

Contour People:

A Parameterized Model of 2D Articulated Human Shape




O. Freifeld S. Zuffi A. Weiss M.J. Black

Brown University




2D generative models of humans

- In wide use in Computer Vision
Pose estimation; tracking; gesture analysis
- Go back a long way
[Fischler & Elschlager 73, Hinton 76, Hogg 76, Ju et al. 96]
- Computationally efficient
Pictorial Structure (PS) and Belief Propagation (BP)
[Felzenszwalb & Huttenlocher, IJCV '05]
[Andriluka et al. CVPR '09]




Problem

- Lack of realism (no **detailed shape**)
- Body shape estimation has many applications
Gaming, clothing industry, security
- A good model of shape can improve pose estimation




Possible solution: SCAPE

- A 3D **graphics** model
[Anguelov et al. Siggraph '05]
- Used in **computer vision** for shape and pose estimation from multiple calibrated cameras
[Balan et al. '07]




Problem

- Inference is computationally **expensive**
- Assumes **calibrated cameras**
- Ambiguous in a **single view**



[Guan et al. ICCV '09]

Goal: The best of both



2D

parameters: 12-40




3D

60-100

Goal: The best of both

- Learn an articulated 2D model of detailed shape that is lower-dimensional than SCAPE



2D

parameters: 12-40



25-60




3D

60-100

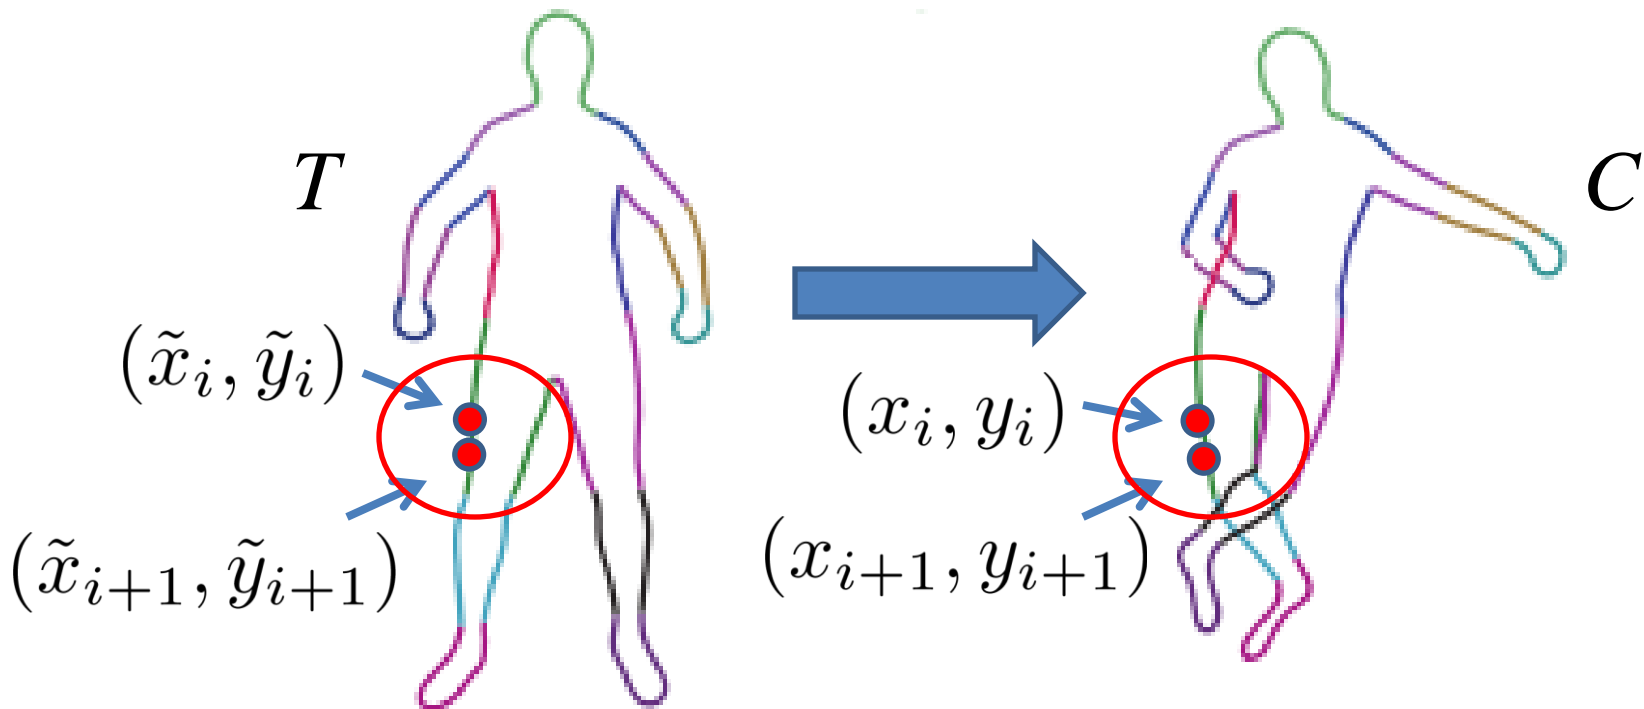
The Contour Person (CP) model

- A part-based articulated **deformable template** model
- It **factorizes** shape, pose and camera variations
- A 2D model that captures the **detailed shape** of the human body
- The model is learned from 3D SCAPE




Deformable template

$$T = (\tilde{x}_1 \quad \tilde{y}_1 \quad \dots \quad \tilde{x}_n \quad \tilde{y}_n)^T \quad C = (x_1 \quad y_1 \quad \dots \quad x_n \quad y_n)^T$$




Deformation of a line segment

$$T = (\tilde{x}_1 \quad \tilde{y}_1 \quad \dots \quad \tilde{x}_n \quad \tilde{y}_n)^T \quad C = (x_1 \quad y_1 \quad \dots \quad x_n \quad y_n)^T$$



Deformation of a line segment

$$T = (\tilde{x}_1 \ \tilde{y}_1 \ \dots \ \tilde{x}_n \ \tilde{y}_n)^T \quad C = (x_1 \ y_1 \ \dots \ x_n \ y_n)^T$$



S_i is the ratio of the lengths

$$D_i^{2 \times 2} = S_i \begin{pmatrix} \cos \theta_i & -\sin \theta_i \\ \sin \theta_i & \cos \theta_i \end{pmatrix}$$

