

Mooney Faces from Photos

Tsung-Wei Ke

Stella X. Yu

David Whitney

UC Berkeley / ICSI

Vision Sciences Society Annual Meeting, 2017

Abstract

Mooney faces are simply black and white pictures that human subjects can still effortlessly recognize as a face, along with many attributes such as age, gender, and facial expressions. How such a remarkable ability is achieved with so little information is key to understanding human visual recognition. However, a very basic question is what makes a Mooney face, as not every binarized face image appears as a Mooney face. Current Mooney face generation is ad-hoc, requiring manual grayscale image editing and human subjective judgment, a time and labor-consuming process. Our goal is to generate Mooney faces from photos automatically.

We take a deep learning approach based on very limited Mooney faces but many more face and non-face photos. We first train a gray-scale face classifier using FaceScrub images (91,712 faces, 630 identities) as positive examples and random ImageNet samples (91,712 images, 1000 classes) as negative examples. We then train a black-white Mooney face classifier by fine-tuning the grayscale model on both thresholded face/nonface images and a small training set of Mooney face images. We assume that any non-face images, thresholded at any level, positive or negative, should be classified as non-faces, whereas face images should be classified as Mooney faces only in positive but not in negative contrast. Such instances are generated in balanced mini-batches to train our Mooney face classifier.

Given a grayscale face image, we generate a set of two-tone versions by smoothing and thresholding the image with a range of parameters. The highest scoring image predicted by our Mooney face classifier becomes the Mooney face. If the highest score is below probability 0.5, we reject the Mooney face and declare the grayscale image unfit as a Mooney face candidate. This approach yields a large-scale Mooney face dataset, potentially useful for future computer vision, psychophysics, and neurophysiology experiments.

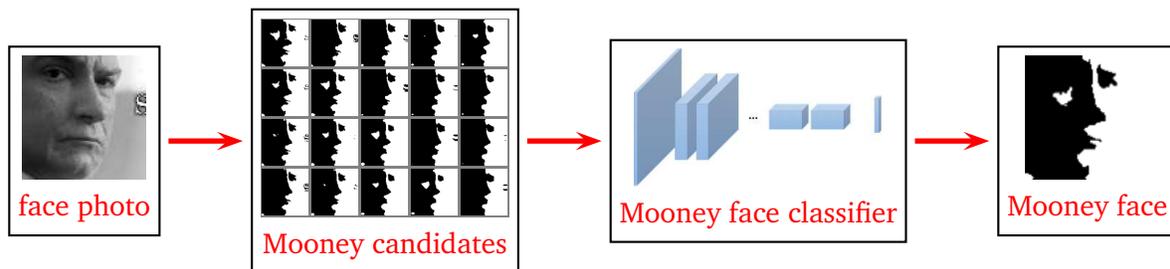


Figure 1: Our algorithm generates a Mooney face from a photo automatically. Given a grayscale image, we generate a set of black/white versions by smoothing and thresholding, each of which is evaluated by our Mooney classifier and the highest scorer is the final Mooney face. Our Mooney face classifier is a CNN (GoogleNet) trained from only 115 Mooney faces but hundreds of thousands of photos.

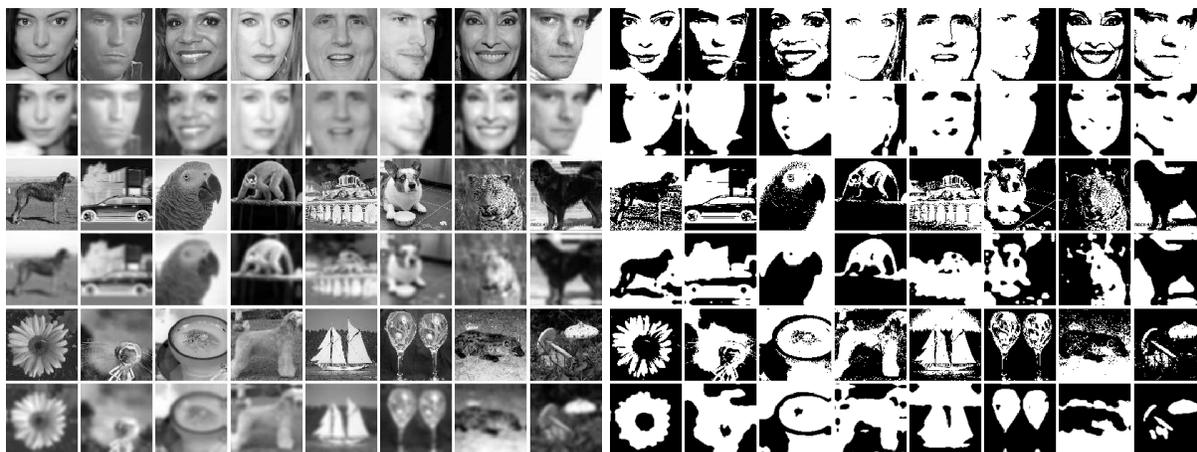


Figure 2: Binarization of grayscale photos and their smoothed versions. Most are not Mooney.



Figure 3: Grayscale images and their best binarized versions ranked high by our Mooney classifier.



Figure 4: Grayscale images and their best binarized versions ranked low by our Mooney classifier.