Automatic Landmarking for Non-cooperative 3D Face Recognition

Clement Creusot

Department of Computer science
THE UNIVERSITY of York

Past research presentation, November 2012

PLAN:

- Background
- PhD Motivation
- Problem(s)
- Keypoint Detection
- Labeling
- Model Generation (if we have time)
- Conclusion

Background

Background

- Born in France
- Studied till highschool in Grasse near Cannes (General Science Baccalaureat)
- Prepa Math and Physics at Lycee Massena in Nice

Background

- Born in France
- Studied till highschool in Grasse near Cannes (General Science Baccalaureat)
- Prepa Math and Physics at Lycee Massena in Nice

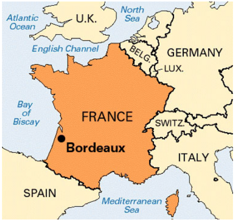

- ENSEIRB engineering school in Bordeaux toward a MEng
- University of Bordeaux toward a MSc (dual-diploma)
- Interships in web apps automatic testing (Epistema) and aeronautics (C-S)

Background

ENSEIRB

- ENSEIRB engineering school in Bordeaux toward a MEng
- University of Bordeaux toward a MSc (dual-diploma)
- Interships in web apps automatic testing (Epistema) and aeronautics (C-S)

Background

UNIVERSITÉ BORDEAUX 1 Sciences Technologies

- ENSEIRB engineering school in Bordeaux toward a MEng
- University of Bordeaux toward a MSc (dual-diploma)
- Interships in web apps automatic testing (Epistema) and aeronautics (C-S)

Background

Background
PhD Motivation
Problem(s)
Keypoint detection
Labeling
Model Generation
Conclusion

- Final MEng thesis (CEA-DAM)
- Final MSc thesis (CEA-DAM)
- University of York (UK) toward a PhD
- Final PhD thesis on “Automatic 3D Landmarking for Non-Cooperative 3D Face Recognition”

Background

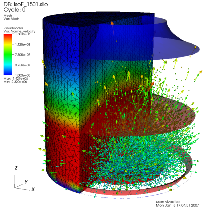

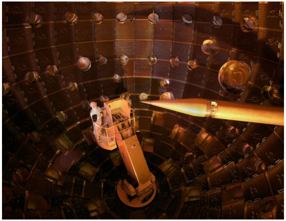

- Final MEng thesis (CEA-DAM)
- Final MSc thesis (CEA-DAM)
- University of York (UK) toward a PhD
- Final PhD thesis on "Automatic 3D Landmarking for Non-Cooperative 3D Face Recognition"

THE UNIVERSITY of York

- Final MEng thesis (CEA-DAM)
- Final MSc thesis (CEA-DAM)
- University of York (UK) toward a PhD
- Final PhD thesis on “Automatic 3D Landmarking for Non-Cooperative 3D Face Recognition”

CYBULA
> high performance pattern recognition systems

- Cybula Ltd (FaceEnforce ic3D Camera)
  - Capture moving subject
  - From 3.5 to 5 meters
  - Up to 16 frames per second
  - Near Infrared pulse laser projection
  - Works in full sunlight
- EVAN project (FP6 Marie Curie Fellowship)

Background

EVAN
setting landmarks in science

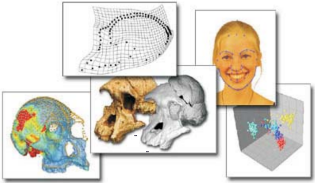

- Cybula Ltd (FaceEnforce ic3D Camera)
  - Capture moving subject
  - From 3.5 to 5 meters
  - Up to 16 frames per second
  - Near Infrared pulse laser projection
  - Works in full sunlight
- EVAN project (FP6 Marie Curie Fellowship)

PhD Motivation

Non-cooperative Recognition at a distance

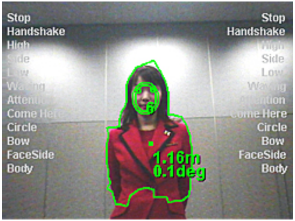

- Application:
  - Surveillance
  - Human-Machine Interaction
- Problems:
  - Pose
  - Occlusion
  - Speed

Non-cooperative Recognition at a distance

From [Savran et al., 2008]

From [Savran et al., 2008]

- Application:
  - Surveillance
  - Human-Machine Interaction
- Problems:
  - Pose
  - Occlusion
  - Speed

Modality

- Non-Cooperative $\supset$ Anti-cooperative
- Proved possible for big database


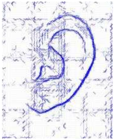

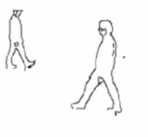

[Proenca, 2008] [Yan and Bowyer, 2007] [Phillips et al., 2005] [Havasi et al., 2007]

Modality

- Non-Cooperative $\supset$ Anti-cooperative
- Proved possible for big database

[Proenca, 2008]

[Yan and Bowyer, 2007]

[Phillips et al., 2005]

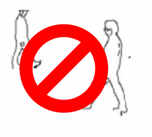

[Havasi et al., 2007]

Modality

- Non-Cooperative $\supset$ Anti-cooperative
- Proved possible for big database




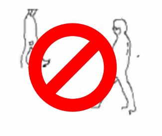

[Proenca, 2008] [Yan and Bowyer, 2007] [Phillips et al., 2005] [Havasi et al., 2007]

- 2D or 3D ?

From [Liu et al., 2007]

3D over 2D

[Phillips et al., 2011]

3D over 2D

[Phillips et al., 2011]

- Verification experiment $\rightarrow$ all existing systems fail
- Is the problem solvable with 2D data only?

Difficult Cases

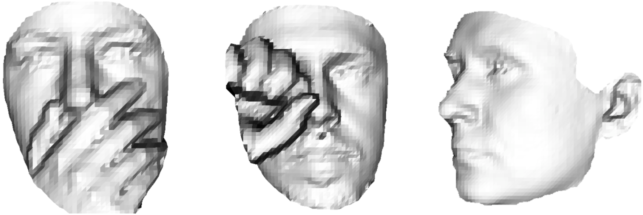

- Recognition:
  - Holistic methods → require Registration
  - Feature based methods → require Feature Localisation
- Will often fail at preprocessing
  - Naive methods for feature detection (expert systems)
  - Strong assumptions on the input (cannot be used for non-cooperative cases)
- Conclusion: Feature detection seems to be a bottleneck

Difficult Cases

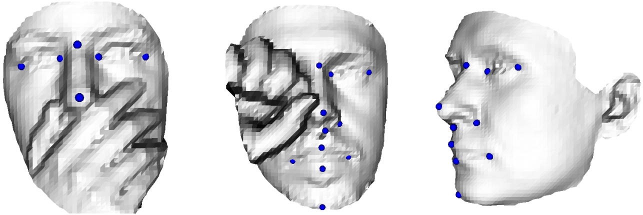

- Recognition:
  - Holistic methods → require Registration
  - Feature based methods → require Feature Localisation
- Will often fail at preprocessing
  - Naive methods for feature detection (expert systems)
  - Strong assumptions on the input (cannot be used for non-cooperative cases)
- Conclusion: Feature detection seems to be a bottleneck

Review

Background
PhD Motivation
Problem(s)
Keypoint detection
Labeling
Model Generation
Conclusion

- Almost all methods expect non occluded frontal face
- A few that don't:
  - Allow some pose variation (All expert systems):
    - [Colbry et al., 2005]: Curvature + ICP + Relaxation
    - [Lu and Jain, 2006]: Directional Maximum
    - [Faltemier et al., 2008a]: Rotated Profile Signature
  - Allow some occlusions (Machine Learning System):
    - [Zhao et al., 2011]: 2D + 3D SFAM

[Colbry et al., 2005]

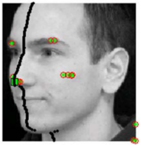

[Lu and Jain, 2006]

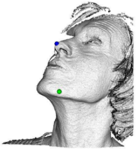

[Faltemier et al., 2008a]