Template contour $T = \left(\tilde{x}_1 \quad \tilde{y}_1 \quad \dots \quad \tilde{x}_n \quad \tilde{y}_n \right)^T$

New contour $C = \left(x_1 \quad y_1 \quad \dots \quad x_n \quad y_n \right)^T$

Connectivity matrix $E = \begin{pmatrix} -1 & 0 & +1 & 0 & \dots & 0 & 0 & 0 \\ 0 & -1 & 0 & +1 & \dots & 0 & 0 & 0 \\ 0 & 0 & -1 & 0 & \dots & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & \dots & -1 & 0 & +1 \\ +1 & 0 & 0 & 0 & \dots & 0 & -1 & 0 \\ 0 & +1 & 0 & 0 & \dots & 0 & 0 & -1 \end{pmatrix}$

Template contour $T = \left(\tilde{x}_1 \quad \tilde{y}_1 \quad \dots \quad \tilde{x}_n \quad \tilde{y}_n \right)^T$

New contour $C = \left(x_1 \quad y_1 \quad \dots \quad x_n \quad y_n \right)^T$

Connectivity matrix $E = \begin{pmatrix} -1 & 0 & +1 & 0 & \dots & 0 & 0 & 0 \\ 0 & -1 & 0 & +1 & \dots & 0 & 0 & 0 \\ 0 & 0 & -1 & 0 & \dots & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & \dots & -1 & 0 & +1 \\ +1 & 0 & 0 & 0 & \dots & 0 & -1 & 0 \\ 0 & +1 & 0 & 0 & \dots & 0 & 0 & -1 \end{pmatrix}$



Template line segments $E T = \left(\tilde{x}_2 - \tilde{x}_1 \quad \tilde{y}_2 - \tilde{y}_1 \quad \dots \quad \tilde{x}_1 - \tilde{x}_n \quad \tilde{y}_1 - \tilde{y}_n \right)^T$

New line segments $E C = \left(x_2 - x_1 \quad y_2 - y_1 \quad \dots \quad x_1 - x_n \quad y_1 - y_n \right)^T$


Deformation of T

$$D(\Theta) = \begin{pmatrix} D_1^{2 \times 2} & 0 & \dots & 0 \\ 0 & D_2^{2 \times 2} & \dots & 0 \\ 0 & 0 & \ddots & 0 \\ 0 & 0 & \dots & D_n^{2 \times 2} \end{pmatrix}$$

$$D_i^{2 \times 2} = S_i \begin{pmatrix} \cos \theta_i & -\sin \theta_i \\ \sin \theta_i & \cos \theta_i \end{pmatrix}$$

Deformation of T


$$D(\Theta) = \begin{pmatrix} D_1^{2 \times 2} & 0 & \dots & 0 \\ 0 & D_2^{2 \times 2} & \dots & 0 \\ 0 & 0 & \ddots & 0 \\ 0 & 0 & \dots & D_n^{2 \times 2} \end{pmatrix}$$



$$E C = D(\Theta) E T$$


Deformation of T

$$D(\Theta) = \begin{pmatrix} D_1^{2 \times 2} & 0 & \dots & 0 \\ 0 & D_2^{2 \times 2} & \dots & 0 \\ 0 & 0 & \ddots & 0 \\ 0 & 0 & \dots & D_n^{2 \times 2} \end{pmatrix}$$




$$E C \approx D(\Theta) E T$$



T





T



ET





$D(\Theta)ET$





$$\min_C \|D(\Theta)ET - EC\|^2$$

$$D(\Theta)ET$$



$$\min_C \|D(\Theta)ET - EC\|^2$$




$$C = E^\dagger D(\Theta)ET$$



Factorization

$$\begin{aligned}\Theta &= (\Theta_{\text{shape}}, \Theta_{\text{pose}}, \Theta_{\text{camera}}) \\ D(\Theta) &= D_{\text{shape}} D_{\text{pose}} D_{\text{camera}} \\ C &= E^\dagger D(\Theta) E T\end{aligned}$$




Factorization

$$\Theta = (\Theta_{\text{shape}}, \Theta_{\text{pose}}, \Theta_{\text{camera}})$$

$$D(\Theta) = D_{\text{shape}} D_{\text{pose}} D_{\text{camera}}$$


$$C = E^\dagger D(\Theta) E T$$




Θ_{shape}^1

Θ_{shape}^2


Θ_{shape}^3



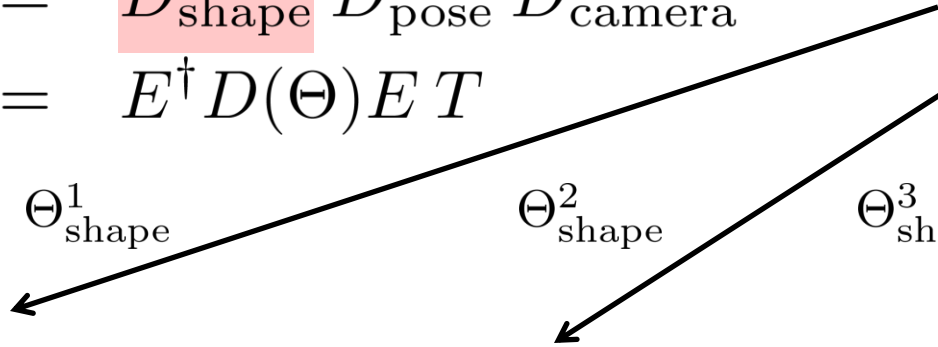
C_1



C_2



C_3




Factorization

$$\Theta = (\Theta_{\text{shape}}, \Theta_{\text{pose}}, \Theta_{\text{camera}})$$

