3D Landmark Model Discovery from a Registered Set of Organic Shapes

Clement Creusot, Nick Pears, Jim Austin

Department of Computer science
The University of York

PCP, CVPRW, Providence (RI), June 2012

Plan

- What/Why: Generalities and Problems
- How: Proposed method
- Results
- Conclusion

Generalities

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?
...

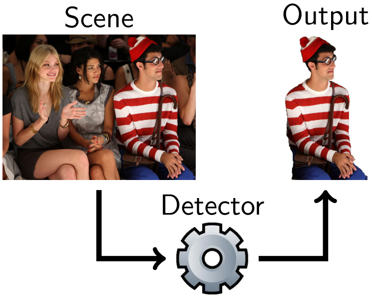

Generalities

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

Model

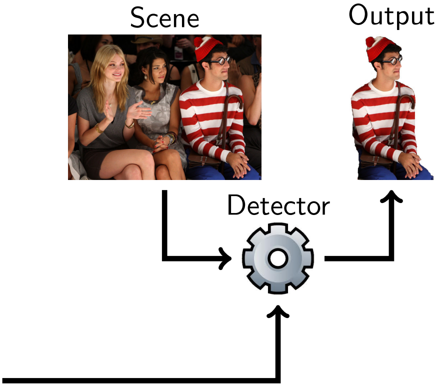

Generalities

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

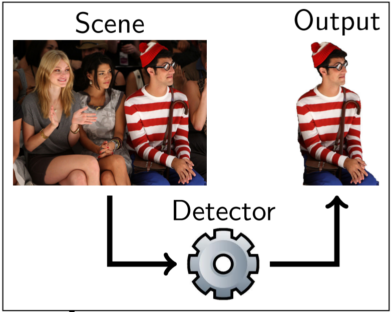

Model

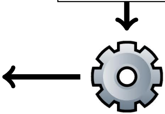

Model Discoverer

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

Model

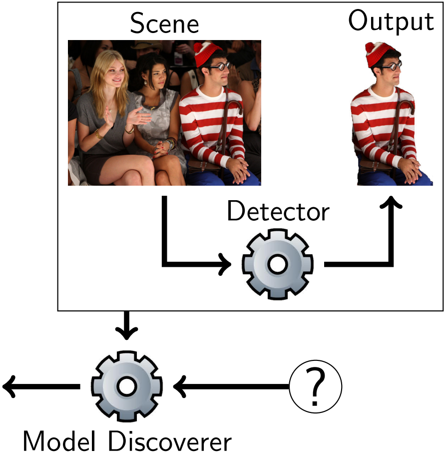

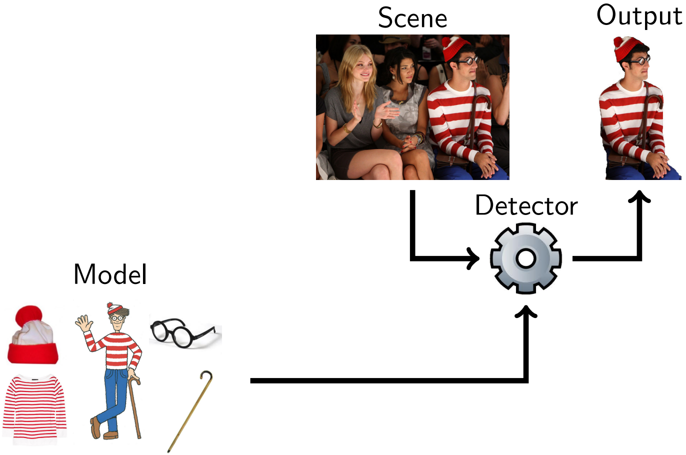
Scene
Output
Detector
Model

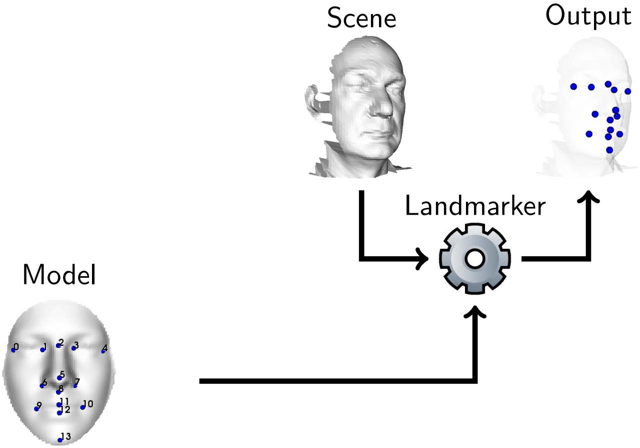

Scene
Output
Landmarker
Model
0
1
2
3
4
5
6
7
8
9
10
11
12
13

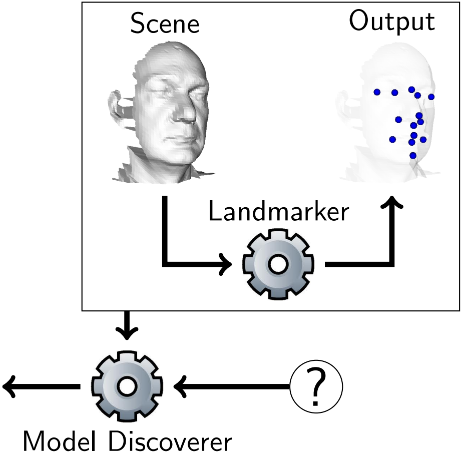
Scene
Output
Landmarker
Model
0
1
2
3
4
5
6
7
8
9
10
11
12
13
?
Model Discoverer

Why? - Gap in Research

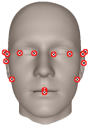