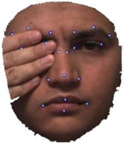

[Zhao et al., 2011]

- Almost all papers expect the nose will be present
- Most papers require two well defined inner corners of the eyes

Gap in research

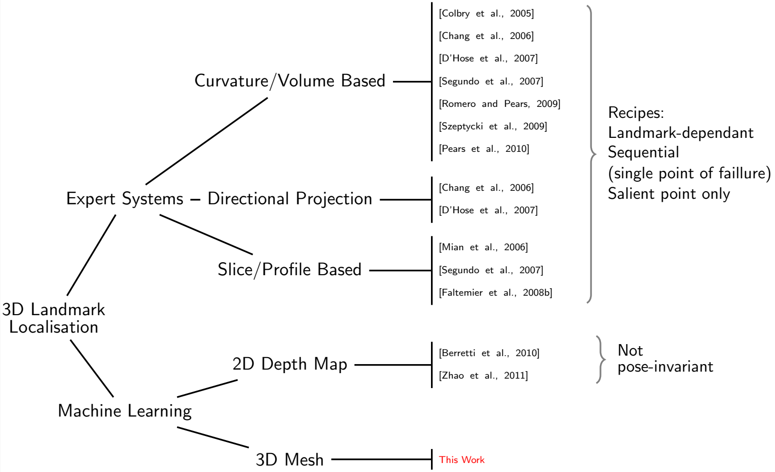

Problem(s)

Problem Breakdown

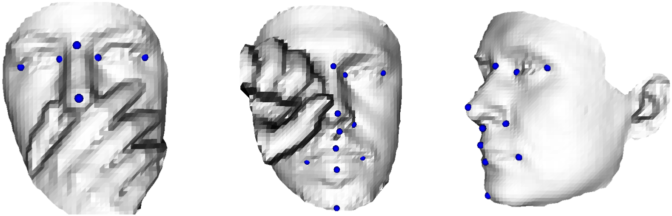

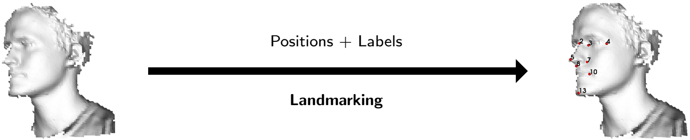

Problem Breakdown

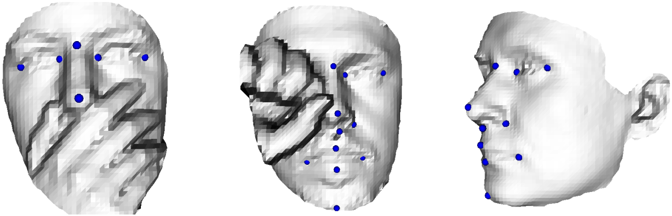

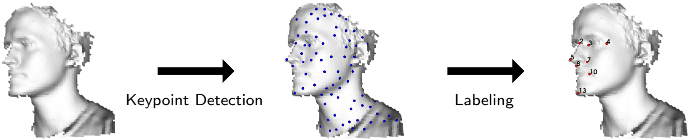

Problem(s)

- Landmarking
  - Keypoint detection
  - Labelling
  - Position refinement
- Hypergraph face representation for recognition
- Model Learning

Problem(s)

In this presentation

- Landmarking
  - Keypoint detection :
    Using On-Manifold Machine Learning techniques
  - Labelling :
    Using Multi-attributed Hypergraph Matching
  - Position refinement :
    Not Discussed Here
- Hypergraph face representation for recognition :
  Not Discussed Here
- Model Learning :
  Using dense registration prior

Databases

- FRGC v2

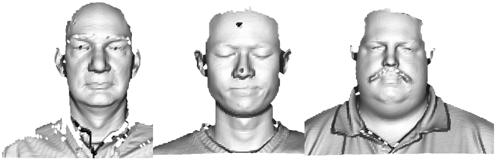

  - 4950 faces from 557 people
  - 200 in training
  - 4750 in test set (3108 Neutral, 1642 Expression)

- Bosphorus

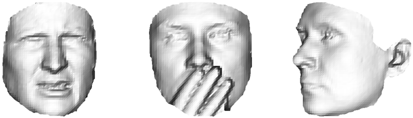

  - 4666 faces from 105 people
  - Occlusion, Expression, Rotation
  - 99 in training (20 for profile)

Learning-based methods for automatic 3D keypoints detection

Aim

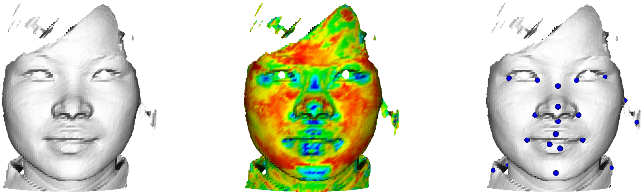

- Keypoints detection (NOT LANDMARKS)
- Similar to any of 14 learnt features (Dictionary of local shapes)

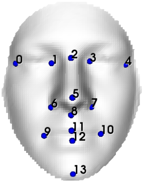

OFFLINE

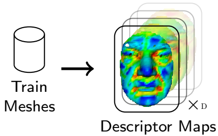

Workflow

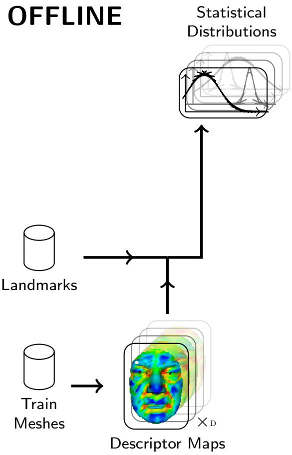
OFFLINE
Statistical Distributions
Landmarks
Train Meshes
× D
Descriptor Maps

Workflow

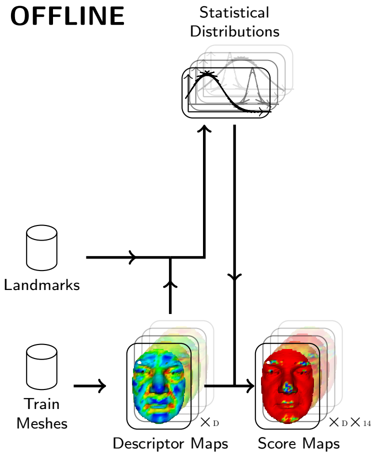
OFFLINE
Statistical
Distributions
Landmarks
Train
Meshes
×D
Descriptor Maps
×D×14
Score Maps

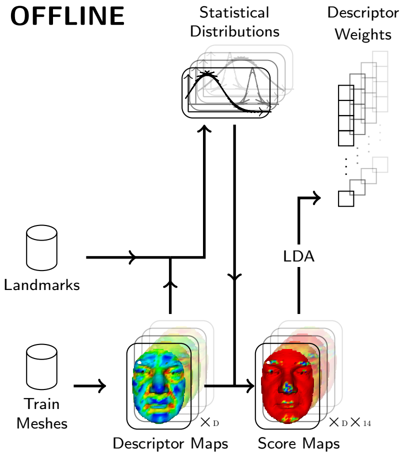
OFFLINE
Statistical Distributions
Descriptor Weights
LDA
Landmarks
Train Meshes
×D
Descriptor Maps
×D×14
Score Maps

