Supplemental Material : Understanding and Mapping Natural Beauty

Scott Workman[1] Richard Souvenir[2] Nathan Jacobs[1]
scott@cs.uky.edu souvenir@temple.edu jacobs@cs.uky.edu
[1]University of Kentucky [2]Temple University

This document contains additional details and experiments related to our methods.

1. Exploring Image Scenicness

1.1. Uncertainty in Scenicness Scores

Figure 1 shows a scatter plot of Shannon entropy, computed per image from ratings, versus scenicness, for all images in the ScenicOrNot (SoN) dataset. Each dot represents an image and is colored by the average beauty score. Unsurprisingly, the images that are rated the least and most scenic have higher consistency in their scenicness ratings.

1.2. Image Captions

Figure 2 shows a word cloud of the most frequent 100 extracted terms from captions and titles associated with non-scenic images in SoN (average rating below 3.0), where the size of the word represents the relative frequency. Example terms that are negatively correlated with scenicness include “road”, “lane”, “house”, and “railway”.

2. Experiments

2.1. Scenicness-Aware Image Cropping

Figure 3 shows additional examples of applying our methods for scenicness-aware image processing. In each image, the inset shows the change in scenicness from the full image to the cropped image.

2.2. Mapping Scenicness

Our method is applicable for generating maps of scenicness at widely varying spatial scales. Figure 4 shows three examples of fine-grained high-resolution maps of scenicness from regions of different sizes around Great Britain. Figure 4 (top) shows a stretch of coast near Holywell, a coastal village in north Cornwall, England. Figure 4 (middle) shows a map centered over Caehopkin, a village in Powys, Wales. Caehopkin sits between Abercraf and Coelbren in the Swansea Valley and lies on the border of the Brecon Beacons National Park to the north. Figure 4 (bottom) shows a map of Greater London. Figure 5 and Figure 6 show additional maps of scenicness computed using our method, alongside several baseline approaches. In all cases, *Cross-View Hybrid (CVH)* combines overhead imagery and nearby ground-level images to more accurately identify scenic and non-scenic locations.