$$D(\Theta) = D_{\text{shape}} D_{\text{pose}} D_{\text{camera}}$$


$$C = E^\dagger D(\Theta) E T$$




Θ_{pose}^1

Θ_{pose}^2


Θ_{pose}^3




C_1



C_2



C_3

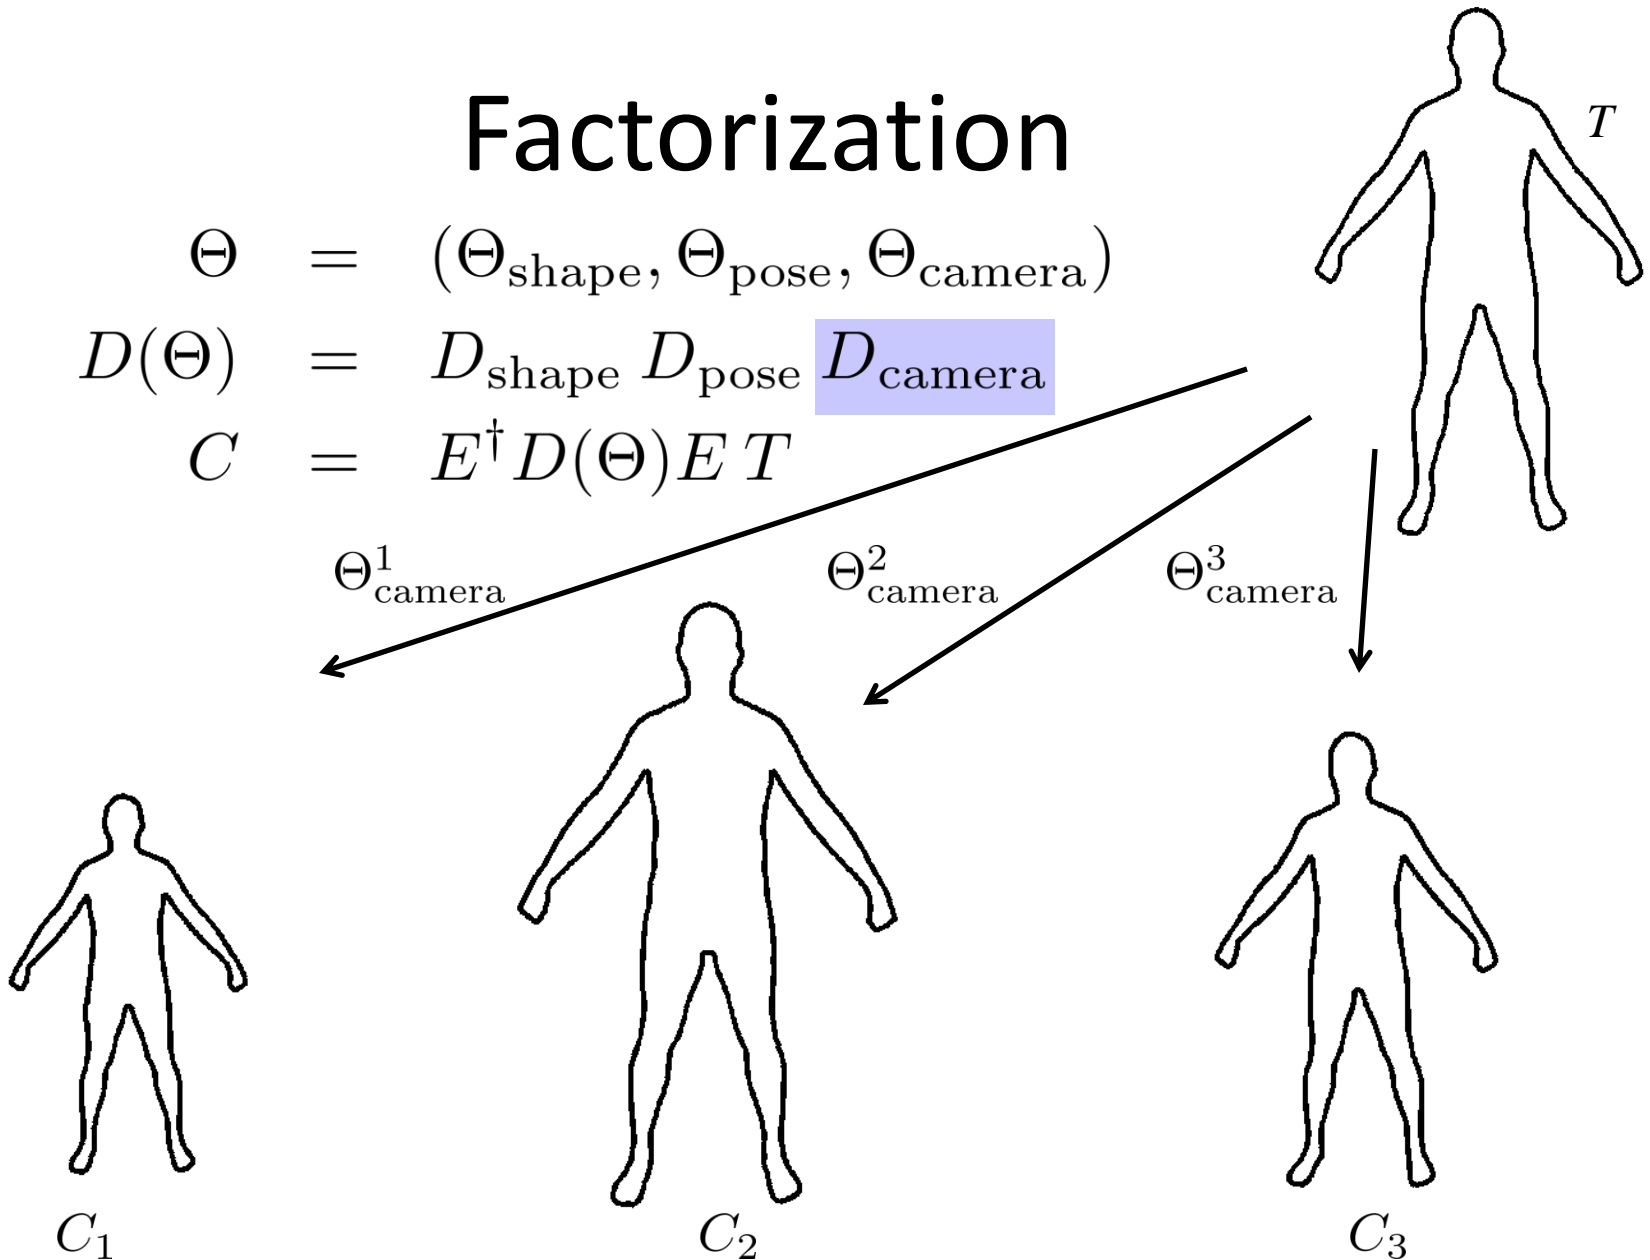


Factorization

$$\Theta = (\Theta_{\text{shape}}, \Theta_{\text{pose}}, \Theta_{\text{camera}})$$