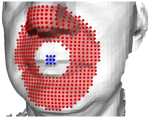

Workflow

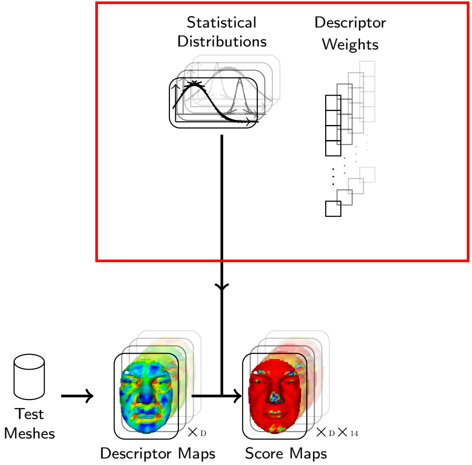

ONLINE

Dictionary of local shapes

Workflow

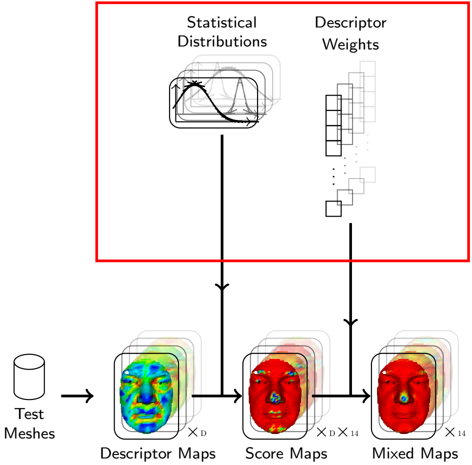

ONLINE

Dictionary of local shapes

Workflow

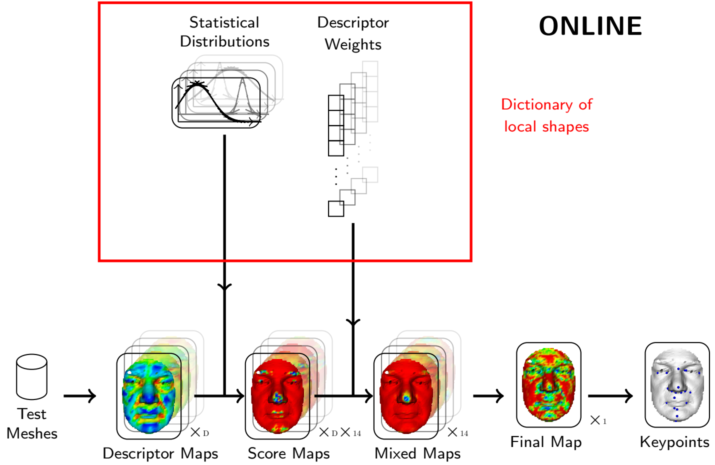
Statistical Distributions
Descriptor Weights
ONLINE
Dictionary of local shapes
Test Meshes
Descriptor Maps
×D
Score Maps
×D×14
Mixed Maps
×14
Final Map
×1
Keypoints

Results

- Sparse selection (max 1%)
- Reapeatable (same subject registration)
  - ∼75% (at 10 mm)
- Close to human hand-placed landmarks
  - average All: ∼85% (at 10 mm)
  - average Nose: ∼99% (at 10 mm)
  - average Eyes: ∼90% (at 10 mm)
- High proportion of the local shapes retreived
  - ∼11.88/14 (at 10 mm)

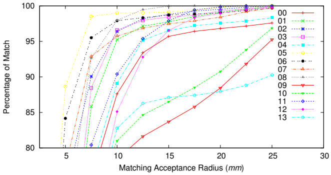

Configuration 1

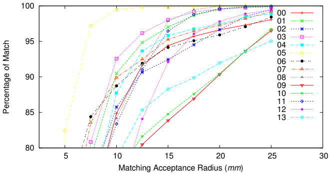

Configuration 2

Examples

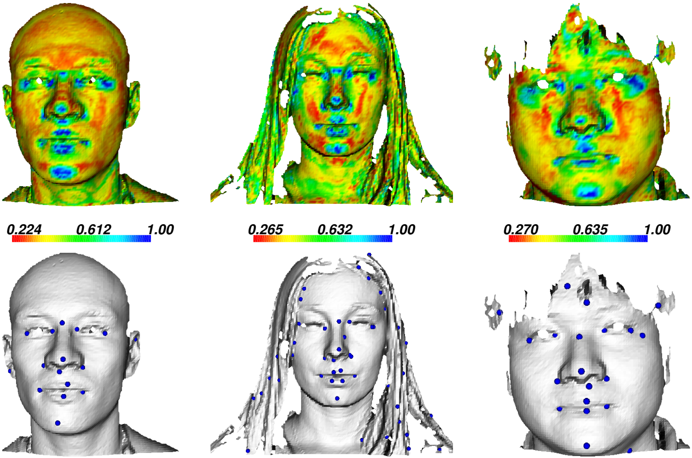

- Good points:
  - Detects "weak" features
  - No single-point-of-failure design
- Limitations:
  - Can be time consuming
    1s (8 desc.)
  - Linear combination of scores

Extending the method

	LDA Scoring	Adaboost Scoring
Nose (05)	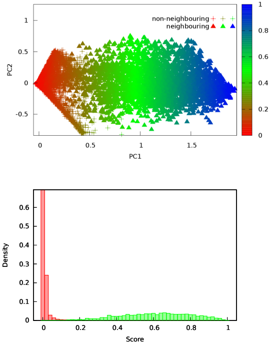	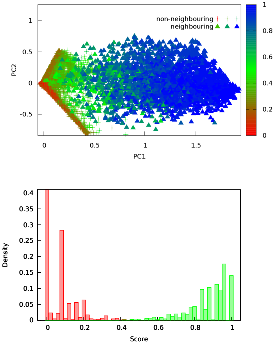

Limitations

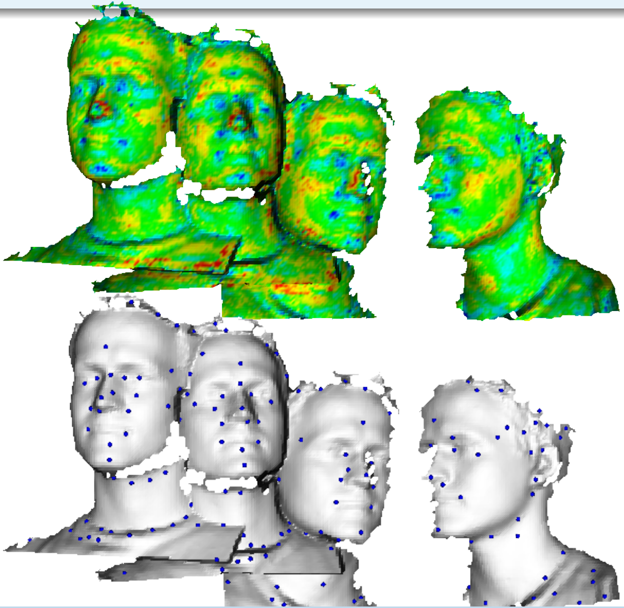

Landmark labeling using multi-attributed hypergraph matching techniques

The landmark Localisation Problem

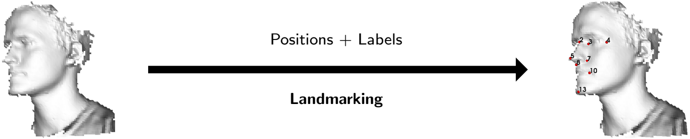

- Landmark = Position + Label
- Two Approaches:
  - Select One Label + Find Corresponding Position
  - Find All Positions + Find Corresponding Labels

The landmark Localisation Problem

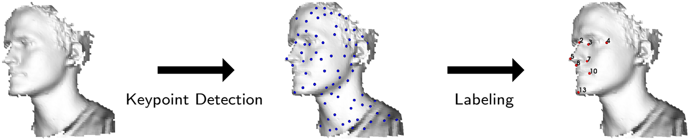

- Landmark = Position + Label
- Two Approaches:
  - Select One Label + Find Corresponding Position
  - Find All Positions + Find Corresponding Labels

The landmark Localisation Problem

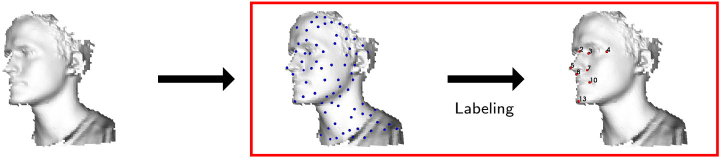

- Landmark = Position + Label
- Two Approaches:
  - Select One Label + Find Corresponding Position
  - Find All Positions + Find Corresponding Labels

