic
- **V. Lepetit**
- F. Moreno-Noguer
- M. Ozuysal
- **J. Pilet**
- **M. Salzmann**
- A. Shaji
- E. Tola
- R. Urtasun
- A. Varol

References

2D Deformable Surfaces

• J. Pilet, V. Lepetit, and P. Fua, **Fast Non-Rigid Surface Detection, Registration and Realistic Augmentation**, International Journal of Computer Vision, Vol. 76, Nr. 2, February 2008.

3D Deformable Surfaces

• M. Salzmann, J.Pilet, S.Ilic, P.Fua, **Surface Deformation Models for Non-Rigid 3--D Shape Recovery**, IEEE Transactions on Pattern Analysis and Machine Intelligence, Vol. 29, Nr. 8, pp. 1481 - 1487, August 2007.

• M. Salzmann and P. Fua, **Linear Local Models for Monocular Reconstruction of Deformable Surfaces**, IEEE Transactions on Pattern Analysis and Machine Intelligence, 2011, In Press.

Point Matching:

• M. Ozuysal, M. Calonder, V. Lepetit and P. Fua, **Fast Keypoint Recognition using Random Ferns**, IEEE Transactions on Pattern Analysis and Machine Intelligence, Vol. 32, Nr. 3, pp. 448 - 461, March 2010

• M. Calonder, V. Lepetit and P. Fua, **BRIEF: Binary Robust Independent Elementary Features**, European Conference on Computer Vision, Heraklion, Greece, 2010.