$$D(\Theta) = D_{\text{shape}} D_{\text{pose}} D_{\text{camera}}$$

$$C = E^\dagger D(\Theta) E T$$

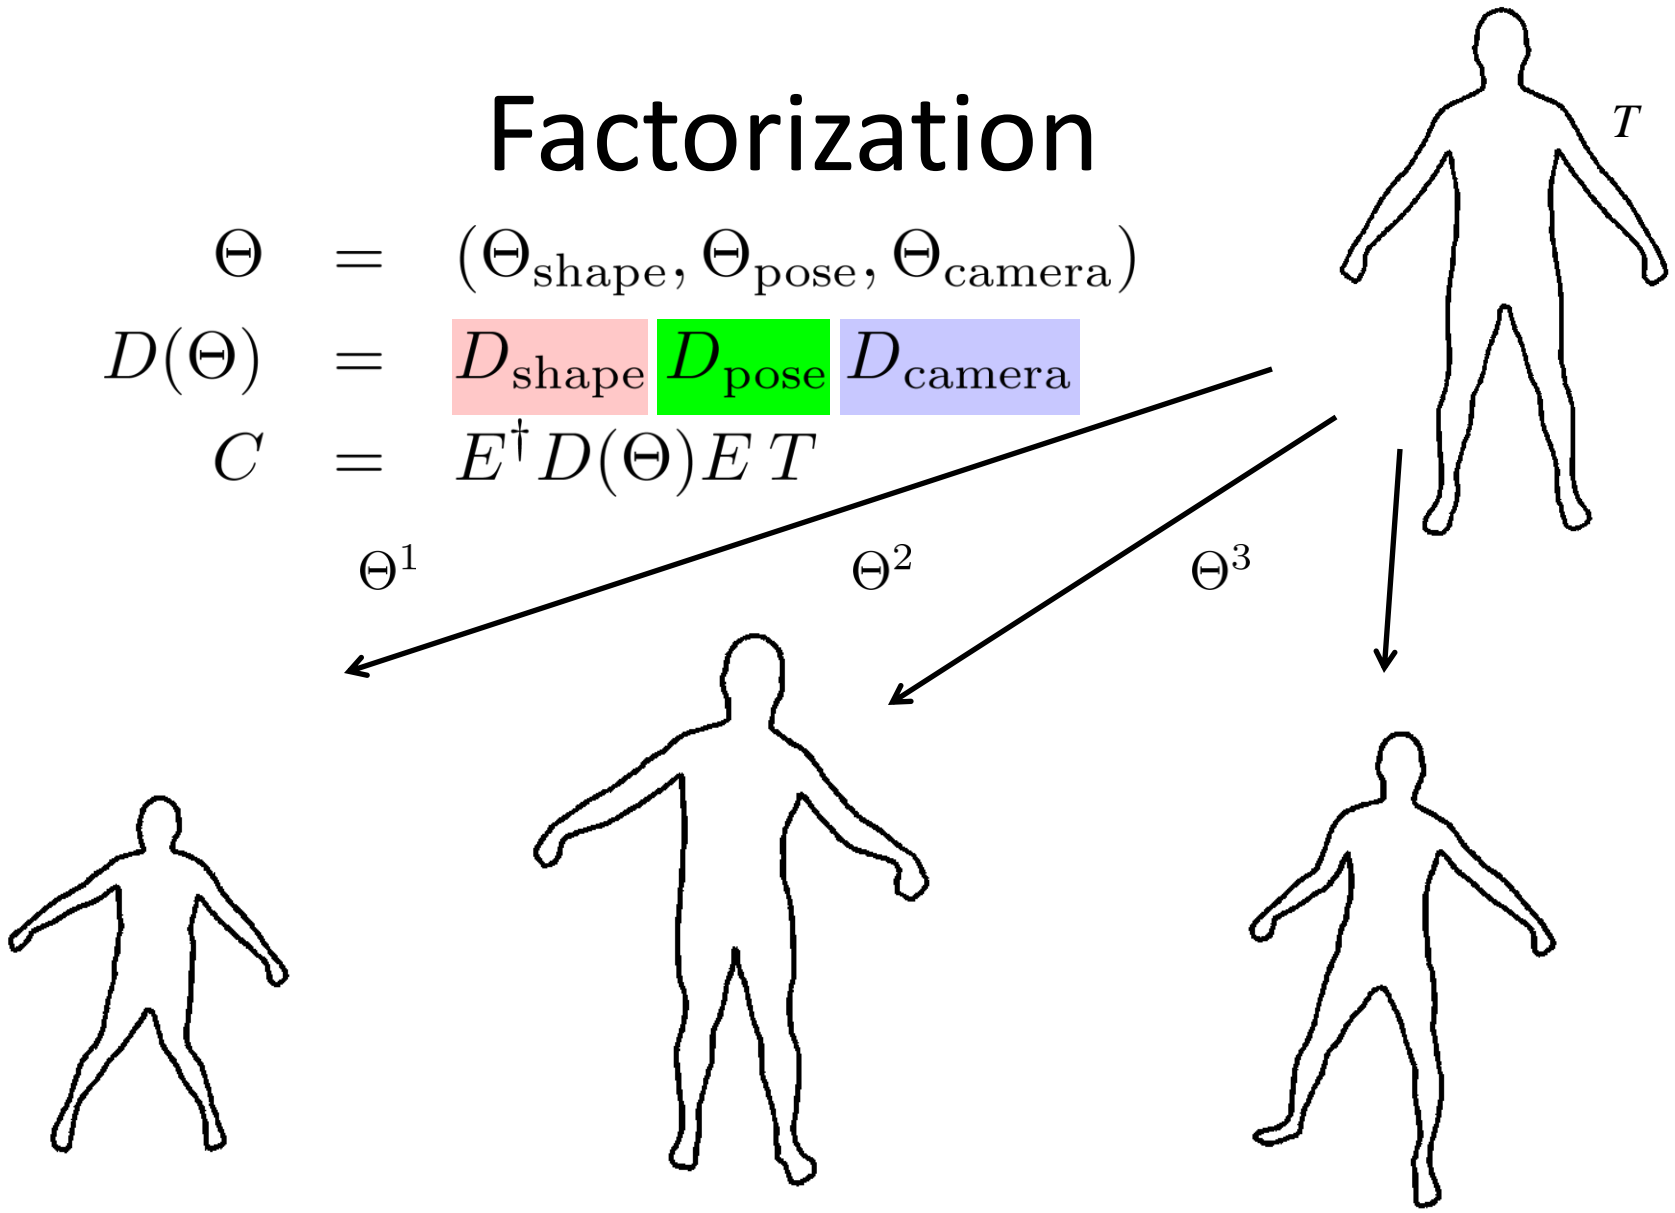


Factorization

$$\Theta = (\Theta_{\text{shape}}, \Theta_{\text{pose}}, \Theta_{\text{camera}})$$