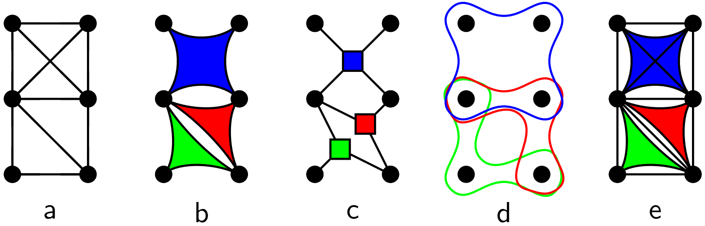

Examples of hypergraph representations:

(a) 2-uniform hypergraph (just a graph)

(b) Non-uniform hypergraph

(c) Bipartite graph representation

(d) Set representation

(e) Non-uniform hypergraph

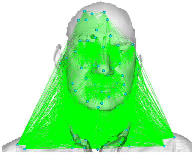

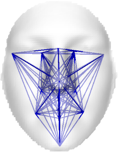

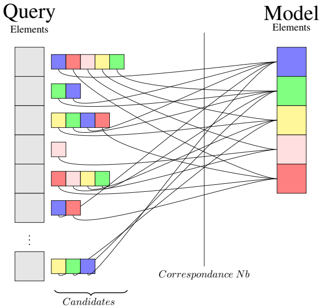

- Structure
  - list of candidates
  - Associated scores
  - time more important than memory

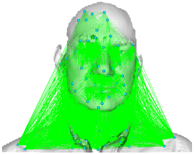

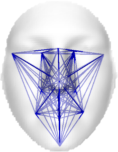

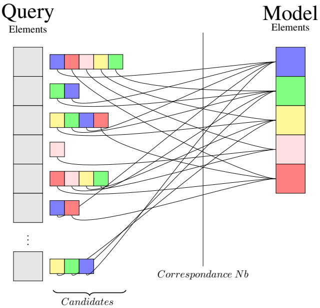

- Structure
  - list of candidates
  - Associated scores
  - time more important than memory
- Objective:
  - Reduce correspondence Nb

Structural Matching

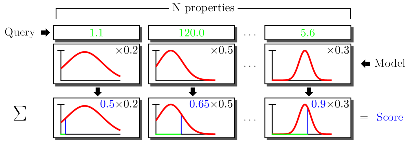

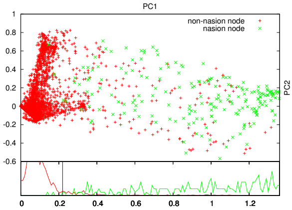

LDA

- Structure
  - list of candidates
  - Associated scores
  - time more important than memory
- Objective:
  - Reduce correspondence Nb
- Seeding
  - Partial scores $\xrightarrow{LDA}$ Score

- Correpsondance Feature Space:
  - Local Decisions:
    - Relaxation on hypergraph ($\neq$ [Christmas et al., 1995])
  - Global Decisions:
    - Convex Optimization [Zass and Shashua, 2008]
    - Tensor power iteration [Duchenne et al., 2009]
    - Randomly Sparsified Spectral method [Chertok and Keller, 2010]
- Correspondance Similarity Space: (post-processing only)
  - Decisions by Clustering (for hyperedges of degree 3):
    - Unit-quaternion clustering technique
    - RANSAC on model registration errors

Structural Matching - approaches

- Correpsondance Feature Space:
  - Local Decisions:
    - Relaxation on hypergraph ($\neq$ [Christmas et al., 1995]) **new**
  - Global Decisions:
    - Convex Optimization [Zass and Shashua, 2008] **modified**
    - Tensor power iteration [Duchenne et al., 2009] **used**
    - Randomly Sparsified Spectral method [Chertok and Keller, 2010] **not used**
- Correspondance Similarity Space: (post-processing only)
  - Decisions by Clustering (for hyperedges of degree 3):
    - Unit-quaternion clustering technique **new**
    - RANSAC on model registration errors **new**

Correspondance through clustering

- It is sometime possible to see the graph matching problem as a clustering problem:

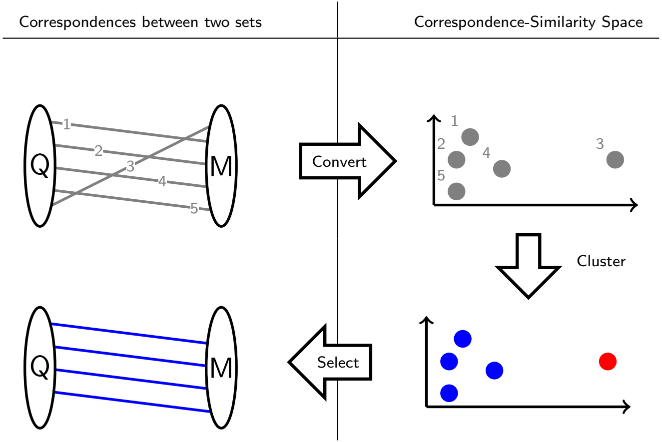

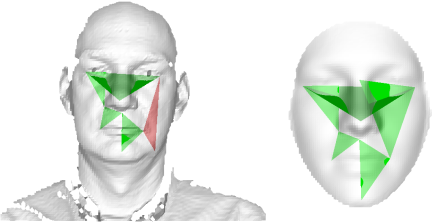

Transformation Matrix 4x4:

$$\left(\begin{array}{c|c} R' & \vec{t} \\ \hline 0 & 1 \end{array}\right) \rightarrow \begin{array}{ll} \dot{q} & \text{Unit Quaternion} \\ \vec{t} & \text{Translation} \\ s & \text{Scale} \end{array}$$

Post-Processing

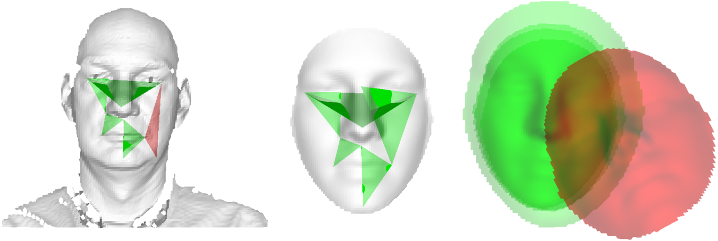

Unit-Quaternion clustering

- Similarity: Angle between quaternions
- Clustering + Mean Rotation
- Final Correspondence (NN)

RANSAC Selection

- Similarity: Mean distance between registered landmarks
- RANSAC Selection

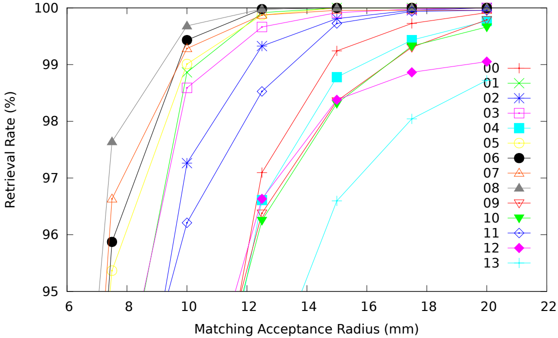

Figure: Landmark retrieval rate for the 14 landmarks on the FRGC test set.