[Amberg et al., 2007]

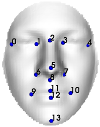

[Creusot et al., 2011]

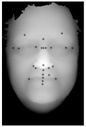

[Gupta et al., 2007]

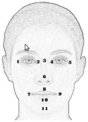

[Romero-Huertas and Pears, 2008]

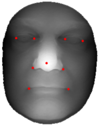

[Szeptycki et al., 2009]

[Zhao et al., 2011]

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

Nature of a model for a 3D-object class

- Sparse
- "Descriptive"

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

Possible Local Features:

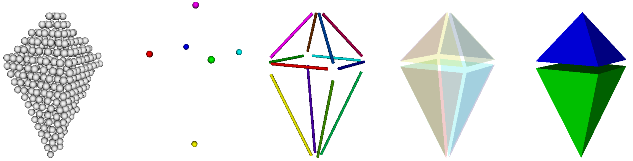

Nature of a model for a 3D-object class

- Sparse
- "Descriptive"

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

Possible Local Features:

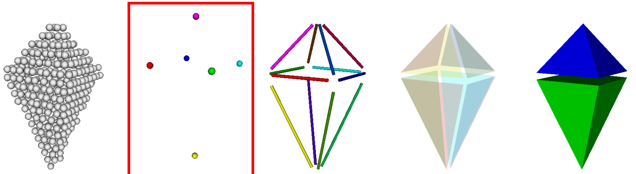

Object Points Curves Surfaces Volumes

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

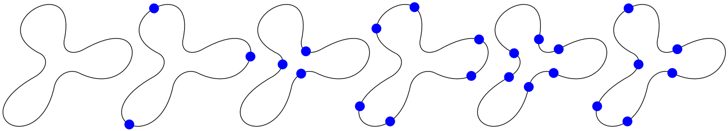

Example of 3D-objects point models

Articulated Models:

- Articulations
- Extremities

Non-Articulated Models:

- ???
- ???

[Shotton et al., 2011]

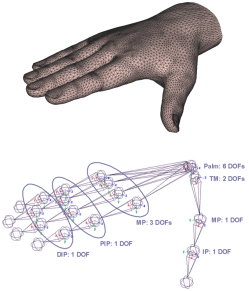

[Bray et al., 2004]

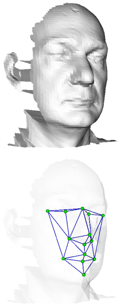

[Creusot et al., 2011]

