SeeTheSeams: Localized Detection of Seam Carving based Image Forgery in Satellite Imagery
– Supplementary Material –

Chandrakanth Gudavalli[†,1], Erik Rosten[†,1], Lakshmanan Nataraj [1], Shivkumar Chandrasekaran [1, 2], and B. S. Manjunath[1,2]

[†]Equal Contributors [1]Mayachitra, Inc.
[2]Electrical and Computer Engineering Department, UC Santa Barbara
Santa Barbara, California, USA

1. Overview

The supplementary material includes details that could not be included in the main paper due to space limitations. In Section 2, we present the experimental setup that is used to train and test the proposed models . In Section 3, we present different seam carving techniques that we used in main paper to test the generalizability capabilities of trained models across several seam carving techniques. In Section 4, we present a number of examples of object removal and displacement in satellite imagery from all three datasets covered in the main paper and examples of images generated using different seam carving methods.

2. Experiments

2.1. Experimental Setup

In this subsection, we cover the experimental setup for training our models. Seam removal and seam insertion pixelwise classifiers for the xView dataset are trained using the Adam optimizer with a learning rate of 0.0003, momentum coefficients $\beta_1 = 0.9$ and $\beta_2 = 0.999$ and a numerical stability constant of $\epsilon = 10^{-6}$. The loss function used is a pixelwise mean squared error. We find that using binary cross entropy as the loss function produces comparable results, but the mean squared error leads to smoother convergence. We use a batch size of 8 where each batch is sampled randomly from the training set without replacement for each epoch. We train for 30 epochs and select the model with the highest validation set accuracy, reducing the learning rate by a factor of 0.2 if there is no validation accuracy improvement for 5 epochs. For training on xBD and Orbview-3 datasets, we keep all hyperparameters the same except for a learning rate adjustment to 0.0001.

Stage 2 binary classifiers are trained using identical hyperparameters as above, except that we use a cross entropy loss, a learning rate of 0.0001, and train for 35 epochs. These hyperparameters are the same for the stage 2 model used in Table 5 and Table 6 of the main paper.

3. Seam Carving Methods

Seam carving is a content-aware image retargetting algorithm proposed in 2007 [2]. A seam is defined as an optimal 8-connected path of pixels from top to bottom or left to right. The optimality of a seam is determined by assigning an energy value to each pixel in a given image. One possible energy function proposed in [2] is the gradient magnitude, also termed as *backward energy*.

$$E(\boldsymbol{I}) = \sqrt{\frac{\partial^2 \boldsymbol{I}}{\partial x^2} + \frac{\partial^2 \boldsymbol{I}}{\partial y^2}} \tag{1}$$

where $\boldsymbol{I}(x, y)$ is an image indexable by row x and column y, and $E(\boldsymbol{I})$ is the energy map of $\boldsymbol{I}$. Then, the optimal vertical seam for an m by n image can be defined as

$$\begin{aligned} s^x &= \{s_i^x\}_{i=1}^m = \{i, x(i)\}_{i=1}^m \\ &\text{s.t. } \forall i, \ \ |x(i) - x(i-1) \leq 1| \end{aligned} \tag{2}$$

The optimal seam is then the seam that minimizes the energy function over all possible seams,

$$s^* = \min_{\boldsymbol{s}} E(\boldsymbol{s}) = \min_{\boldsymbol{s}} \sum_{i=1}^{m} E(\boldsymbol{I}(\boldsymbol{s})) \tag{3}$$

This optimal vertical seam can be found through dynamic programming, computing a minimum cumulative energy matrix $M(x, y)$ for all possible seams through (x, y) by traversing from the second row to the last row:

$$M(x,y) = e(x,y) + \min \begin{cases} M(x-1,y-1) \\ M(x-1,y) \\ M(x-1,y+1) \end{cases}$$

where $e(x,y)$ is an optional additional energy measure.

Once the cumulative energy matrix has been computed, the minimum value in the last row of M gives the index of the end of the optimal seam. The rest of the seam can be found by backtracking through M.

Once we find the optimal seam, we can simply remove it to reduce the width or height of an image. To increase the dimensions of an image, we take the average of the optimal seam and a seam to the immediate right and insert it at the location of the optimal seam. By successively removing and inserting optimal seams, images can be resized without modifying many of the original pixel values of the image. These removed seams have low energy according to Eqn. 1, and often correspond to smooth regions in the image, visually preserving image content and structure.