Results

Authors	[Chang et al., 2006]	[Mian et al., 2006]	[Segundo et al., 2007]	[Romero and Pears, 2009]	[Alyuz et al., 2010]			[Segundo et al., 2010]		This Work, 2011			
#Landmarks	3	1	6	3	5			5		14			
Acceptance Radius	$<?$	$<?$	$<?$	<12	<10	<12	<20	<10	<15	<10	<12	<15	<20
Nose (05)	99.40	98.3	99.95	99.77	99.62	99.80	99.87	**99.95**	99.95	99.01	**99.81**	**100.0**	**100.0**
Eye Inner Corners (01,03)	–	–	99.83	96.82	96.59	98.54	99.54	**99.02**	99.64	98.73	**99.71**	**99.96**	**100.0**
Nose Corners (06,07)	–	–	99.76	–	98.60	99.29	99.87	99.35	99.95	**99.36**	**99.87**	**99.98**	**99.98**
Subnasale (08)	–	–	99.98	–	–	–	–	–	–	99.68	99.98	100.0	100.0
Mouth Corners (09,10)	–	–	–	–	–	–	–	–	–	91.33	95.63	98.34	99.73
Eye Outer Corners (00,04)	–	–	–	–	–	–	–	–	–	89.84	95.92	99.01	99.84
Nasion (02)	–	–	–	–	–	–	–	–	–	97.26	99.07	99.81	100.0
Upper Lip (11)	–	–	–	–	–	–	–	–	–	96.21	98.21	99.73	99.96
Lower Lip (12)	–	–	–	–	–	–	–	–	–	92.04	96.00	98.38	99.05
Chin (13)	–	–	–	–	–	–	–	–	–	84.94	91.96	96.60	98.72
Candidate Selection	ES	ES	ES	ES	ES			ES		ML			
Independence	no	n/a	no	yes	no			no		yes			
Test Size	4,485	4,950	4,007	4,013	4,007			4,007		4,750			
Train Size	–	–	–	–	–			–		200			
Pre-processing	S,C[1]	∅	H,C	S,H	S,H,C			S,H,C		∅			
Pre-processing Time	–	–	1.1s	–	–			1.0s		0s			
Processing Time	–	–	0.4s	–	–			0.3s		1.18s			

ES: Expert System, ML: Machine Learning, C: Cropped/Segmented, H: Hole Filling, S: Spike Removal

Examples

Examples

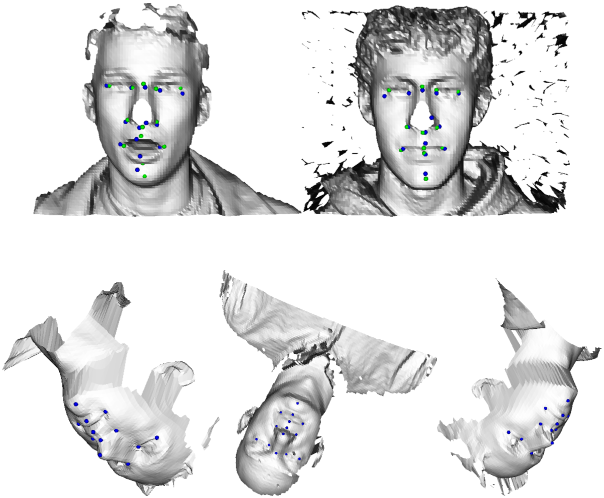

Examples

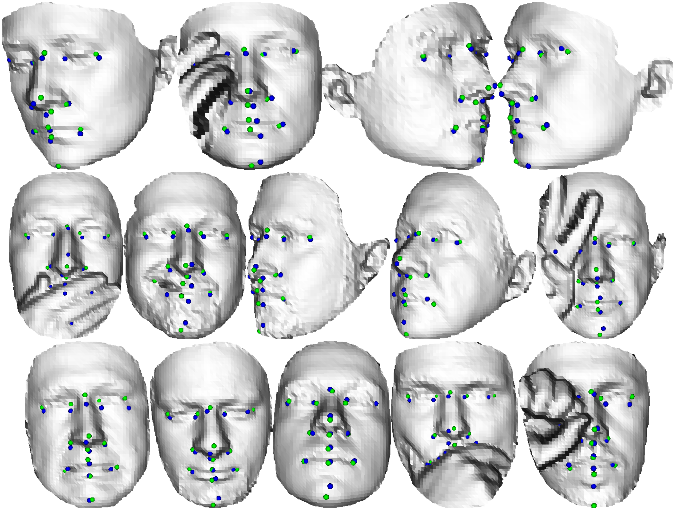

Model Generation

Where is Wally?
Waldo?
Charlie?
Walter?
ウォーリー?
威利?
...

Scene

Generalities

Where is Wally?
Waldo?
Charlie?
Walter?
ウォーリー?
威利?
. . .

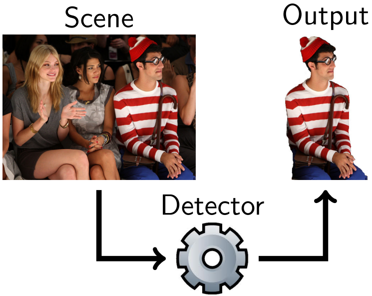

Generalities

Where is Wally?
Waldo?
Charlie?
Walter?
ウォーリー?
威利?
...

Scene

Output

Detector

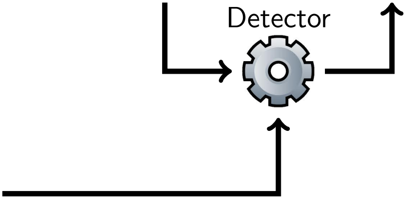

Model

Generalities

Where is Wally?
Waldo?
Charlie?
Walter?
ウォーリー?
威利?
...

Model

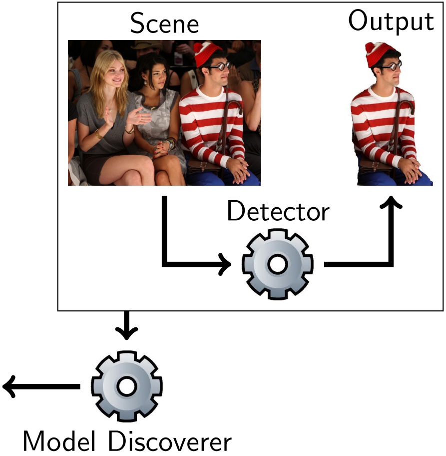

Generalities

Where is Wally?
Waldo?
Charlie?
Walter?
ウォーリー?
威利?
. . .