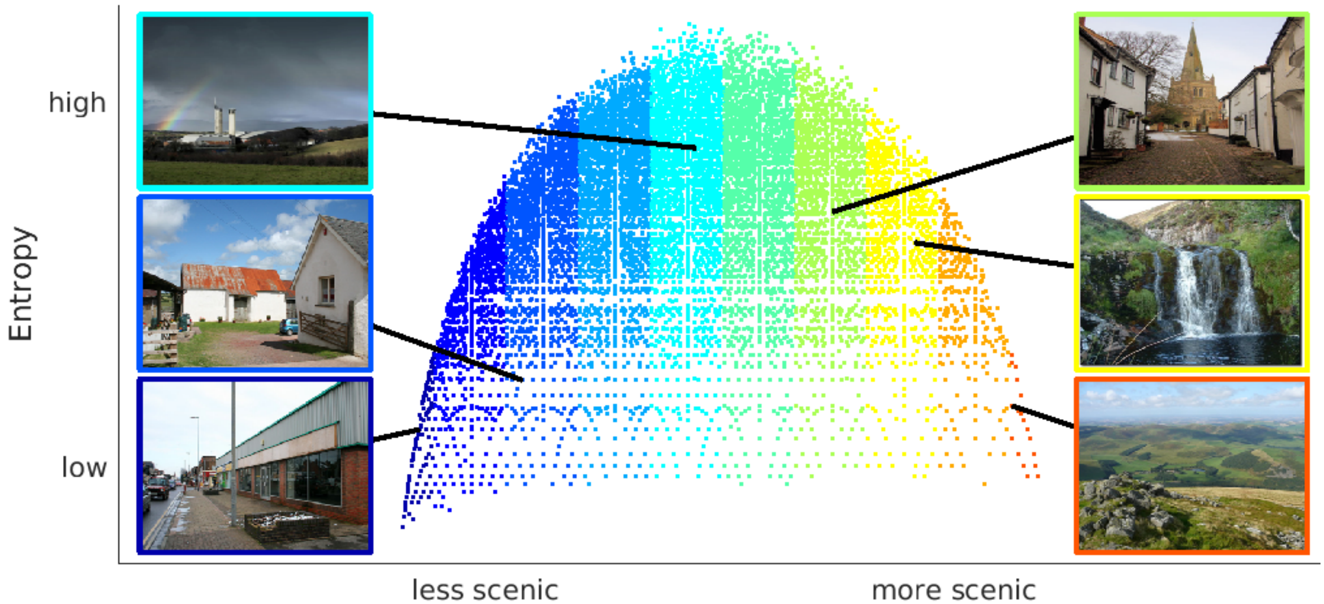

Figure 1: A scatter plot of per-image entropy vs. scenicness, computed from human ratings, for each image in the ScenicOrNot dataset.

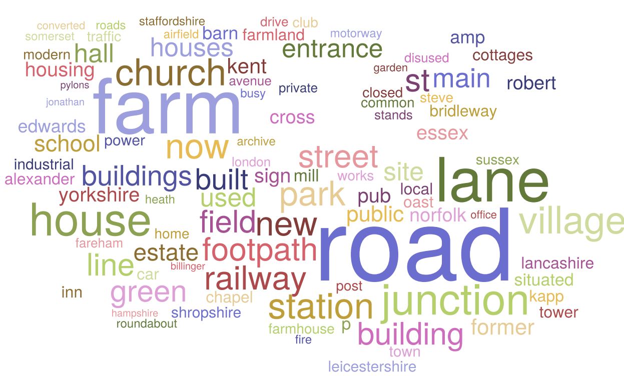

Figure 2: The word cloud depicts the relative frequency of title and caption terms found in non-scenic images from the ScenicOrNot dataset.

Figure 3: Scenicness-aware image processing. For each image, the green bounding box shows the image crop that maximizes scenicness. The predicted scenicness scores for both the entire image and the cropped region are shown in the inset.

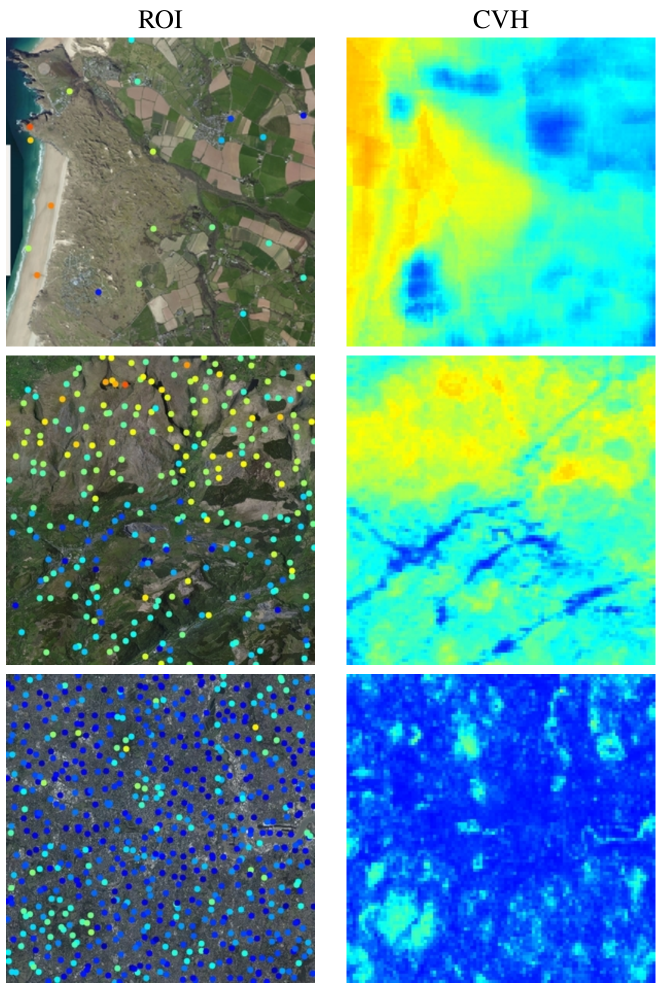

Figure 4: Varying spatial resolutions. The first column shows an overhead image where dots correspond to geo-tagged ground-level imagery, colored by average scenicness rating (warmer colors correspond to more scenic images). The second column shows a false-color image that reflects the average scenicness predicted by our method.