Hypothesis

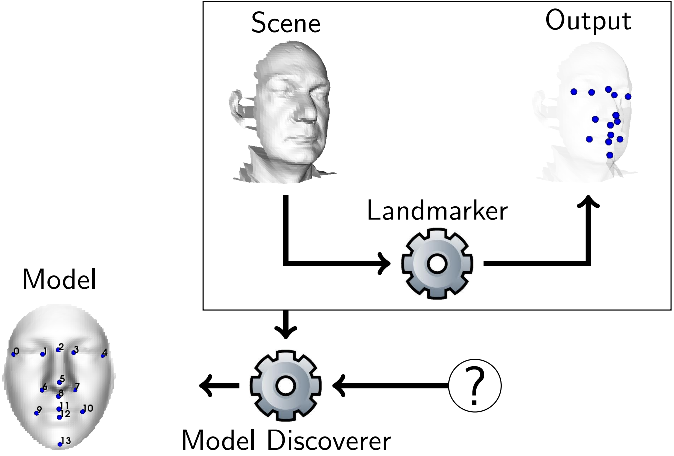
Scene
Output
Landmarker
Model
0
1
2
3
4
5
6
7
8
9
10
11
12
13
Model Discoverer
?

Hypothesis

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

Model

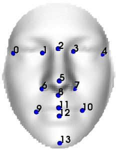

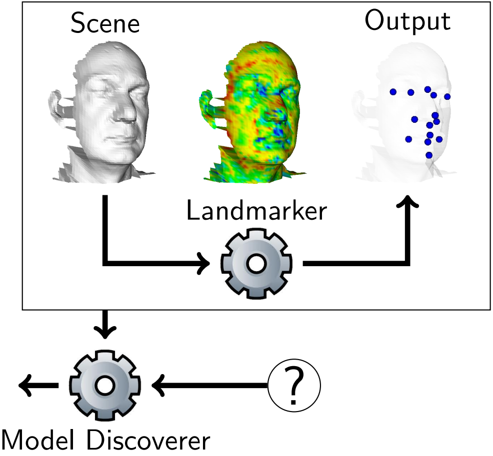

Our Approach

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

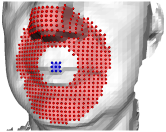

- Test as many models as there are vertices in the template mesh (~ 2000)
- Define two cost functions for each model:
  - **Saliency**: Different from its neighborhood (good)
  - **Ubiquity**: Ubiquitous over the face (bad)

Databases

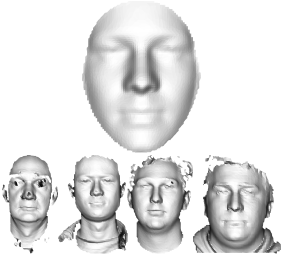

FRGC (real)
(Coarse Correspondence)

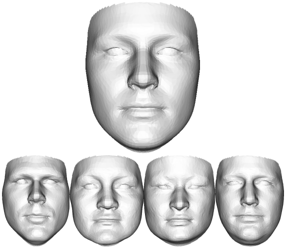

BFM (synthetic)
(Fine Correspondence)