Model

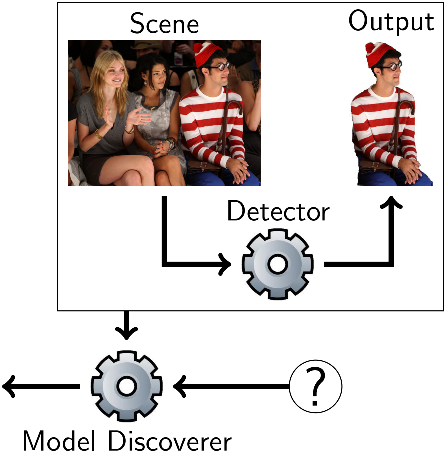

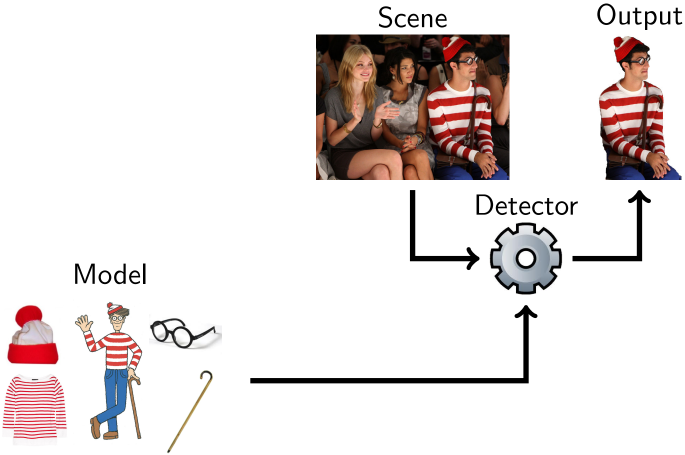
Scene
Output
Detector
Model

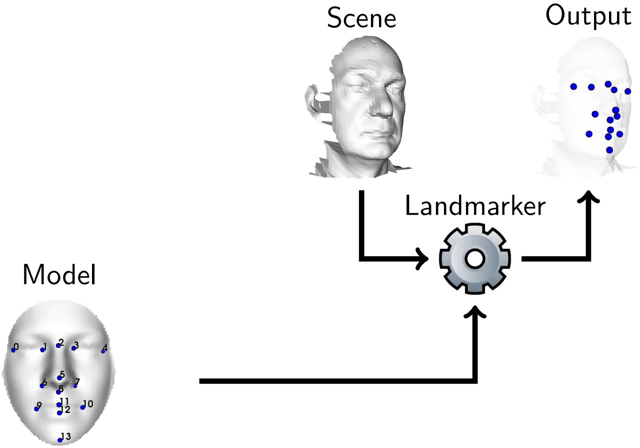

Scene
Output
Landmarker
Model

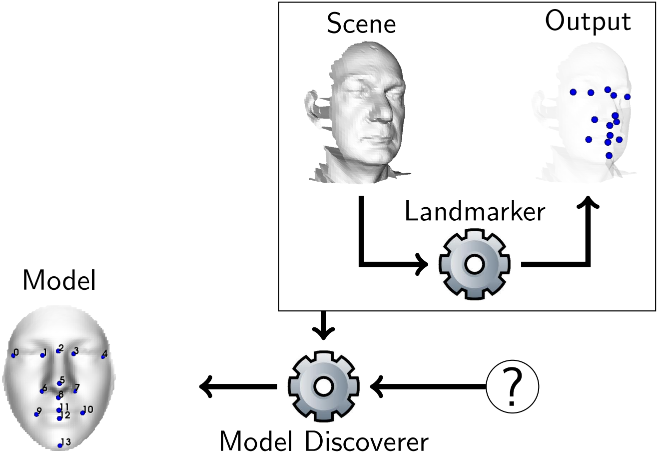
Scene
Output
Landmarker
Model
0
1
2
3
4
5
6
7
8
9
10
11
12
13
?
Model Discoverer

Why? - Gap in Research

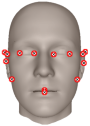

[Amberg et al., 2007]

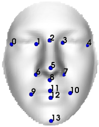

[Creusot et al., 2011]

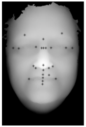

[Gupta et al., 2007]

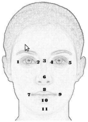

[Romero-Huertas and Pears, 2008]

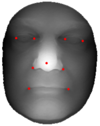

[Szeptycki et al., 2009]

[Zhao et al., 2011]

- Easy to label or explain to an operator
- Linked to 2D projections and plane symmetries
- Overall arbitrary

- Sparse
- "Descriptive"

- Featural/Local information (nodes)
- Structural/Global information (edges/hyperedges)

Possible Local Features:

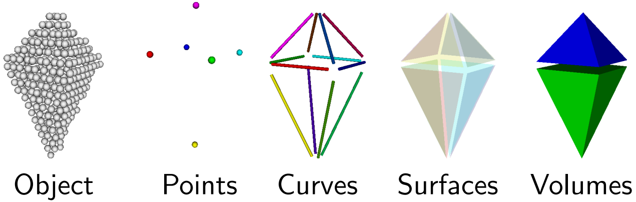

- Sparse
- "Descriptive"

- Featural/Local information (nodes)
- Structural/Global information (edges/hyperedges)

Possible Local Features:

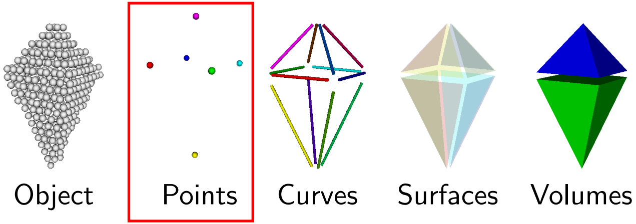

- More possible point-models than geometric shapes
- Less intuition about what model is good

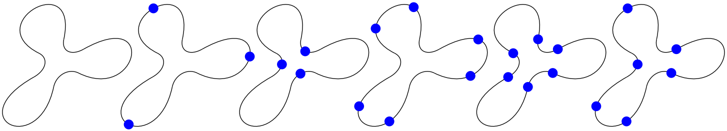

Example of 3D-objects point models

Articulated Models:

- Articulations
- Extremities

Non-Articulated Models:

- ???
- ???

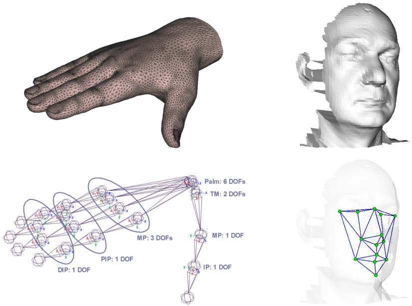

[Shotton et al., 2011] [Bray et al., 2004] [Creusot et al., 2011]

Hypothesis

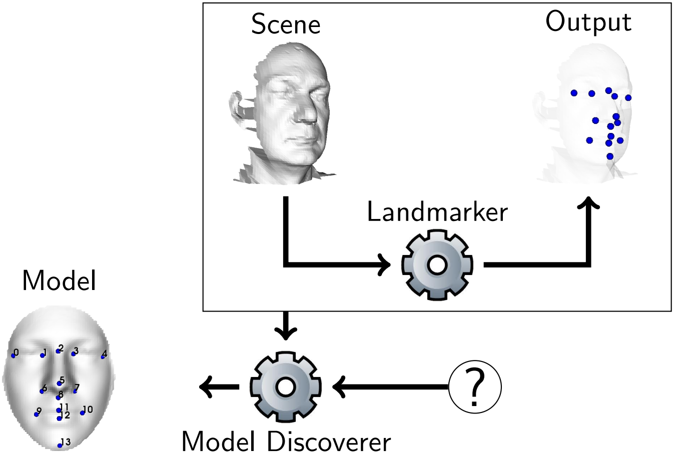

Scene
Output
Landmarker
Model
0
1
2
3
4
5
6
7
8
9
10
11
12
13
?
Model Discoverer

Hypothesis

- “Probabilistic” response map available
- One point per model

Model

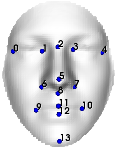

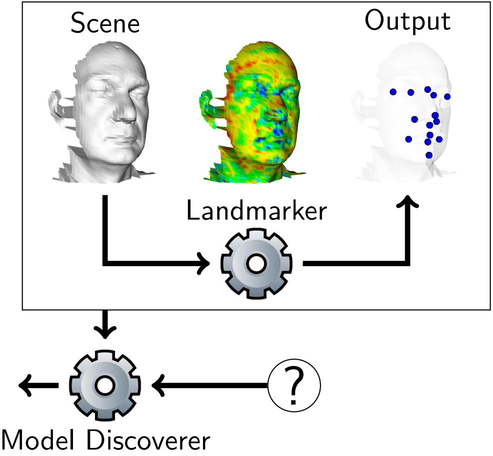

Our Approach

- Use Detector and Neighborhood definition from [Creusot et al., 2011]
  - 8 Local Descriptors
  - Gaussian Distributions
  - Linear Combination (LDA based)

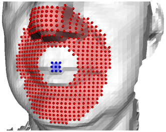

- Test as many models as there are vertices in the template mesh (~ 2000)
- Define two cost functions for each model:
  - **Saliency**: Different from its neighborhood (good)
  - **Ubiquity**: Ubiquitous over the face (bad)

Databases

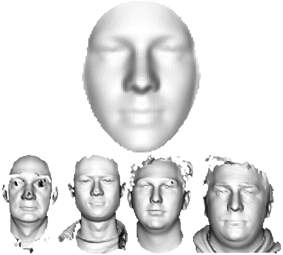

FRGC (real)
(Coarse Correspondence)

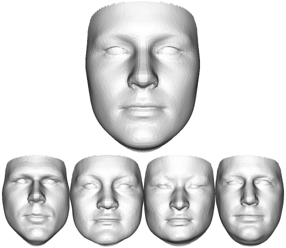

BFM (synthetic)
(Fine Correspondence)