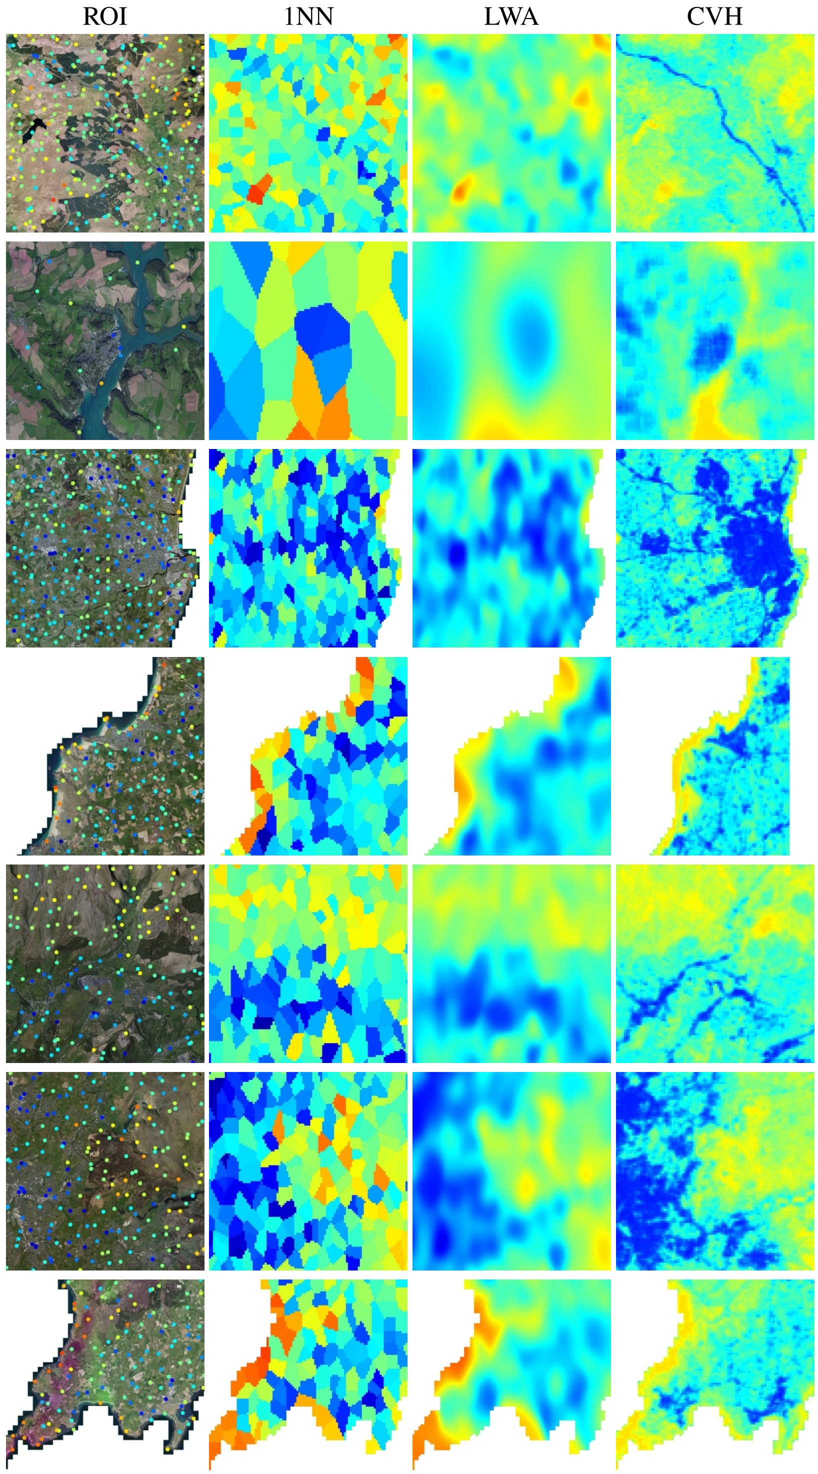

Figure 5: Scenicness maps. The first column shows an overhead image where dots correspond to geotagged ground-level imagery, colored by average scenicness rating (warmer colors correspond to more scenic images). The remaining columns show false-color images that reflect the average scenicness predicted by each method.

