Data-Driven Beautification

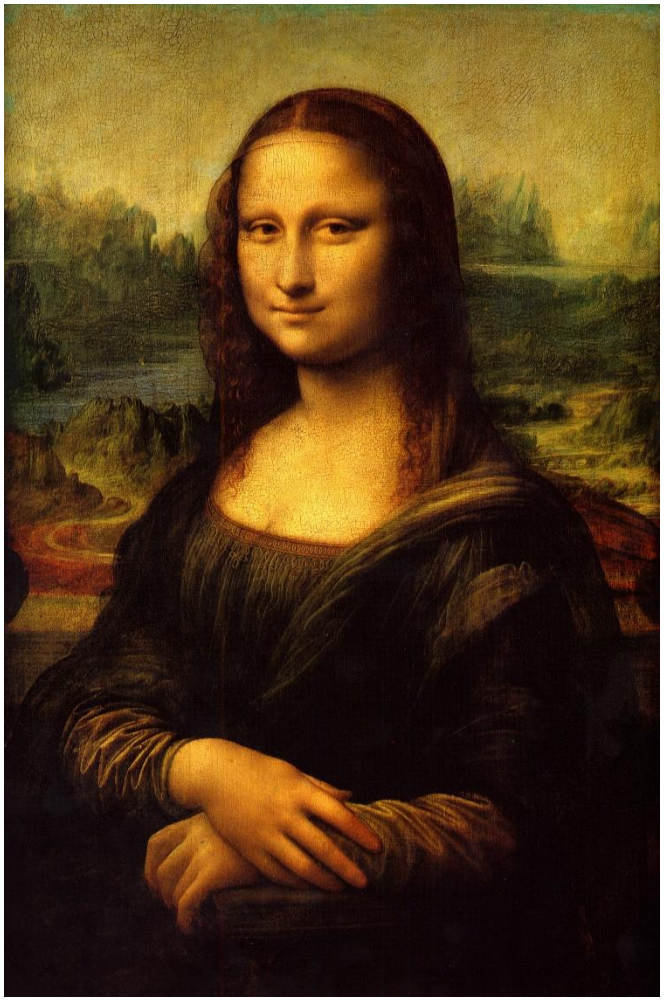

Leonardo da Vinci, “Mona Lisa”

CSC320: Introduction to Visual Computing
Michael Guerzhoy

Data-Driven Enhancement of Facial Attractiveness

Tommer Leyvand, Tel-Aviv University
Daniel Cohen-Or, Tel-Aviv University
Gideon Dror, Academic College of Tel-Aviv-Yaffo
Dani Lischinski, The Hebrew University

Abstract

When human raters are presented with a collection of shapes and asked to rank them according to their aesthetic appeal, the results often indicate that there is a statistical consensus among the raters. Yet it might be difficult to define a succinct set of rules that capture the aesthetic preferences of the raters. In this work, we explore a data-driven approach to aesthetic enhancement of such shapes. Specifically, we focus on the challenging problem of enhancing the aesthetic appeal (or the *attractiveness*) of human faces in frontal photographs (portraits), while maintaining close similarity with the original.

The key component in our approach is an automatic facial attractiveness engine trained on datasets of faces with accompanying facial attractiveness ratings collected from groups of human raters. Given a new face, we extract a set of distances between a variety of facial feature locations, which define a point in a high-dimensional "face space". We then search the face space for a nearby point with a higher predicted attractiveness rating. Once such a point is found, the corresponding facial distances are embedded in the plane and serve as a target to define a 2D warp field which maps the original facial features to their adjusted locations. The effectiveness of our technique was experimentally validated by independent rating experiments, which indicate that it is indeed capable of increasing the facial attractiveness of most portraits that we have experimented with.

Keywords: facial attractiveness, machine learning, optimization, warping

Figure 1: Input facial images (left) and the adjusted images generated by our method (right). The changes are subtle, yet their impact is significant.

1 Introduction

SIGGRAPH 2008

Step 1: Learn to predict humans' evaluation of beauty

- Collect lots of faces
- Extract a set of *feature points* on the face (NOTE: this is not easy)
- The *features* used in training are the 234 distances b/w the feature points
- Have humans rate beauty to construct a training set
- Learn to predict human evaluation

Step 2: Finding a Nearby More Beautiful Face

- For an input face, compute the features (by first extracting the feature points, computing the distances between them)
- Find the k nearest neighbours with similar features, and compute their weighted average, giving more weight to the beautiful features
  - Hopefully, the average still represents a face
- Now, using optimization, find a configuration of *feature points* which produces the beautiful set of features (i.e., distances) obtained above
- Warp the face onto the new feature points