$$D(\Theta) = D_{\text{shape}} D_{\text{pose}} D_{\text{camera}}$$

$$C = E^\dagger D(\Theta) E T$$




Learning

$$\begin{aligned}\Theta &= (\Theta_{\text{shape}}, \Theta_{\text{pose}}, \Theta_{\text{camera}}) \\ D(\Theta) &= D_{\text{shape}} D_{\text{pose}} D_{\text{camera}} \\ C &= E^\dagger D(\Theta) E T\end{aligned}$$




- How do we learn this parametric model?
 - Create random shapes and poses from SCAPE
 - Create random cameras
 - Project 3D body to 2D (but keep the body-part segmentation)

Shape training examples




Example #1




$D_{\text{shape}}^{(1)}$

$C_{\text{shape}}^{(1)}$

Shape training examples




Example #170




$D^{(170)}$
shape

$C^{(170)}$
shape

Shape training examples




Example #519



$D^{(519)}$
shape

$C^{(519)}$
shape

Pose training examples




Example #1




$D_{\text{pose}}^{(1)}$

$C_{\text{pose}}^{(1)}$

Pose training examples



Example #22




$D_{\text{pose}}^{(22)}$


An arrow points from the target pose T to the shape contour $C_{\text{shape}}^{(22)}$, with the label $D_{\text{pose}}^{(22)}$ placed next to the arrow.

$C_{\text{shape}}^{(22)}$

Pose training examples




Example #450




$D_{\text{pose}}^{(450)}$

$C_{\text{pose}}^{(450)}$

Camera training examples




Example #1




$D_{\text{camera}}^{(1)}$

$C_{\text{camera}}^{(1)}$

Camera training examples




Example #71




$D_{\text{camera}}^{(71)}$

$C_{\text{camera}}^{(71)}$

Camera training examples



Example #645



$D_{\text{camera}}^{(645)}$

$C_{\text{camera}}^{(645)}$

Training set

$$D_i^{2 \times 2} = S_i \begin{pmatrix} \cos \theta_i & -\sin \theta_i \\ \sin \theta_i & \cos \theta_i \end{pmatrix}$$

$$s_i = \log(S_i)$$

$$A_{\text{shape}} = \begin{pmatrix} \dots & s_1^{(j)} & \dots \\ \dots & \theta_1^{(j)} & \dots \\ \vdots & \vdots & \vdots \\ \dots & s_n^{(j)} & \dots \\ \dots & \theta_n^{(j)} & \dots \end{pmatrix}$$



$$D_{\text{shape}}^{(j)}$$

The deformation
of the j^{th}
contour person

$n = \# \text{line segments}$



Shape examples



Principal Component Analysis

- PCA on A_{shape} :
$$\theta_i = \bar{\theta}_i + \sum \beta_k \theta_i^k$$
$$s_i = \bar{s}_i + \sum \beta_k s_i^k$$
$$\Theta_{\text{shape}} = (\beta_1, \dots, \beta_K)$$

k stands for the k^{th} eigenvector

Principal Component Analysis

- PCA on A_{shape} :
$$\theta_i = \bar{\theta}_i + \sum \beta_k \theta_i^k$$
$$s_i = \bar{s}_i + \sum \beta_k s_i^k$$
$$\Theta_{\text{shape}} = (\beta_1, \dots, \beta_K)$$

k stands for the k^{th} eigenvector

- This defines

$$D_{\text{shape}} = D_{\text{shape}}(\Theta_{\text{shape}})$$

- We do the same for A_{camera} and this defines


$$D_{\text{camera}} = D_{\text{camera}}(\Theta_{\text{camera}})$$

Principal Component Analysis

- Why use angle and log-scale? Interpretation:
PCA in a Lie-algebra [Vaillant et al. '04]

$$S_i \begin{pmatrix} \cos \theta_i & -\sin \theta_i \\ \sin \theta_i & \cos \theta_i \end{pmatrix} = \exp \left(\left(\begin{pmatrix} s_i & 0 \\ 0 & s_i \end{pmatrix} + \begin{pmatrix} 0 & -\theta_i \\ \theta_i & 0 \end{pmatrix} \right) \right)$$

