Restricted Coloring using Saliency-based Image Segmentation

BROWN

Marc ten Bosch

Seong Jae Lee

Overview

We describe two different methods of restricting the number of colors used to draw a certain image while preserving detail in salient regions.

Inspired by the method described by DeCarlo and Santella, we segment images based on saliency data, assigning larger regions to less salient areas of the image, and filling each region with its mean color. Unlike their method, which measures saliency using an eye-tracker, we define saliency using a hand-drawn saliency map.

Original image and saliency map.

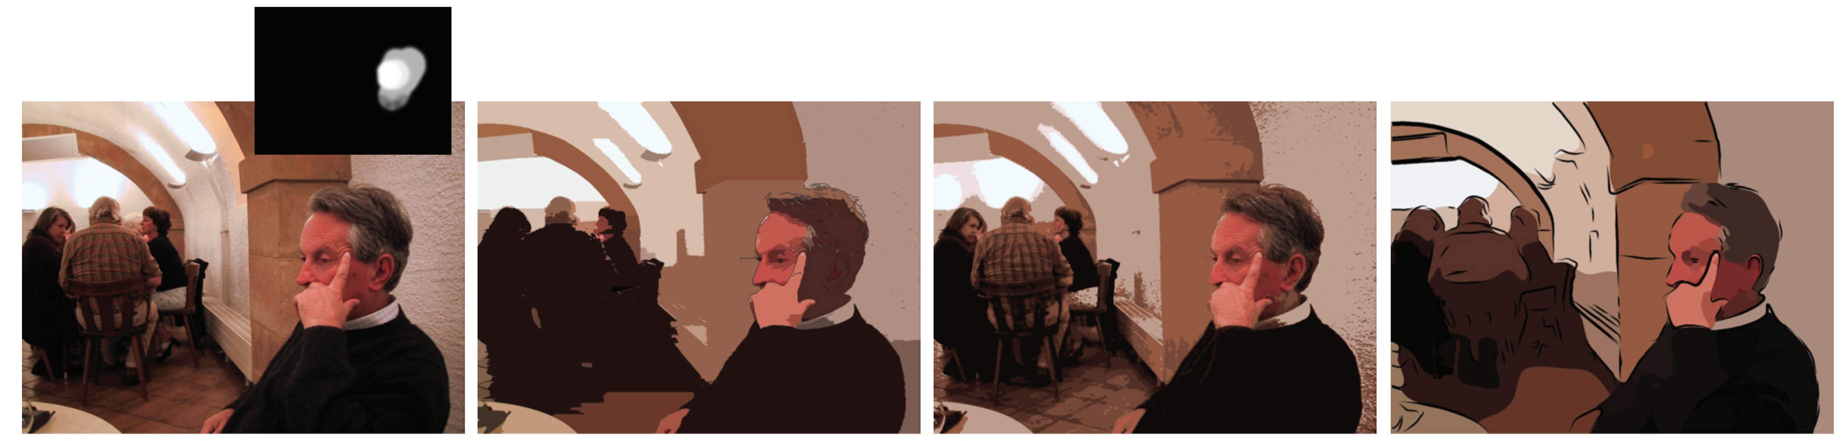

Left to right: Orignal Image/our saliency map (DeCarlo uses eye-tracker data); CCL method; KMC method; DeCarlo method. DeCarlo and Santella apply Mean Shift segmentation at multiple scales and create a segmentation tree by assigning to each region a parent region at a coarser scale. This tree is pruned based on saliency. In contrast, our methods operates directly on the finer scale.

The saliency map can also be used to overlay edges computed at different scale on top of the image. A Gaussian pyramid is constructed, and edges are computed at every scale using a Sobel filter. For each pixel in the final image we select a particular scale based on saliency, with less salient regions corresponding to coarser scale.

Running Recursive 2-means on an image sequence with hand-drawn saliency. Maximum recusion depth is 2 for non-salient regions and around 6 for salient regions.

Connected Components Labeling Algorithm

In our second method, to allow for more cohesive regions, we introduce a new image segmentation algorithm based on connected components labeling (CCL) of images, which assigns a different ID to every region of connected pixels which have the same value. We loosen this definition such that we assign a pixel to a region if its distance in color space from the mean of the current region is under a certain threshold. This threshold can be used to control the number of regions created. The saliency map can also be used to vary the threshold across the image.