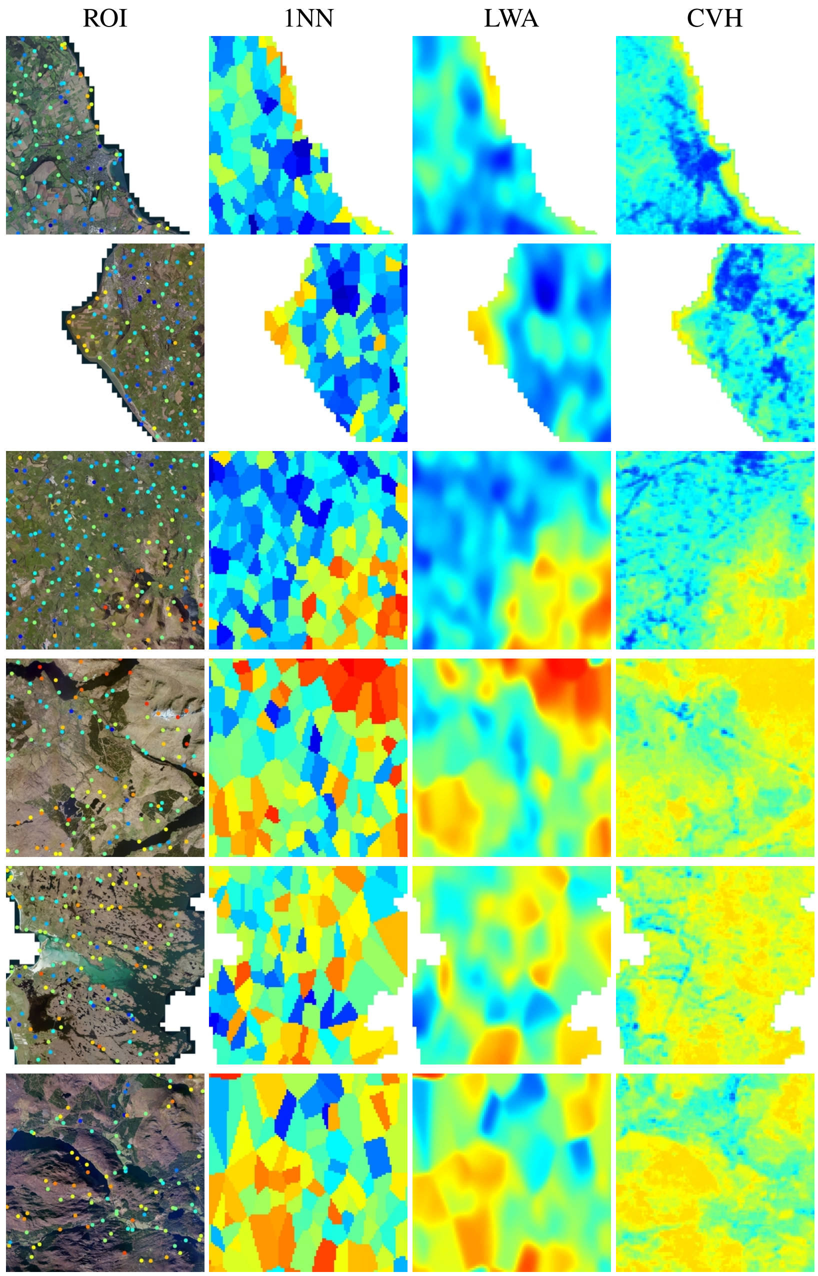

Figure 6: Scenicness maps. The first column shows an overhead image where dots correspond to geotagged ground-level imagery, colored by average scenicness rating (warmer colors correspond to more scenic images). The remaining columns show false-color images that reflect the average scenicness predicted by each method.