Seam carving can be applied to remove objects by setting energy map values from a user-provided mask to a "low" value, and successively removing seams until the object is removed. The user can insert seams to restore the original image size, as well as preserve objects through a mask that tags regions to have a "high" energy. Figure 2 in the main paper is one such example of this *object removal* strategy and perceptually it is very difficult for a human to deduce that objects have been removed in this image. Seam carving can also be utilized for *object displacement* in a satellite image by marking the region on one side of an object for seam removal and then inserting seams on the other side of the object after the seams have been removed (Figure 4 in main paper is one such example). In this way, objects can be "shifted" to a different GPS location with very little perceptual difference.

3.1. Other Seam Carving Methods

There have been several modifications to the original seam carving algorithm,. One such modification was proposed in 2008 [4] that considers the energy introduced into the resized image by bringing previously non-adjacent pixels together by removing seams. This formulation, termed as *forward energy*, computes the energy map by looking at the differences in pixel values depending on the direction of the potential seam, defining three possible cases: up, up and to the left, and up and to the right.

$$\begin{aligned} C_L &= |\boldsymbol{I}(x,y+1) - \boldsymbol{I}(x,y-1)| + |\boldsymbol{I}(x-1,y) - \boldsymbol{I}(x,y-1)| \\ C_U &= |\boldsymbol{I}(x,y+1) - \boldsymbol{I}(x,y-1)| \\ C_R &= |\boldsymbol{I}(x,y+1) - \boldsymbol{I}(x,y-1)| + |\boldsymbol{I}(x-1,y) - \boldsymbol{I}(x,y+1)| \end{aligned}$$

Then, the cumulative energy matrix can be updated as:

$$M(x,y) = e(x,y) + \min \begin{cases} M(x-1,y-1) + \ C_L(x,y) \\ M(x-1,y) + \ C_U(x,y) \\ M(x-1,y+1) + \ C_R(x,y) \end{cases}$$

where $e(x,y)$ is an optional additional energy measure.

Another seam carving variation proposes using an importance map based on *salient region detection* [1]. In this particular formulation, salient regions are uniformly identified considering global contrast as opposed to local edges in the two original methods outlined above. This prevents the necessity to recompute the energy map after each optimal seam is found. The saliency map used in place of the energy map also incorporates color information, which the two methods above omit. To calculate the saliency map, the image is converted into Lab color space. Then, the final map is the Euclidean distance between the average pixel vector and a Gaussian blurred version, approximated by a 5 x 5 binomial filter (both in Lab space):

$$E_{Lab}(x,y) = ||\boldsymbol{I}_\mu - \boldsymbol{I}_{Gauss}(x,y)|| \qquad (4)$$

The final variation of seam carving that we explore is called *seam merging* [3]. This method merges a two-pixel-width seam element into one new pixel during image reduction and inserts a new pixel between the two pixels during image enlargement. This algorithm utilizes importance and structure energies to define seam optimality, as well as an additional energy term that suppresses artifacts generated by excessive reduction or enlargement from repeated merging or inserting.

4. Visual Examples

Besides content aware image resizing, seam carving can be used to remove or displace objects in a given image. In Figure 1 and Figure 2, we present the sample manipulations that one can do using seam carving and visualize the localization results of stage 1 pixelwise classifiers on manipulated images.

Figure 1 has five examples illustrating the application of seam carving to remove objects in satellite imagery while making sure that the seam carved satellite image looks authentic to the human eye and retains its original dimensions. In Figure 1a, truck under the removal mask (red) is taken out and seams are inserted to restore the original image, all while leaving pixels in the protective mask (green) are left undisturbed. This is achieved by setting the energy map values at the removal mask locations to a low energy value, forcing seam carving algorithm to pass through. When inserting seams to restore the original dimensions, pixels at the locations of the protective mask are set to a high energy value, ensuring that the seam carving algorithm ignores them. Figure 1b has an example in which an excavator under the red mask is removed from the image. Similar examples are shown in Figures 1c, 1d, and 1e, where pixels under the red colored mask are removed while pixels under the green colored mask (if present) are left undisturbed.

Figure 2 has five examples illustrating seam carving based manipulations to displace objects in a given image while retaining the visual authenticity and size of the original image. In Figure 2a, the white colored building at the center of the image is displaced by 50 pixels to left. This is achieved by forcing the seam carving algorithm to remove 50 seams from left side of the object and insert back same number of seams to right using removal and protective masks. In Figure 2b, the road is displaced by 50 pixels to left, whereas in Figure 2c, road is displaced by 50 pixels to right. An excavator is moved to right by 50 pixels in Figure 2d, and aeroplanes are moved to left by 40 pixels in Figure 2e.