The method can be improved by overlaying the result of applying a Canny edge detector over the image. The Canny edges increase contrast along important edges in the image, preventing regions from spanning across them.

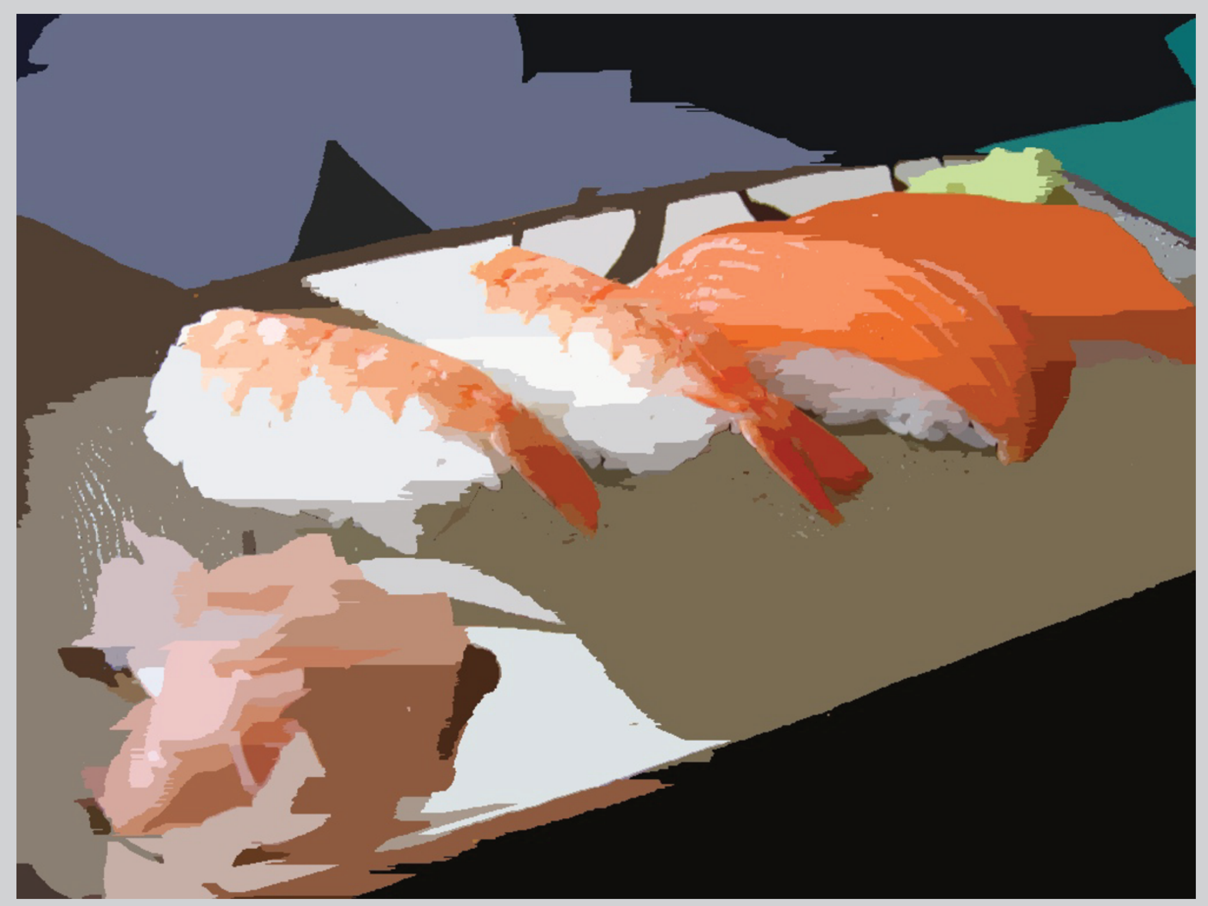

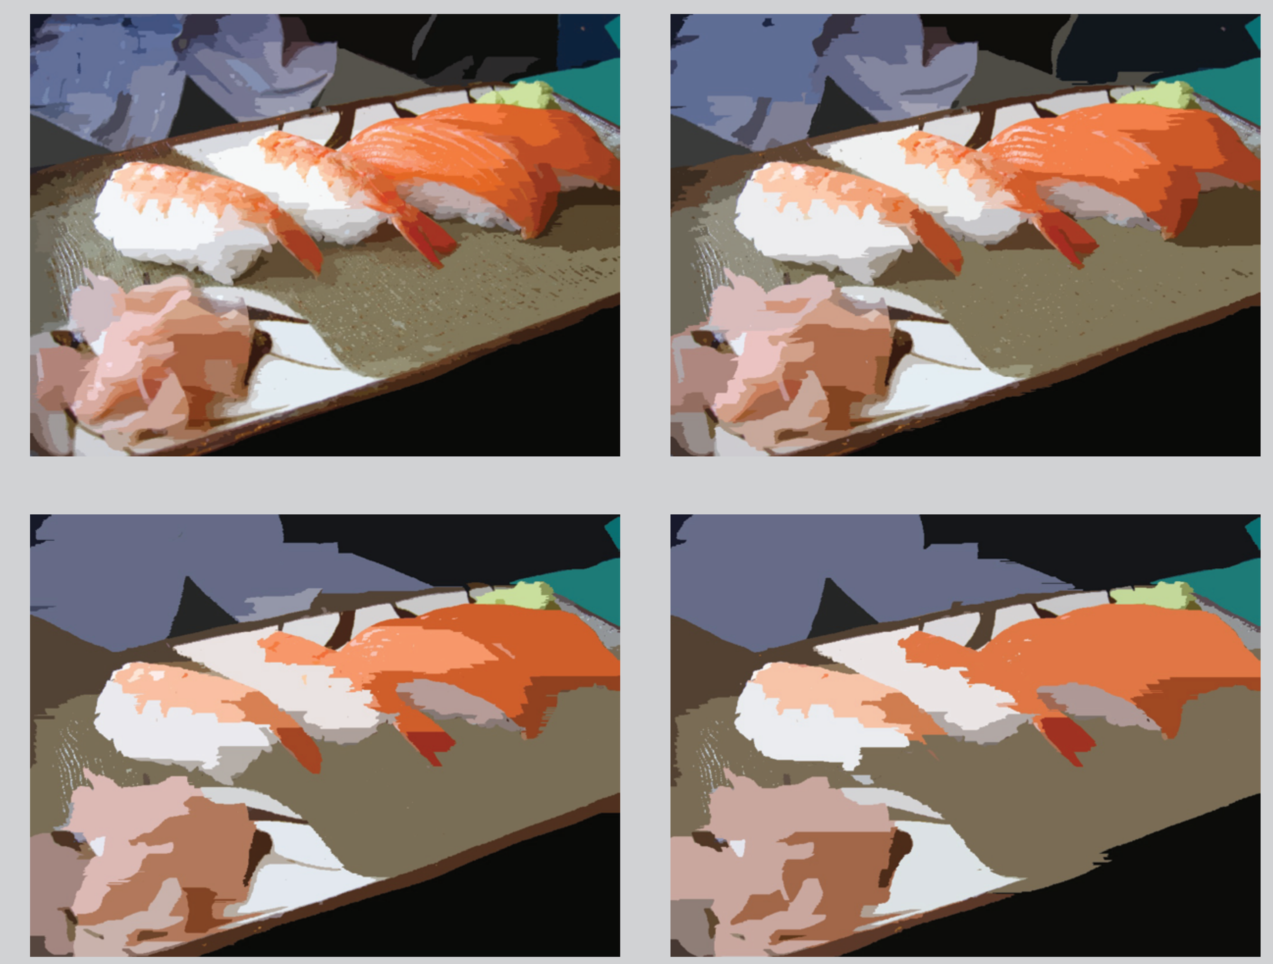

Right: images generated using increasing thresholds. Above: image with threshold varied by saliency map.

K-Means with Connectivity Constraint Algorithm

In our first method, we recursively segment the images using a K-Means with Connectivity Constraint (KMC) algorithm. KMC tends to assign pixels within a close distance to the same group by defining the distance metric as the weighted sum of the distances in LUV color space, pixel location and saliency. Beginning from one parent region, we run KMC recursively on each region if its mean saliency is less than a threshold and the maximum depth has not been reached (K more segments are created and pushed into the segment list). We generally set K between 2 and 4.

In order to run K-Means on image sequences, we can segment the entire volume by considering a distance in time as well. To reduce computation we calculate the means on a scaled down volume, and later apply them to the full size sequence.

