Supplementary Material for ECCV 2014 Paper: Distance estimation of an unknown person from a single portrait

Xavier P. Burgos-Artizzu[1,2] Matteo Ruggero Ronchi[2] Pietro Perona[2]

[1]Technicolor Research - Cesson Sévigné, France
[2]California Institute of Technology - Pasadena, CA, USA
xavier.burgos@technicolor.com mronchi,perona@caltech.edu

Abstract. This document accompanies the paper "Distance estimation of an unknown person from a single portrait". We provide further insight on some of the parameter/method choices made in the main paper and additional physiognomy interpretation of the results.

1 Parameter Sweeps

To validate the implementation choices discussed in the main paper, we illustrate the impact of performing the regression over the direct distance value of an image or its inverse, Figure 1(a), as well as the effect of using several regression methods, Figure 1(b). We report our findings only for the distance regression task.

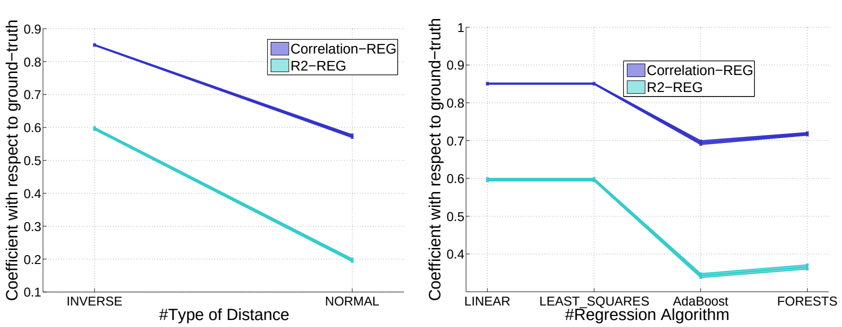

(a) (b)

Fig. 1. How various method parameters affect the performance. The vertical axis represents the correlation between ground truth and the estimate of distance obtained by our algorithm in the regression task (higher corresponds to a better prediction). (a) Effect of using normal vs inverse distance. (b) Performance of different regression methods.

Based on the fact that performance grows without saturating when increasing the number of subjects in the Training Set (as discussed in Figure 9(a) of the main paper), we carried out these experiments in the optimal setting of using a lave-one-out validation scheme on the entire dataset. For each subject, a linear regressor was trained using the pictures of all the other 52 subjects in the dataset and then tested only on its seven distances.

- **Figure 1(a).** As mentioned in the paper, regression on the inverse distance greatly improves performance compared to regular distance. This is due to the saturation of the signal at high distances.
- **Figure 1(b).** We tested multiple regression methods. As shown, the best performances were achieved with the simplest methods, such as multivariate linear regression. We attribute this to the low number of training examples which causes classification algorithms (AdaBoost and Random Forests) to overfit the data.

2 Physiognomy Interpretation

One of the main conclusions drawn from the experiments reported in the paper is that physiognomy biases systematically the estimate. Therefore, we measured for all the faces in the dataset their average bias in the estimated distance across all distances. We show our findings over the whole dataset in Figure 2. Furthermore, we investigated the correlation between facial properties of the subjects, such as the width / height ratio and the bias in estimated distance, however no evident pattern was found, as shown in Figure 3.

Further understanding the relationship and patterns between subject physiognomy and distance estimate bias could allow us to cluster subjects into templates based on their appearance and have a great impact on the performance of our algorithm. This will be one of our main lines of investigation in the future.

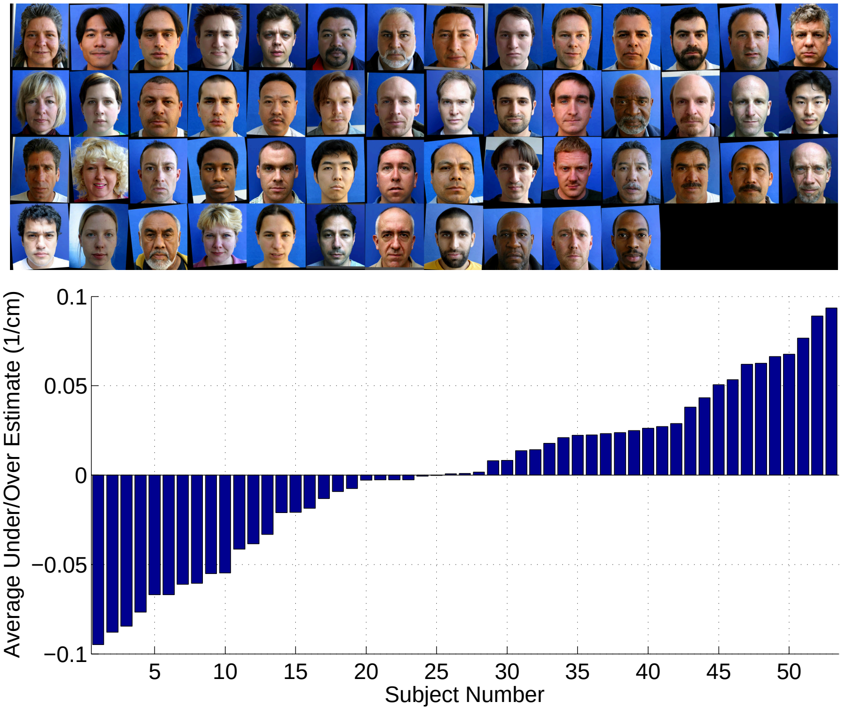

Fig. 2. (Top) All the subjects belonging to the dataset ordered by terms of the bias in their distance estimation, averaged over all the seven images of a subject. (Bottom) The value of the bias in the estimate of the distance of an image over all the distances for a certain subject. The correspondence between images and bars start with the top-left image and leftmost column and continues over the rows of the dataset montage

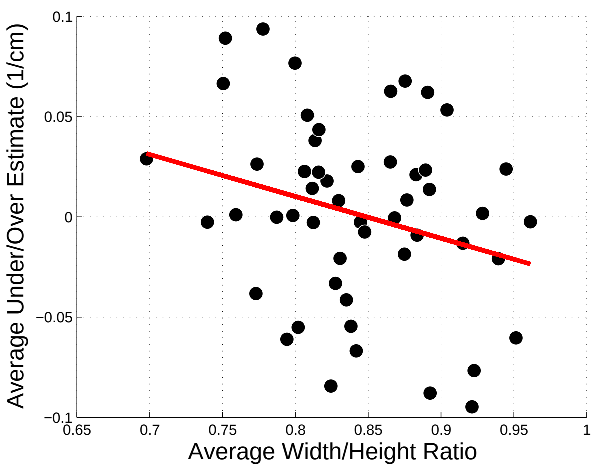

Fig. 3. There is no apparent correlation between the width / height ratio of a subject's face and it's overall bias in the distance estimation. Very squared (or round) and elongated faces are respectively plotted on the right and left extreme of the x-axis. On the y-axis, faces are displayed from those whose distance is over estimated on average (top) to those with under estimated distance (bottom)