Show me your Body: Gender Classification from Still Images

Ioannis A. Kakadiaris[1], Nikolaos Sarafianos[1] and Christophoros Nikou[1,2]

[1]Computational Biomedicine Lab, Dept. of Computer Science, University of Houston, Houston, TX
[2]Department of Computer Science and Engineering, University of Ioannina, Ioannina, Greece

Overview

Training & Testing
Observable information:
- Arm length
- Knee height
- Waist height
- Hip breadth

Training Only
Privileged information:
- Hip circumference
- Chest circumference
- Ankle circumference

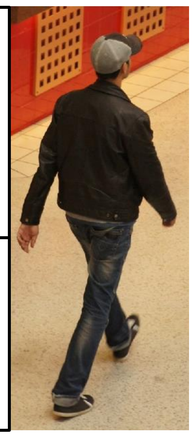

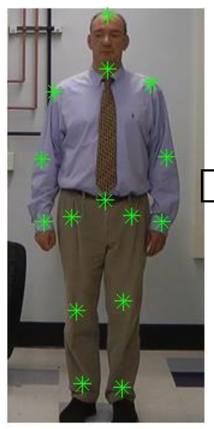

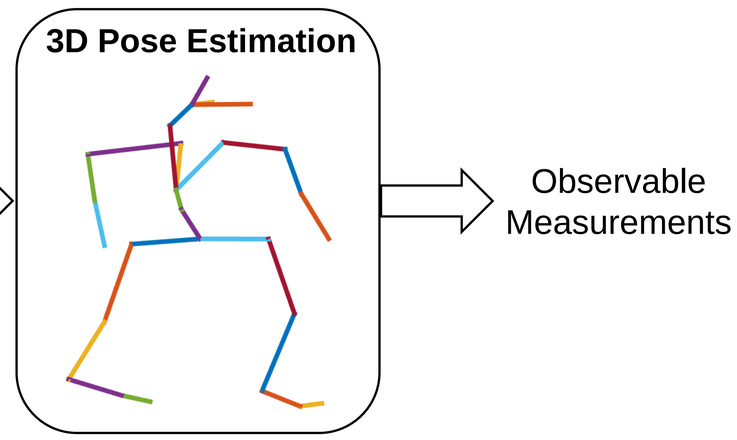

Introduction

Problem Statement
- Predict the gender using human metrology

Motivation
- Explore the use of ratios of anthropometric measurements for gender estimation
- Exploit privileged information available during training

Background
- Observable features: Information available at both training and testing
- Privileged features: Information available only at training time

Method

Ratios of Anthropometric Measurements
- Original features: $X = NxM_1, M_1 = \frac{F_1 x (F_1 - 1)}{2}$, F_1 the number of original measurements
- Privileged features $X^* = NxM_2, M_2 = \frac{F_2 x (F_2 - 1)}{2}$, F_2 the number of privileged measurements
- Features are split to upper body (X_U) and lower body (X_L) sets

Classification

SVM: $\min_{\xi_1,\dots,\xi_N,w,b} \left\{ \frac{1}{2}||w||^2 + C\sum_{i=1}^{N} \xi_i \right\}$

$$\text{s.t. } y_i(\langle w_i, x_i \rangle + b) \geq 1 - \xi_i, \xi_i \geq 0, i = 1, \dots, N$$

SVM+: $\min_{\xi_1,\dots,\xi_N,w,b} \left\{ \frac{1}{2}(||w||^2 + \gamma||w^*||^2) + C\sum_{i=1}^{N} \xi_i(w^*, b^*) \right\}$

$$\text{s.t. } y_i(\langle w_i, x_i \rangle + b) \geq 1 - \xi_i(w^*, b^*)\ \xi_i(w^*, b^*) \geq 0, \quad i = 1, \dots, N$$

Margin Transfer: $\min_{\xi_1,\dots,\xi_N,w,b} \left\{ \frac{1}{2}||w||^2 + C\sum_{i=1}^{N} \xi_i \right\}$

$$\text{s.t. } y_i(\langle w_i, x_i \rangle + b) \geq \rho_i - \xi_i, \xi_i \geq 0, i = 1, \dots, N$$

w the weight vector, b the bias parameter, ξ_i is the slack, and C denotes the penalty parameter

Results

CAESAR dataset - Testing Features	SVM	SVM+
Only observable (X)	97.61 ± 0.44	**98.18 ± 0.56**
Lower body observable (X_L)	95.34 ± 0.74	**95.82 ± 0.81**
Upper body observable (X_U)	**76.69 ± 2.98**	76.54 ± 2.95
Observable & Privileged ($X+X^*$)	**99.10 ± 0.23**	-
Cao *et al.* [1]	99.37	-

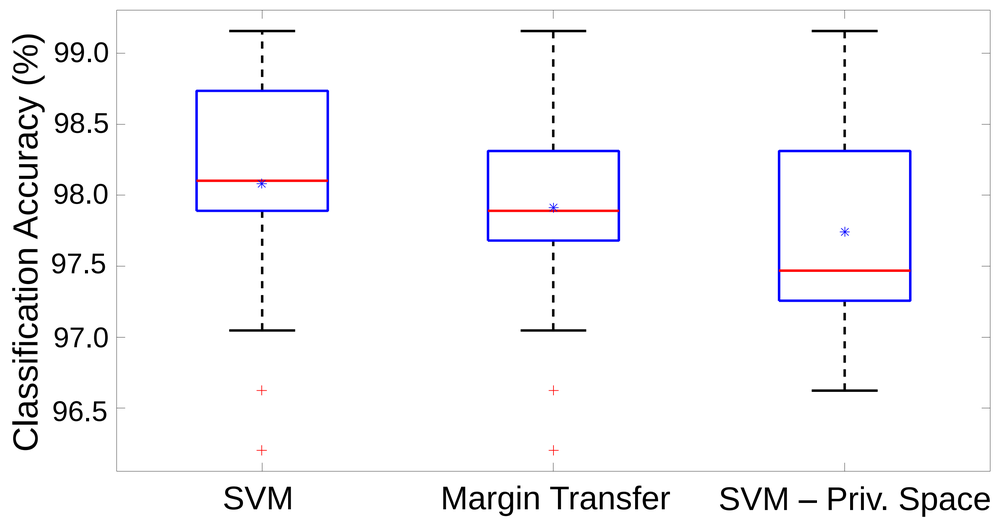

Image Dataset	Set of features: X	X_L	X_U
PaSC	71.37 ± 1.64	57.65 ± 2.82	58.06 ± 2.73
SARC3D	86.00 ± 2.00	78.00 ± 4.00	72.00 ± 4.00

Contributions

- Using privileged information improves the classification accuracy
- Ratios of measurements are as discriminative as the actual values
- Predicting the gender from images using anthropometry is feasible

References

[1] D. Cao, C. Chen, D. Adjeroh and A. Ross, "Predicting gender and weight from human metrology using a copula model," in Proc. 5th IEEE International Conference on Biometrics Theory, Applications and Systems, Washington, DC, USA, Sep. 23-26 2012, pp. 162-169.

Acknowledgements

This research was funded in part by the UH Hugh Roy and Lillie Cranz Cullen Endowment Fund and the European Commission (H2020-MSCA-IF-2014), under grant agreement No 656094.

UNIVERSITY of HOUSTON | CBL Changing the way people look at computers