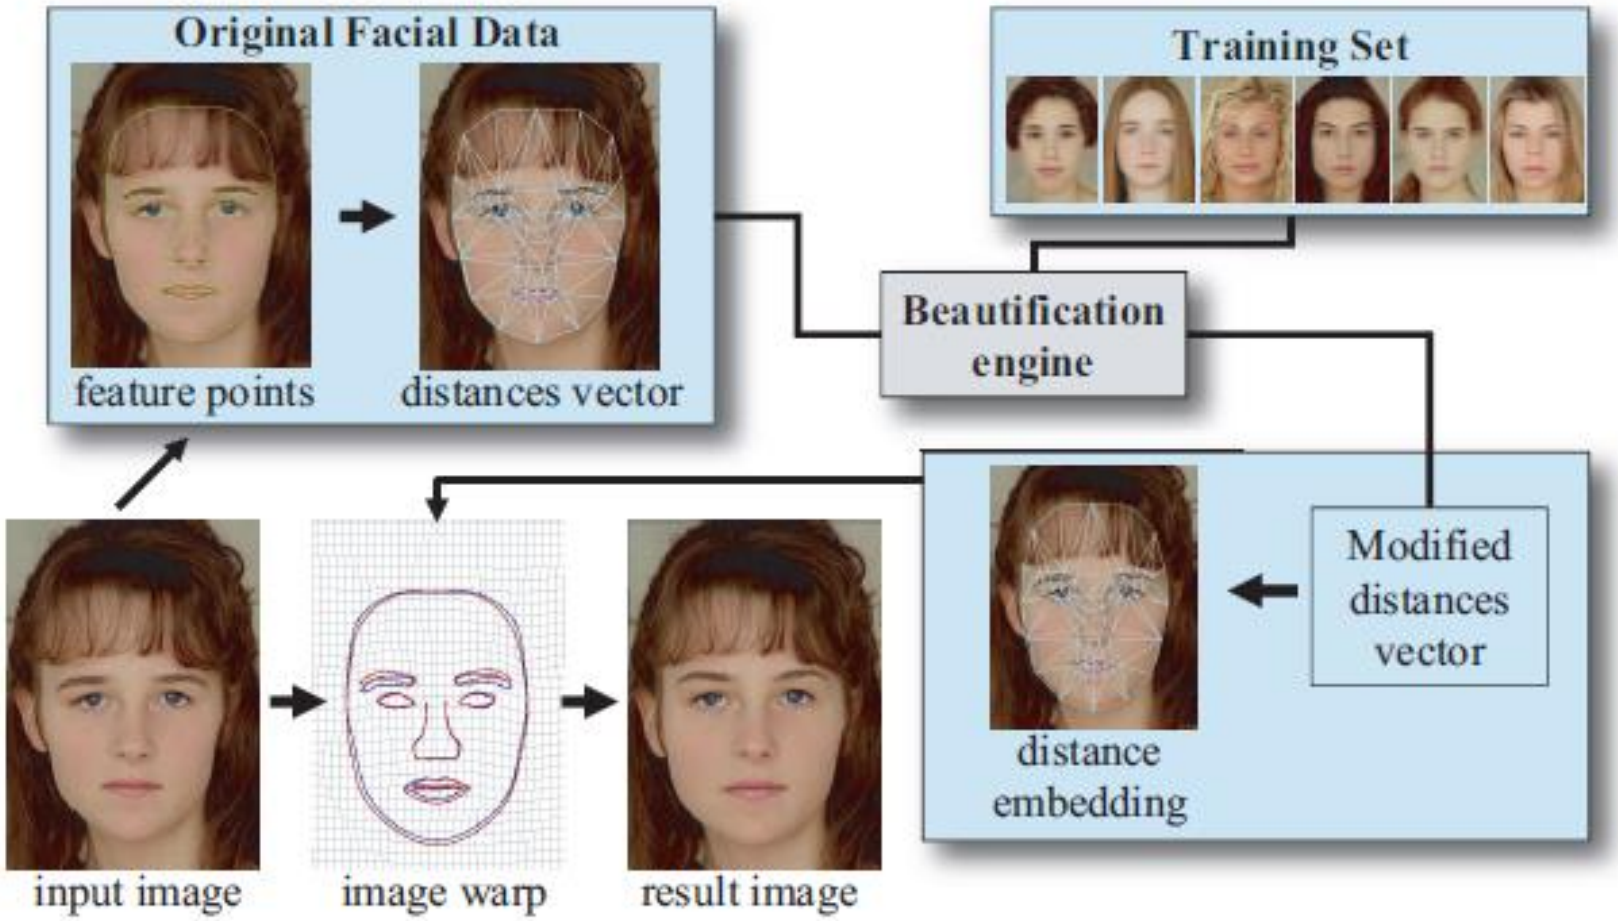
Original Facial Data
feature points
distances vector
Training Set
Beautification engine
Modified distances vector
distance embedding
input image
image warp
result image

Results