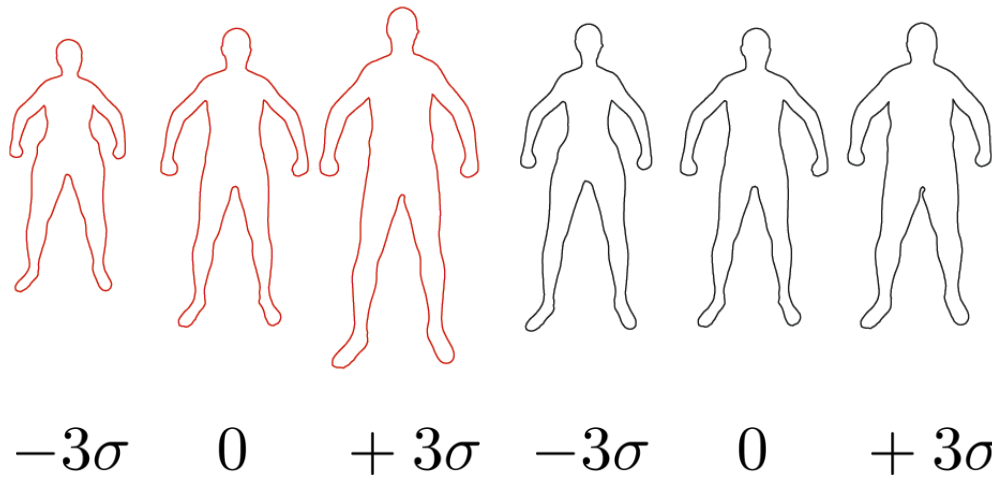
Eigen-shapes



-3σ 0 $+3\sigma$

Principal component #1


Eigen-shapes



Principal component #1

Principal component #2

Eigen-shapes



-3σ

0

$+3\sigma$

-3σ

0

$+3\sigma$

-3σ

0


$+3\sigma$

Principal component #1

Principal component #2

Principal component #3

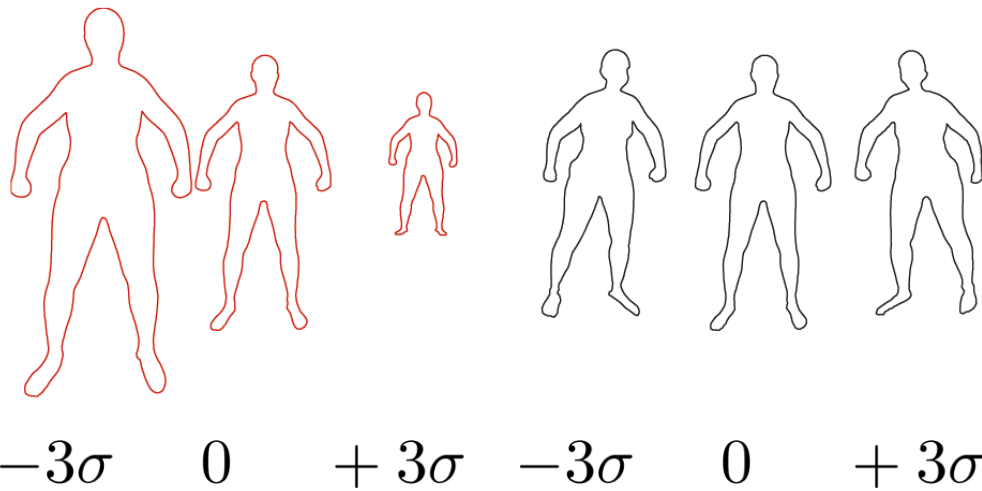
Eigen-cameras



-3σ 0 $+3\sigma$

Principal component #1


Eigen-cameras



Principal component #1

Principal component #2

Eigen-cameras




Principal component #1

Principal component #2

Principal component #3

Pose: rigid + non-rigid deformations


(PS meets SCAPE)



$$D_{\text{pose}} = D_{\text{pose}}(\Theta_{\text{pose}})$$


Pose: rigid + non-rigid deformations

(PS meets SCAPE)



$$\forall i \in \text{arm} \quad \theta_i = \theta_{\text{RIGID}}$$
$$s_i = s_{\text{RIGID}}$$

Pose: rigid + non-rigid deformations (PS meets SCAPE)

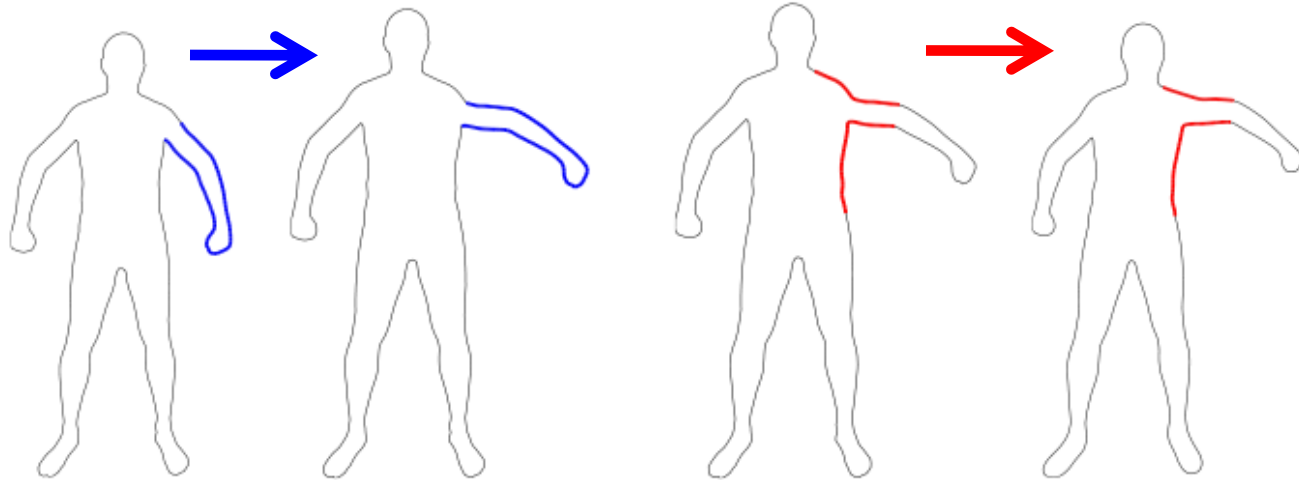


$$\begin{pmatrix} \Delta\theta_i \\ \Delta s_i \end{pmatrix} = H_i(\theta_{\text{RIGID}}, s_{\text{RIGID}}, 1)^T$$

$$\theta_i = \theta_{\text{RIGID}} + \Delta\theta_i$$

$$s_i = s_{\text{RIGID}} + \Delta s_i$$

Pose: rigid + non-rigid deformations (PS meets SCAPE)