Figure 3 visualizes the distribution of removed seams in a given image that is seam carved using different seam carving algorithms described in Section 3. It can be observed that even though the seams are removed using different seam carving algorithms, our seam removal detector in stage 1 is able to predict the locations of the removed seams.

References

[1] R. Achanta and S. Süsstrunk. Saliency detection for content-aware image resizing. In *2009 16th IEEE International Conference on Image Processing (ICIP)*, pages 1005–1008, 2009. 2

[2] Shai Avidan and Ariel Shamir. Seam carving for content-aware image resizing. *ACM Trans. Graph.*, 26(3):10, 2007. 1

[3] Kazu Mishiba and Masaaki Ikehara. Image resizing using improved seam merging. In *2012 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1261–1264. IEEE, 2012. 2

[4] Michael Rubinstein, Ariel Shamir, and Shai Avidan. Improved seam carving for video retargeting. *ACM Transactions on Graphics (SIGGRAPH)*, 27(3):1–9, 2008. 2

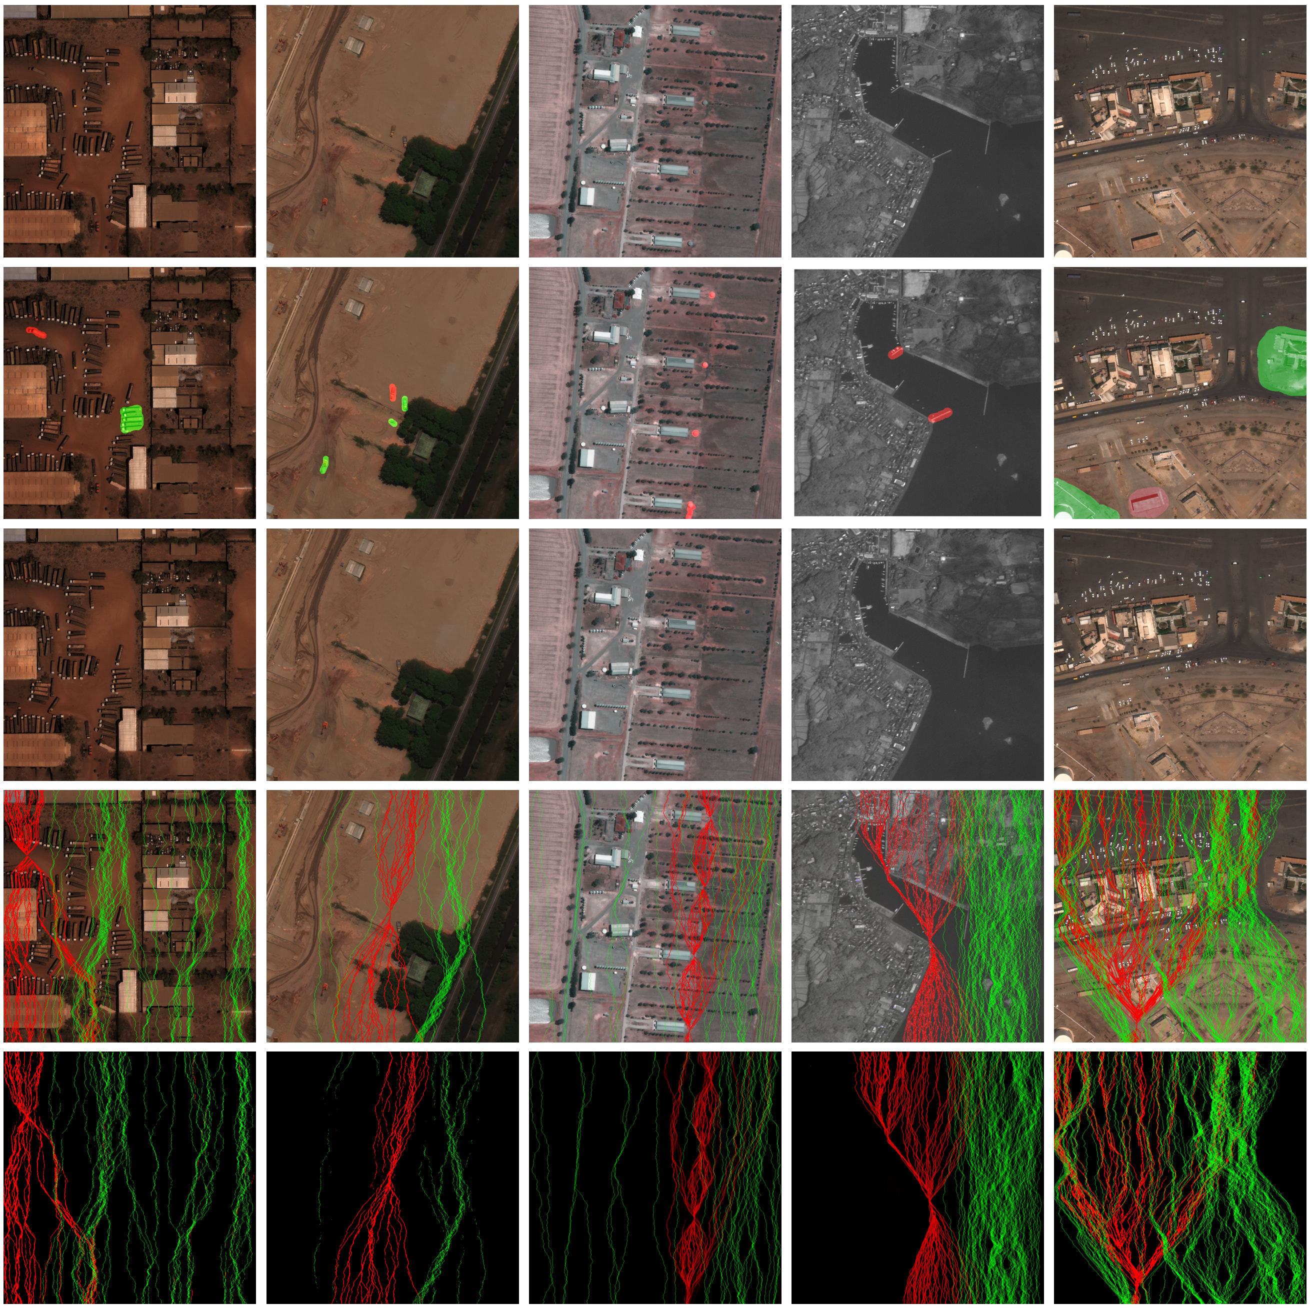

(a) (b) (c) (d) (e)

Figure 1. Object Removal Examples. From top to bottom: Pristine satellite image; Pristine image overlaid on removal (red) and protective (green) masks; Seam carved image with objects removed while retaining the original image size; Ground truth seam mask, with removed (red) and inserted (green) seams, overlaid on seam carved image; Predicted seam mask generated by stage1 seam removal detector (red) and seam insertion detector (green).