Saliency Score per Vertex

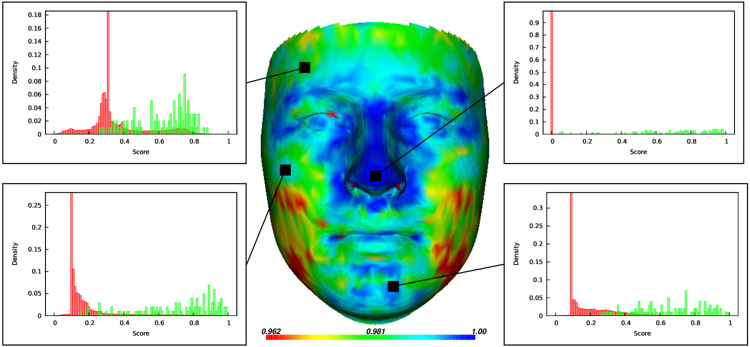

Ubiquity Score per Vertex

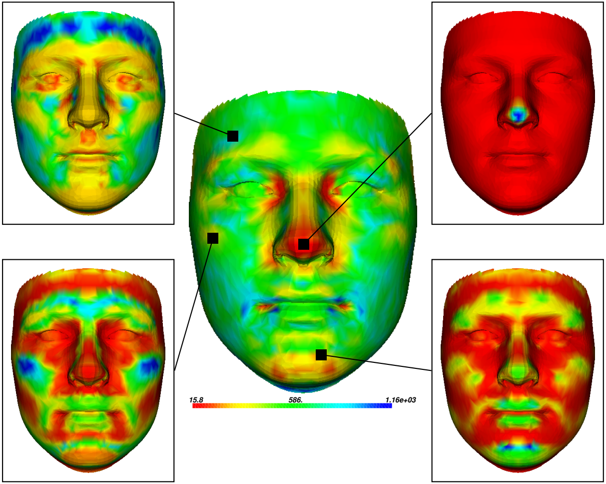

Results

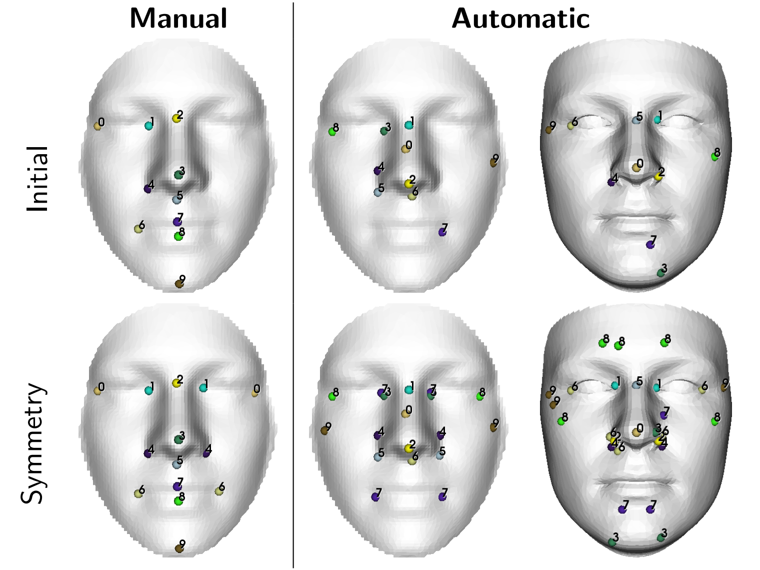

Saliency

	Similar						
ID	8	3	1	2	5	6	7
U	658.37	215.52	287.70	67.452	252.50	393.68	324.40
Automatic	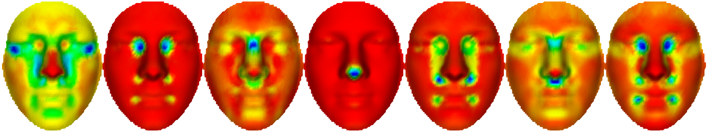						
ID	0	1	2	3	4	5	6
U	683.62	192.43	492.43	87.408	328.08	226.66	206.76
Manual	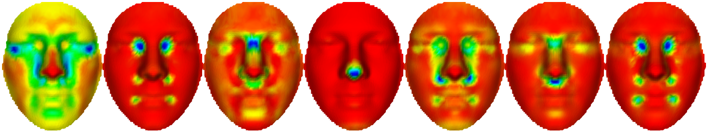						

Saliency

	Different		
	0	4	9
U	325.51	561.31	416.18
Automatic	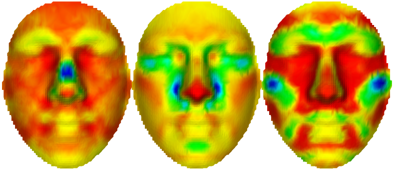		
ID	7	8	9
U	701.64	954.16	444.54
Manual	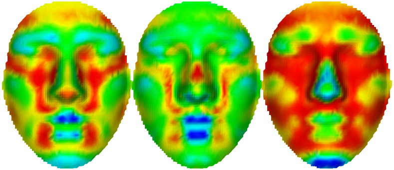		

Problems

- Different answers depending on the registration method:
  - Fine registration on clean data (BFM)
  - Coarse registration on unclean data (FRGC)
  - Fine registration on unclean data (???) Needed
- Optimization method → Depends on the detector used (and its parameters)
- How to include structural information in the model discovery?
- How to project a newly discovered model to unseen training data? (again a registration problem)

Conclusion

- Good:
  - Optimize model for a detector
  - Validate most human-chosen landmarks
  - Give quantifiable measure of landmark quality
- Bad:
  - Only non-articulated objects for now
  - Requires a large set of finely-registered objects
- Brain teaser:
  - How do you extend the idea to multi-dimensional features (curves, area, volumes)?

Conclusion

Main Contributions

- Methodological contributions (New approach):
  - Full 3D (pose invariant) Machine Learning based face landmarking.
  - No domain specific constrains assumed.
  - Independant/parallel search of numerous landmarks.
- Practical contributions:
  - A new framework for keypoint detection (2 methods tested)
  - A study of numerous simple descriptors at different scales on faces
  - A framework for feature labelling using hypergraph matching filters (2 methods tested)
  - A new graph matcher by relaxation alterning between the hypergraph and its dual

Characteristic	Expert Systems Face Landmarkers	This Work
Object type	3D face only	Non-articulated objects
Landmarks number	Fix (often <5)	Arbitrary (tested up to 14)
Individual detection	Landmark dependent	Landmark-independent
Processing order	Sequential	Concurrent
Detections	Map extrema	Score map extrema
Landmark-Map correlations	Manually provided by researcher	Learnt automatically
Pre-processing needed	Yes	No
Local descriptors type	Scalar only	Scalar and histogram
Descriptors number	Fixed (<2)	Arbitrary (tested up to 40)
Descriptors combination	Manually fixed (linear)	Learnt (linear or non-linear)

Limitations

- Local 3D shape descriptors suffer for spurious data
- Local 3D shape descriptors suffer for occlusions (near profiles)
- 3D descriptors are computationally expensive

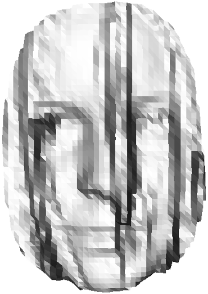

- Interesting challenges:
  - Keypoint detectors for mesh border points (new descriptors).
  - Extend to higher-dimensional features
  - Explore hypergraph representation for non-cooperative face recognition

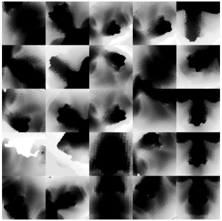

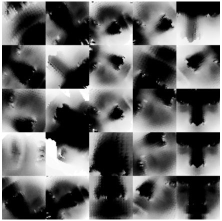

Thank You For Listening!

http://www.cs.york.ac.uk/~creusot

References I

Alyuz, N., Gokberk, B., and Akarun, L. (2010).
Regional registration for expression resistant 3-d face recognition.
Information Forensics and Security, IEEE Transactions on, 5(3):425–440.