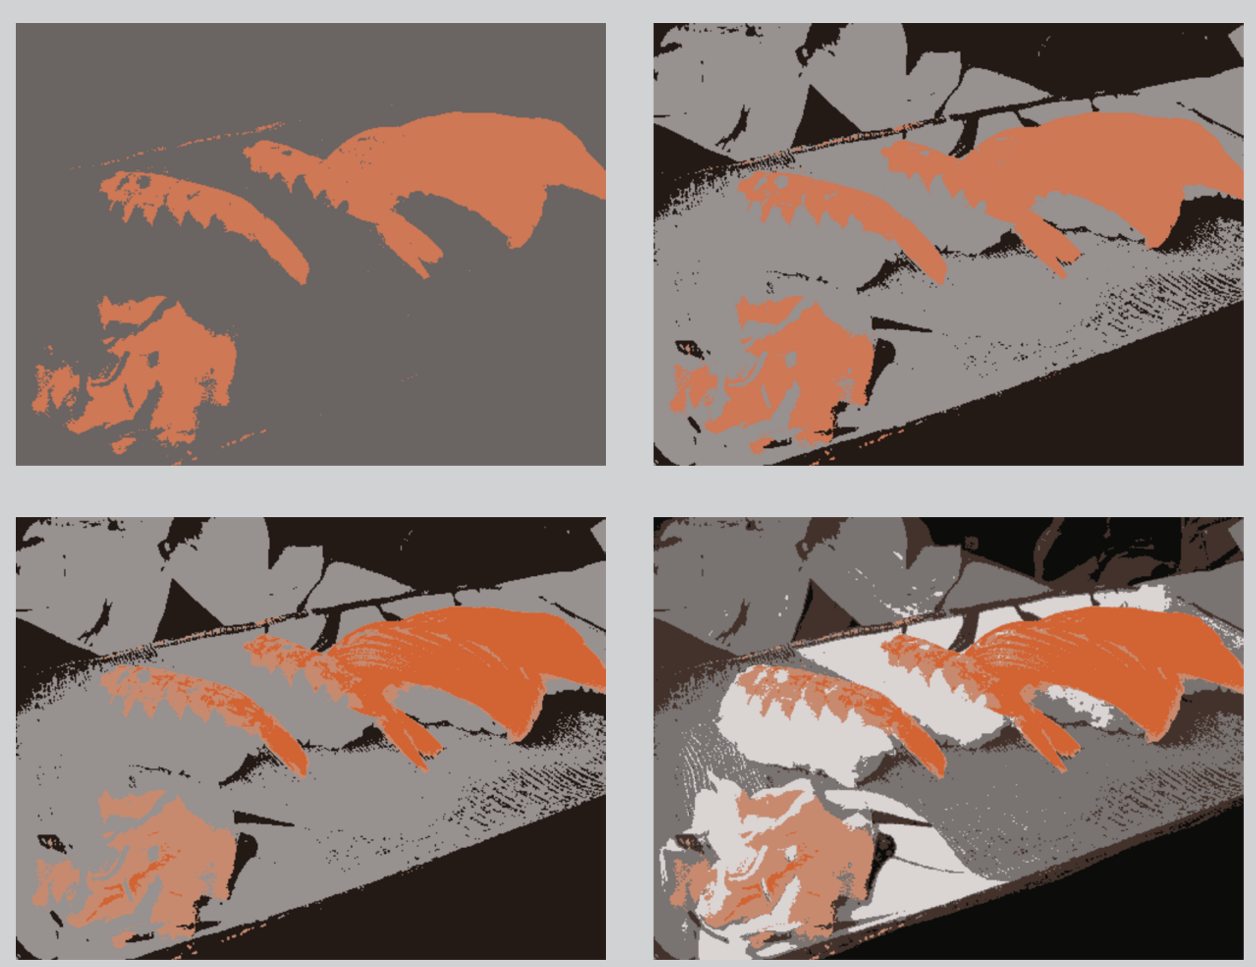

Left: images generated using increasing recursion depths based on saliency map. Above: final image. For all images K=2.

Overlaying Canny edges on top of the image before applying the CCL segementation algorithm increases contrast, in this case insuring the tree branch is well separated from the leaves.