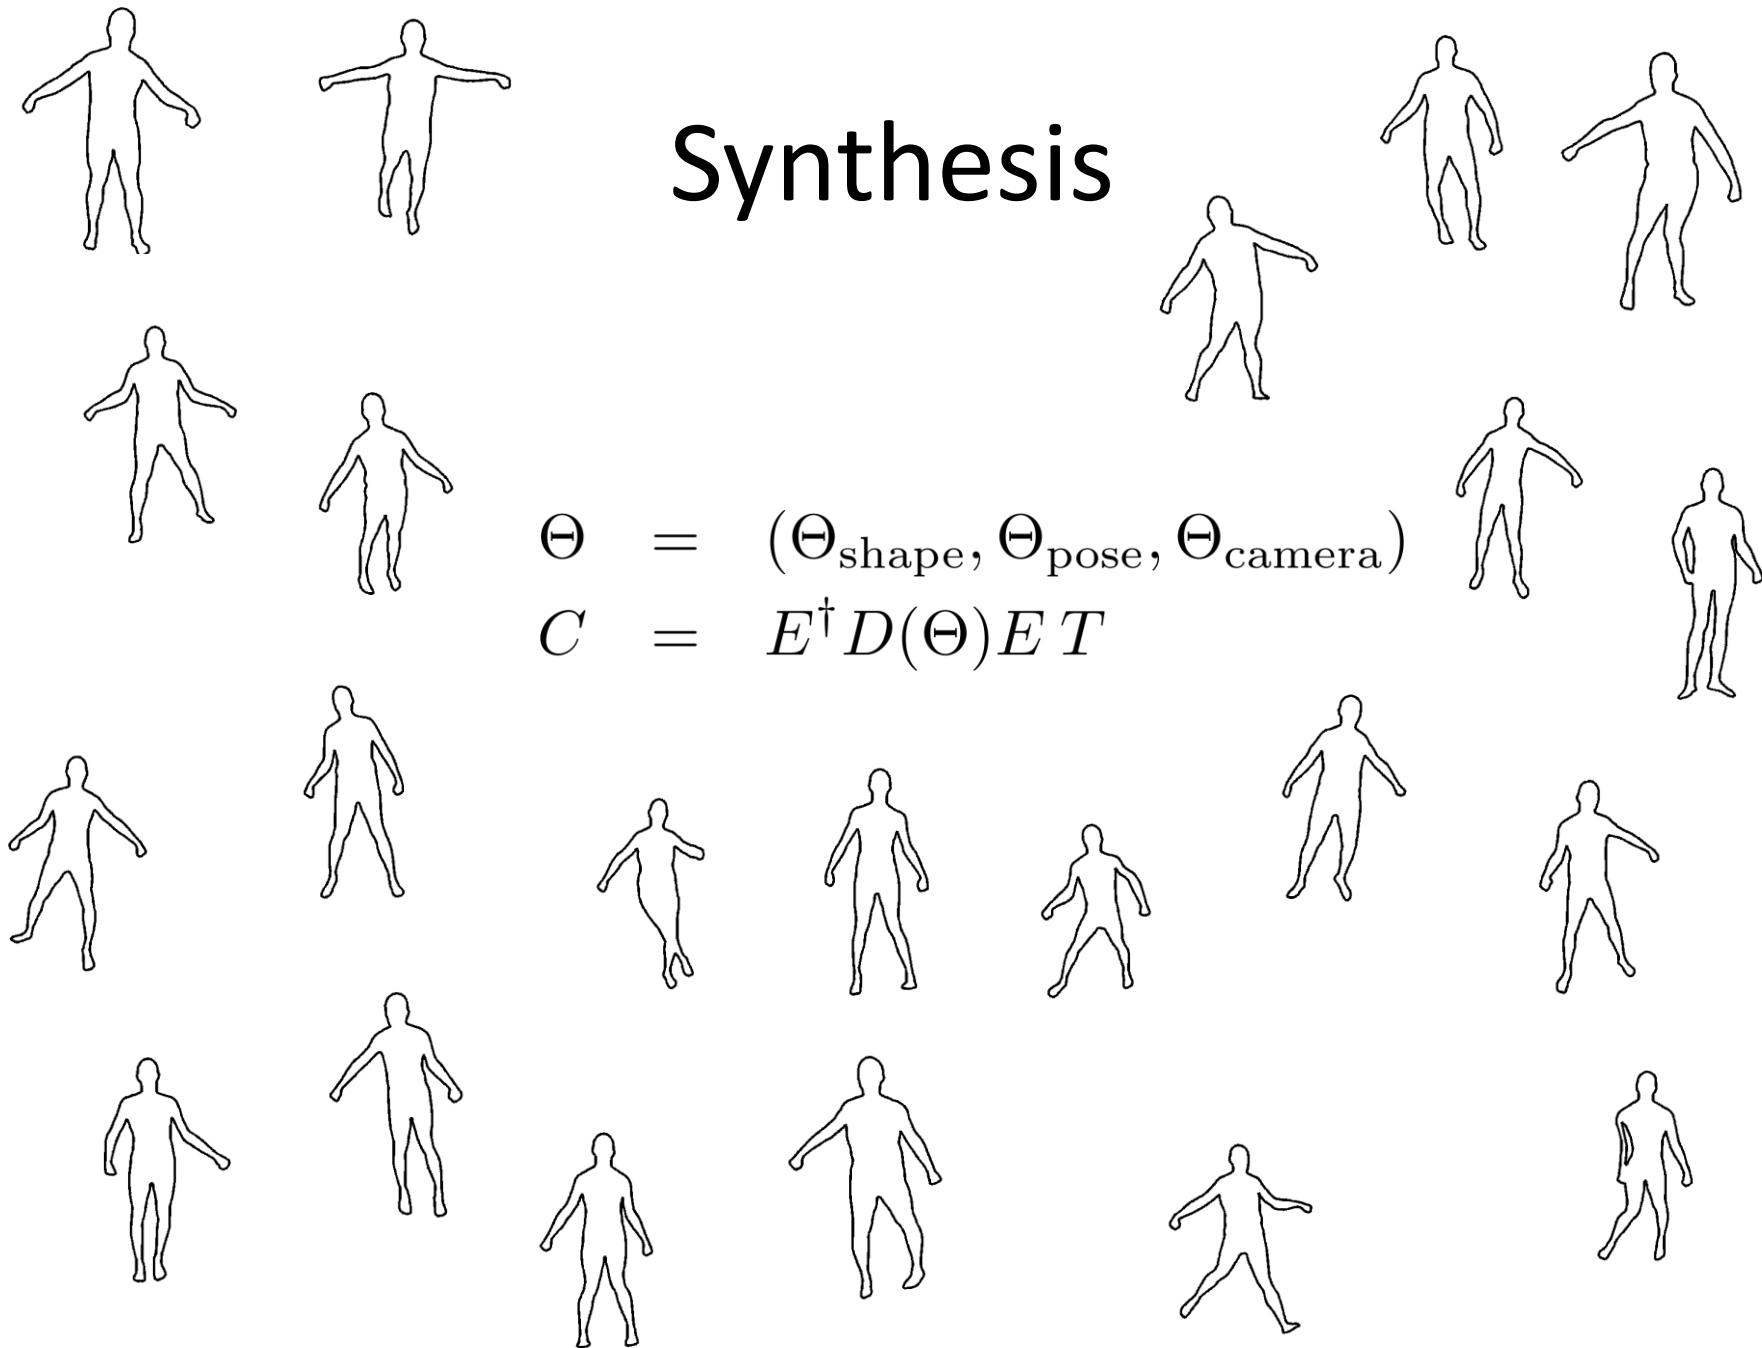
$$\begin{pmatrix} \Delta\theta_i \\ \Delta s_i \end{pmatrix} = H_i(\theta_{\text{RIGID}}, s_{\text{RIGID}}, 1)^T$$