Amberg, B., Romdhani, S., and Vetter, T. (2007).
Optimal step nonrigid icp algorithms for surface registration.
In *Computer Vision and Pattern Recognition, 2007. CVPR '07. IEEE Conference on*, pages 1 –8.

Berretti, S., Bimbo, A. D., and Pala., P. (2010).
Recognition of 3d faces with missing parts based on profile networks.
In *1st ACM Workshop on 3D Object Retrieval (ACM 3DOR'10)*, pages 81–86, Firenze, Italy.

Bray, M., Koller-Meier, E., Mueller, P., Gool, L. V., and Schraudolph, N. N. (2004).
3d hand tracking by rapid stochastic gradient descent using a skinning model.
In Chambers, A. and Hilton, A., editors, *1st European Conference on Visual Media Production (CVMP)*, pages 59–68. IEE.

Chang, K. I., Bowyer, K., and Flynn, P. (2006).
Multiple nose region matching for 3d face recognition under varying facial expression.
Pattern Analysis and Machine Intelligence, IEEE Transactions on, 28(10):1695–1700.

Chertok, M. and Keller, Y. (2010).
Efficient high order matching.
Pattern Analysis and Machine Intelligence, IEEE Transactions on, 32(12):2205–2215.

Christmas, W. J., Kittler, J., and Petrou, M. (1995).
Structural matching in computer vision using probabilistic relaxation.
IEEE Trans. Pattern Anal. Mach. Intell., 17(8):749–764.

Colbry, D., Stockman, G., and Jain, A. (2005).
Detection of anchor points for 3D face veri.cation.
In *Computer Vision and Pattern Recognition - Workshops, 2005. CVPR Workshops. IEEE Computer Society Conference on*, pages 118–118.

Creusot, C., Pears, N., and Austin, J. (2011).
Automatic keypoint detection on 3D faces using a dictionary of local shapes.
In *3D Imaging, Modeling, Processing, Visualization and Transmission (3DIMPVT), 2011 International Conference on*, pages 204–211.

D'Hose, J., Colineau, J., Bichon, C., and Dorizzi, B. (2007).
Precise localization of landmarks on 3D faces using Gabor wavelets.
Biometrics: Theory, Applications, and Systems, 2007. BTAS 2007. First IEEE International Conference on, pages 1–6.

Duchenne, O., Bach, F., Kweon, I., and Ponce, J. (2009).
A tensor-based algorithm for high-order graph matching.
In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, pages 1980 –1987.

Faltemier, T., Bowyer, K., and Flynn, P. (2008a).
A region ensemble for 3-D face recognition.
Information Forensics and Security, IEEE Transactions on, 3(1):62–73.

References III

Faltemier, T., Bowyer, K., and Flynn, P. (2008b).
Rotated profile signatures for robust 3D feature detection.
In *Automatic Face Gesture Recognition, 2008. FG '08. 8th IEEE International Conference on*, pages 1–7.

Gupta, S., Markey, M. K., Aggarwal, J., and Bovik, A. C. (2007).
Three dimensional face recognition based on geodesic and euclidean distances.
In *IS&T/SPIE Symposium on Electronic Imaging: Vision Geometry XV, Proceedings of the SPIE*.

Havasi, L., Szlavik, Z., and Sziranyi, T. (2007).
Detection of gait characteristics for scene registration in video surveillance system.
Image Processing, IEEE Transactions on, 16(2):503–510.

Liu, C.-C., Dai, D.-Q., and Yan, H. (2007).
Local discriminant wavelet packet coordinates for face recognition.
J. Mach. Learn. Res., 8:1165–1195.

Lu, X. and Jain, A. (2006).
Automatic feature extraction for multiview 3D face recognition.
In *Automatic Face and Gesture Recognition, 2006. FGR 2006. 7th International Conference on*, pages 585–590.

Mian, A. S., Bennamoun, M., and Owens, R. A. (2006).
Automatic 3D face detection, normalization and recognition.
In *3DPVT*, pages 735–742.

PEARS, N., HESELTINE, T., AND ROMERO, M. (2010).
FROM 3D POINT CLOUDS TO POSE-NORMALISED DEPTH MAPS.
International Journal of Computer Vision, 89(2):152–176.

PHILLIPS, P., BEVERIDGE, J., DRAPER, B., GIVENS, G., O'TOOLE, A., BOLME, D., DUNLOP, J., LUI, Y. M., SAHIBZADA, H., AND WEIMER, S. (2011).
AN INTRODUCTION TO THE GOOD, THE BAD, AMP; THE UGLY FACE RECOGNITION CHALLENGE PROBLEM.
IN *Automatic Face Gesture Recognition and Workshops (FG 2011), 2011 IEEE International Conference on*, PAGES 346–353.

PHILLIPS, P., FLYNN, P., SCRUGGS, T., BOWYER, K., CHANG, J., HOFFMAN, K., MARQUES, J., MIN, J., AND WOREK, W. (2005).
OVERVIEW OF THE FACE RECOGNITION GRAND CHALLENGE.
Computer Vision and Pattern Recognition, 2005. CVPR 2005. IEEE Computer Society Conference on, 1:947–954.

PROENCA, H. (2008).
IRIS RECOGNITION: A METHOD TO SEGMENT VISIBLE WAVELENGTH IRIS IMAGES ACQUIRED ON-THE-MOVE AND AT-A-DISTANCE.
Advances in Visual Computing, PAGES 731–742.

ROMERO, M. AND PEARS, N. (2009).
LANDMARK LOCALISATION IN 3D FACE DATA.
IN *Advanced Video and Signal Based Surveillance, 2009. AVSS '09. Sixth IEEE International Conference on*, PAGES 73–78.

References V

Romero-Huertas, M. and Pears, N. (2008).
3D facial landmark localisation by matching simple descriptors.
In *Biometrics: Theory, Applications and Systems, 2008. BTAS 2008. 2nd IEEE International Conference on*, pages 1–6.

Savran, A., Alyüz, N., Dibeklioğlu, H., Çeliktutan, O., Gökberk, B., Sankur, B., and Akarun, L. (2008).
Bosphorus database for 3D face analysis.
In *Biometrics and Identity Management: First European Workshop, BIOID 2008, Roskilde, Denmark, May 7-9, 2008.*, pages 47–56, Berlin, Heidelberg. Springer-Verlag.

Segundo, M., Queirolo, C., Bellon, O., and Silva, L. (2007).
Automatic 3D facial segmentation and landmark detection.
Image Analysis and Processing, 2007. ICIAP 2007. 14th International Conference on, pages 431–436.

Segundo, M., Silva, L., Bellon, O. R. P., and Queirolo, C. C. (2010).
Automatic face segmentation and facial landmark detection in range images.
Systems, Man, and Cybernetics, Part B: Cybernetics, IEEE Transactions on, 40(5):1319–1330.

Shotton, J., Fitzgibbon, A., Cook, M., Sharp, T., Finocchio, M., Moore, R., Kipman, A., and Blake, A. (2011).
Real-time human pose recognition in parts from single depth images.
In *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*, pages 1297 –1304.

References VI

Szeptycki, P., Ardabilian, M., and Chen, L. (2009).
A coarse-to-fine curvature analysis-based rotation invariant 3D face landmarking.
In *International Conference on Biometrics: Theory, Applications and Systems*, pages 32–37.

Yan, P. and Bowyer, K. (2007).
Biometric recognition using 3d ear shape.
Pattern Analysis and Machine Intelligence, IEEE Transactions on, 29(8):1297–1308.

Zass, R. and Shashua, A. (2008).
Probabilistic graph and hypergraph matching.
In *Computer Vision and Pattern Recognition, 2008. CVPR 2008. IEEE Conference on*, pages 1–8.

Zhao, X., Dellandré anda, E., Chen, L., and Kakadiaris, I. A. (2011).
Accurate landmarking of three-dimensional facial data in the presence of facial expressions and occlusions using a three-dimensional statistical facial feature model.
Systems, Man, and Cybernetics, Part B: Cybernetics, IEEE Transactions on, 41(5):1417–1428.