Saliency Score per Vertex

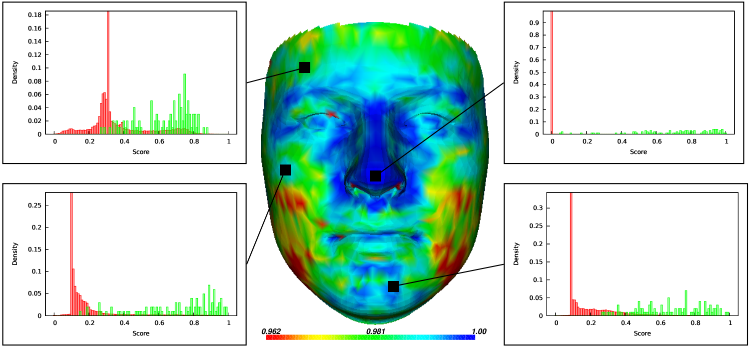

Ubiquity Score per Vertex

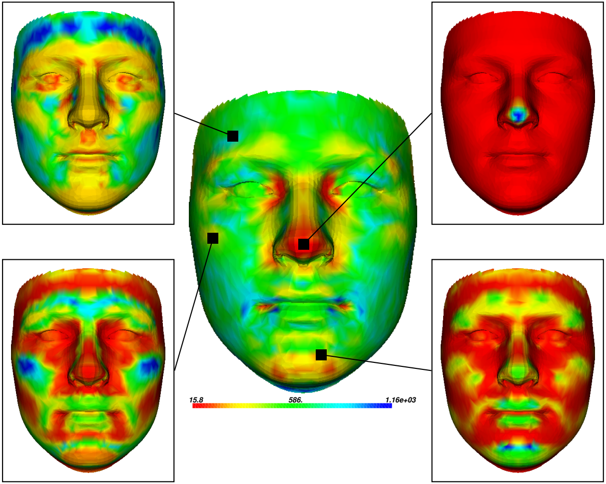

Results

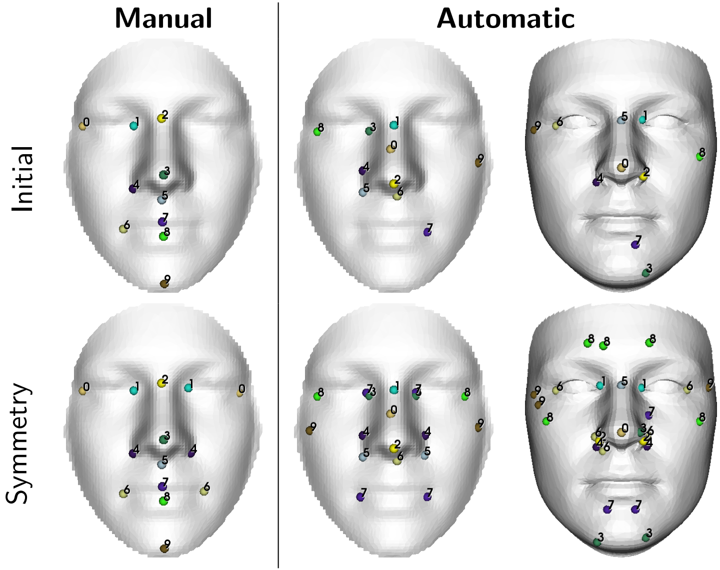

Saliency

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

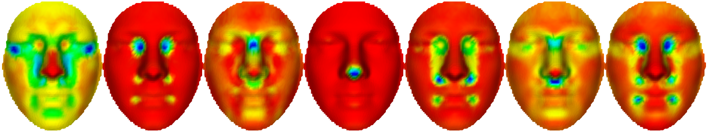

ID	0	1	2	3	4	5	6
U	683.62	192.43	492.43	87.408	328.08	226.66	206.76

Manual

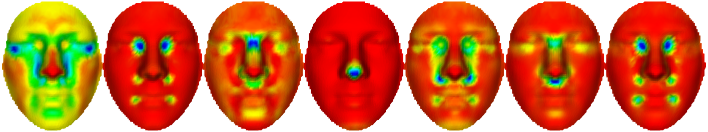

Saliency

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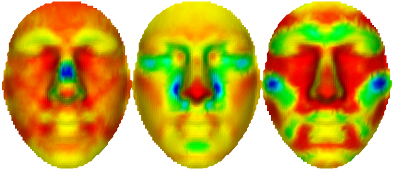		
ID	7	8	9
U	701.64	954.16	444.54
Manual	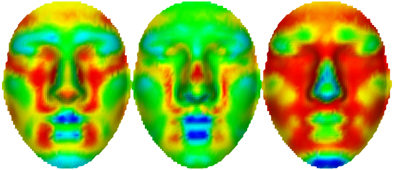		

Problems

- 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 $\rightarrow$ 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 (Do you have one to share?)
- Questions to you:
  - How do you learn a model structure in your application domain?
  - Are there applications where you think this might help?
  - Brain teaser: How do you extend the idea to multi-dimensional features (curves, area, volumes)?

References I

Amberg, B., Romdhani, S., and Vetter, T. (2007).
Optimal step nonrigid icp algorithms for surface registration.
In *IEEE Int Conf. CVPR*.

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.

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

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 Symp. on Electronic Imaging: Vision Geometry XV*.

Romero-Huertas, M. and Pears, N. (2008).
3D facial landmark localisation by matching simple descriptors.
In *IEEE Int. Conf. BTAS*, pages 1–6.

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 II

Szeptycki, P., Ardabilian, M., and Chen, L. (2009).
A coarse-to-fine curvature analysis-based rotation invariant 3D face landmarking.
In *IEEE Int. Conf. BTAS*, pages 32–37.

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.
IEEE Trans. Syst. Man, and Cybernetics, Part B: Cybernetics, 41(5):1417–1428.