Face Painting: querying art with photos

Elliot J. Crowley
elliot@robots.ox.ac.uk
Omkar M. Parkhi
omkar@robots.ox.ac.uk
Andrew Zisserman
az@robots.ox.ac.uk

Visual Geometry Group
Department of Engineering Science
University of Oxford

We study the problem of matching photographs of people to similar looking paintings in a large corpus.

Of course, the extent to which a person in a photo resembles a different person in a painting is subjective, and very difficult to quantify. So, instead, we consider the question: given *photographs* of a person, can we retrieve *paintings* of that same person? If we can, we should then be able to find different, but similar looking, people in paintings, starting from a photo.

Initially, one might be skeptical over whether retrieving paintings of a person starting from a photo is achievable. Photographs and paintings have very different low level statistics and to make matters worse, painted portraits are prone to large variations in style: politicians are often highly caricatured, Hollywood icons of the past frequently get the Andy Warhol treatment and are transformed into pop art. This problem is essentially one of domain adaptation from faces in photos to those in paintings; learning how to overcome both the low-level and stylistic differences.

Contributions. (i) we introduce new datasets for learning and evaluating this adaptation problem (figure 1); (ii) we show that, depending on the face representation used, performance can be improved substantially by learning – either by a linear projection matrix common across identities, or by a per-identity classifier (figure 2); (iii) we compare Fisher Vector [6] and Convolutional Neural Network (CNN) representations [5] for this task – some successful painting retrievals from photos using CNN features are provided in figure 3. Finally, (iv) using the learnt descriptors, we show that, given a photo of a person, we are able to find their doppelgänger in a large dataset of oil paintings (figure 4).

Figure 1: **Painting Datasets.** Top row – Paintings from the National Portrait Gallery, these are largely photo-realistic in nature. Bottom row – Paintings crawled from DeviantArt [2], which are much more varied in style. The contrast between these datasets allows us to observe what effect large variations in style have on retrieval. Identities (a)-(q) are given in the paper.

Figure 2: **Photo-painting pairs used for learning.** In each case the photo is on the left and the painting is on the right. The photos are obtained from Google Image Search and the paintings from DeviantArt. These photos and paintings are used to learn either (i) a linear projection on the features using discriminative dimensionality reduction (an idea first explored for faces in [3]) or (ii) exemplar-like [4] face-specific classifiers to boost retrieval performance.

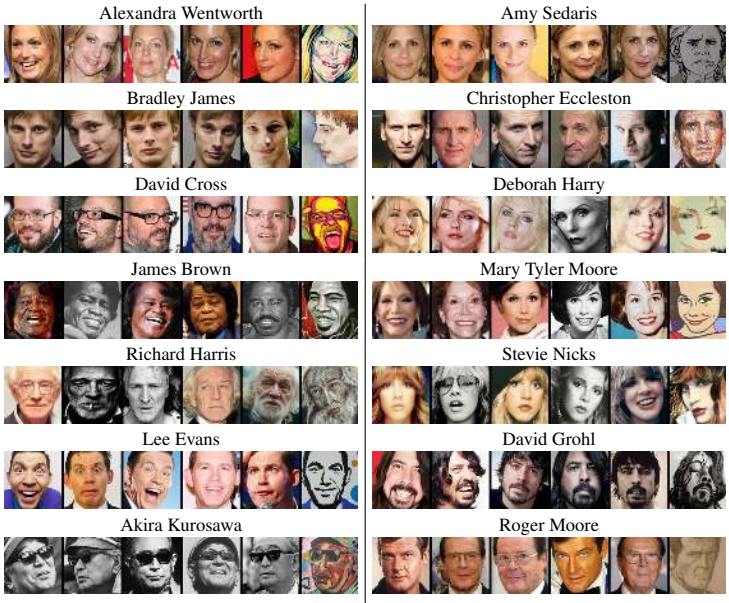

Figure 3: **Successful retrievals using a CNN representation.** In each case, five photos are used to make a query. These are shown beside the top retrieved painting.

Figure 4: **Celebrities and their doppelgängers.** Each photo is matched to the closest portrait in the 'Your Paintings' dataset of oil paintings [1] using Euclidean distance between CNN features.

[1] BBC – Your Paintings. http://www.bbc.co.uk/arts/yourpaintings/.

[2] DeviantArt. http://www.deviantart.com/.

[3] M. Guillaumin, J. Verbeek, and C. Schmid. Is that you? Metric learning approaches for face identification. In *Proc. ICCV*, 2009.

[4] T. Malisiewicz, A. Gupta, and A. A. Efros. Ensemble of exemplar-SVMs for object detection and beyond. In *Proc. ICCV*, 2011.

[5] O. M. Parkhi, A. Vedaldi, and A. Zisserman. Deep face recognition. In *Proc. BMVC*, 2015.

[6] F. Perronnin, J. Sánchez, and T. Mensink. Improving the Fisher kernel for large-scale image classification. In *Proc. ECCV*, 2010.