The matrix H_i is learned from examples

Synthesis

$$\Theta = (\Theta_{\text{shape}}, \Theta_{\text{pose}}, \Theta_{\text{camera}})$$


$$C = E^{\dagger} D(\Theta) E T$$




Fitting CP to silhouettes

- We minimize a cost function of the form
$$F(\text{mask}, \Theta) = d(\text{mask}, E^\dagger D(\Theta) E T)$$


E.g., d can be the bi-directional Chamfer distance




mask



T



Selected results



Combine pose/shape estimation with segmentation

- Similar to PoseCut and ObjCut

[Bray et al. ECCV '06, Pawan et al. PAMI '10]

- We minimize a cost function of the form

$$F(I, \Theta) = F_{\text{region}}(I, \Theta) + F_{\text{edge}}(I, \Theta) + F_{\text{prior}}(\Theta)$$

- We used a PS algorithm as an initialization

[Andriluka et al. CVPR '09]

Selected results

PS

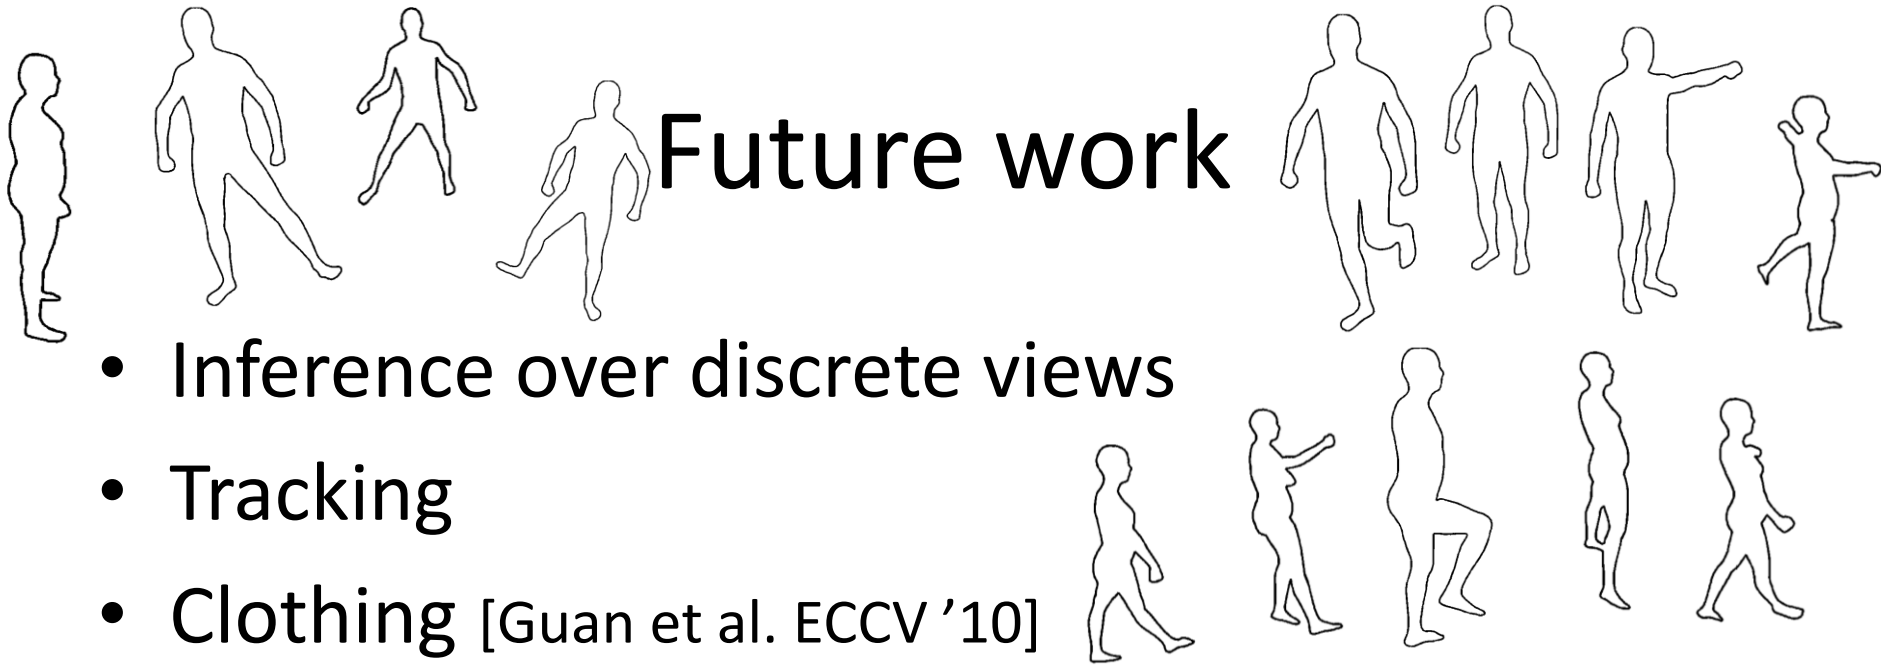



CP




Future work

- Inference over discrete views
- Tracking
- Clothing [Guan et al. ECCV '10]
- Coarse-to-fine inference (PS to CP)








Future work




- Inference over discrete views
- Tracking
- Clothing [Guan et al. ECCV '10]
- Coarse-to-fine inference (PS to CP)






Future work




- Inference over discrete views
- Tracking
- Clothing [Guan et al. ECCV '10]
- Coarse-to-fine inference (PS to CP)



Big question

- PS-like inference using a part-based structure?



Conclusions

- A new 2D part-based generative model of humans
 - Beyond previous deformable template models [Cootes et al. CVIU '95, Baumberg & Hogg ECCV '94, and many others...]
- Factors shape, pose, and camera deformations
- Has advantages of PS models with the detail of a SCAPE-like model
- Initial application: Human-specific image segmentation

Acknowledgments

- This work was funded in part by
 - NIH EUREKA 1R01NS066311-01 grant
 - NSF IIS-0812364 grant
- We thank D. Hirshberg and A. Balan for their assistance with the code
- We thank Andriluka et al. for their PS implementation