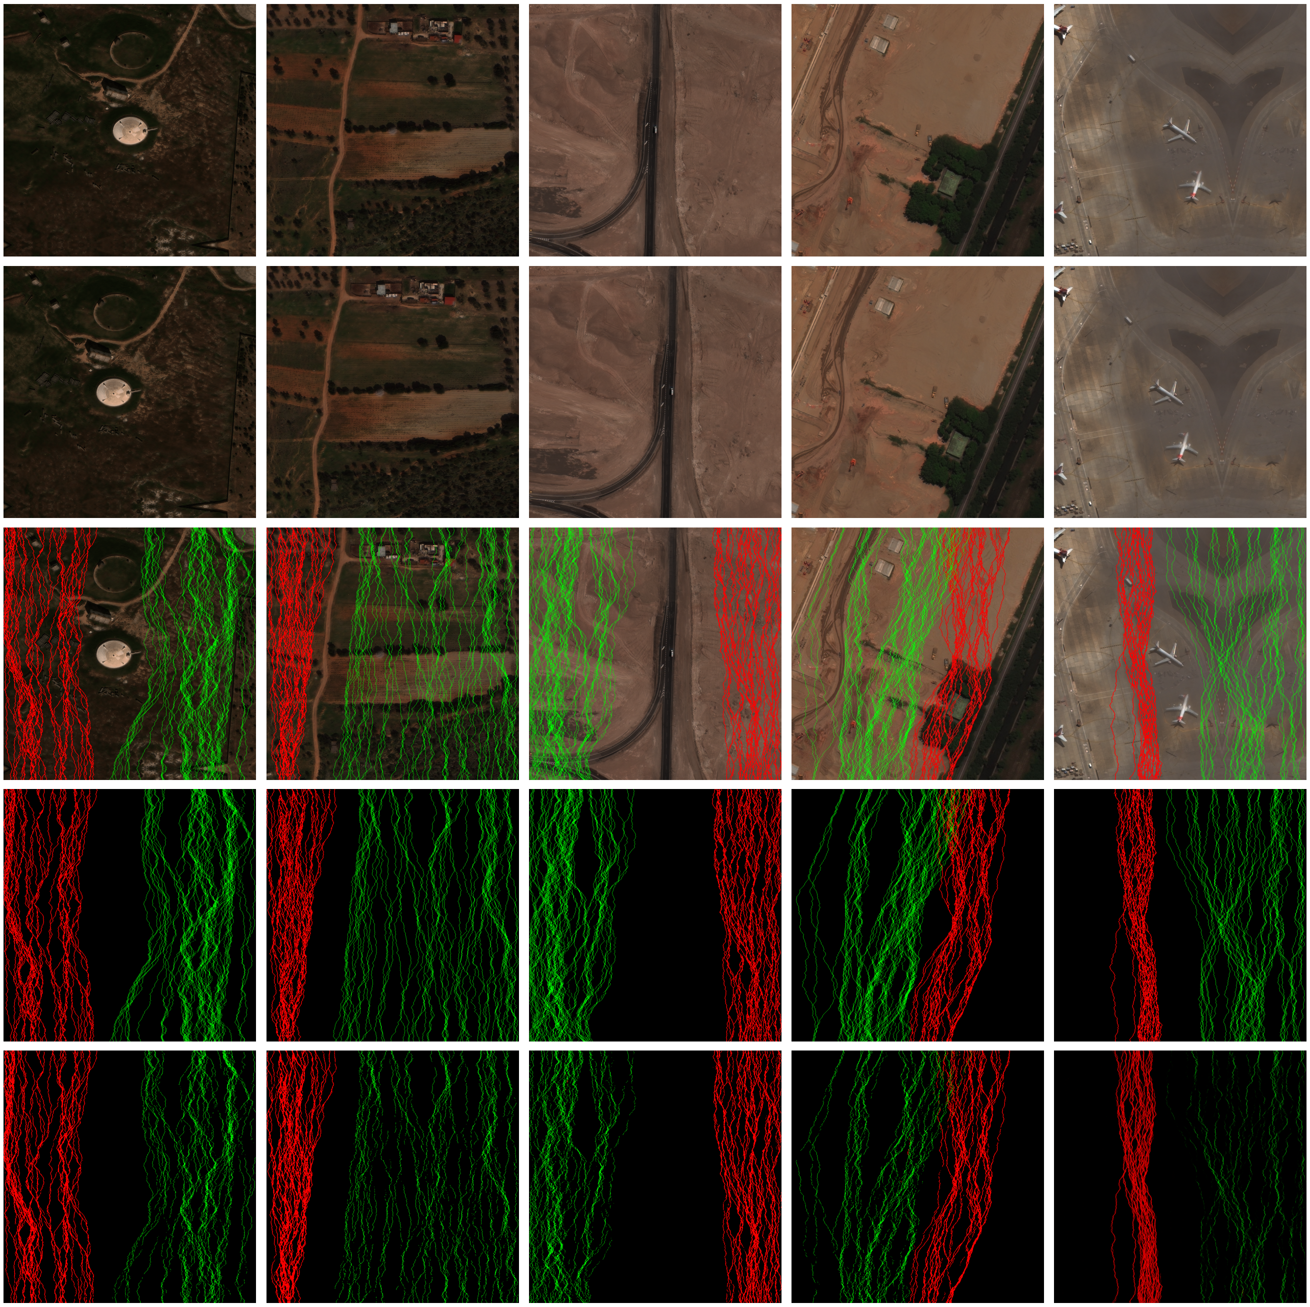

(a) (b) (c) (d) (e)

Figure 2. Object Dislocation Examples. From top to bottom: Pristine satellite image; Seam carved image with objects dislocated while retaining the original image size; Ground truth seam mask, with removed (red) and inserted (green) seams, overlaid on the seam carved image; Ground truth seam mask; Predicted seam mask generated by a stage 1 seam removal detector (red) and seam insertion detector (green).

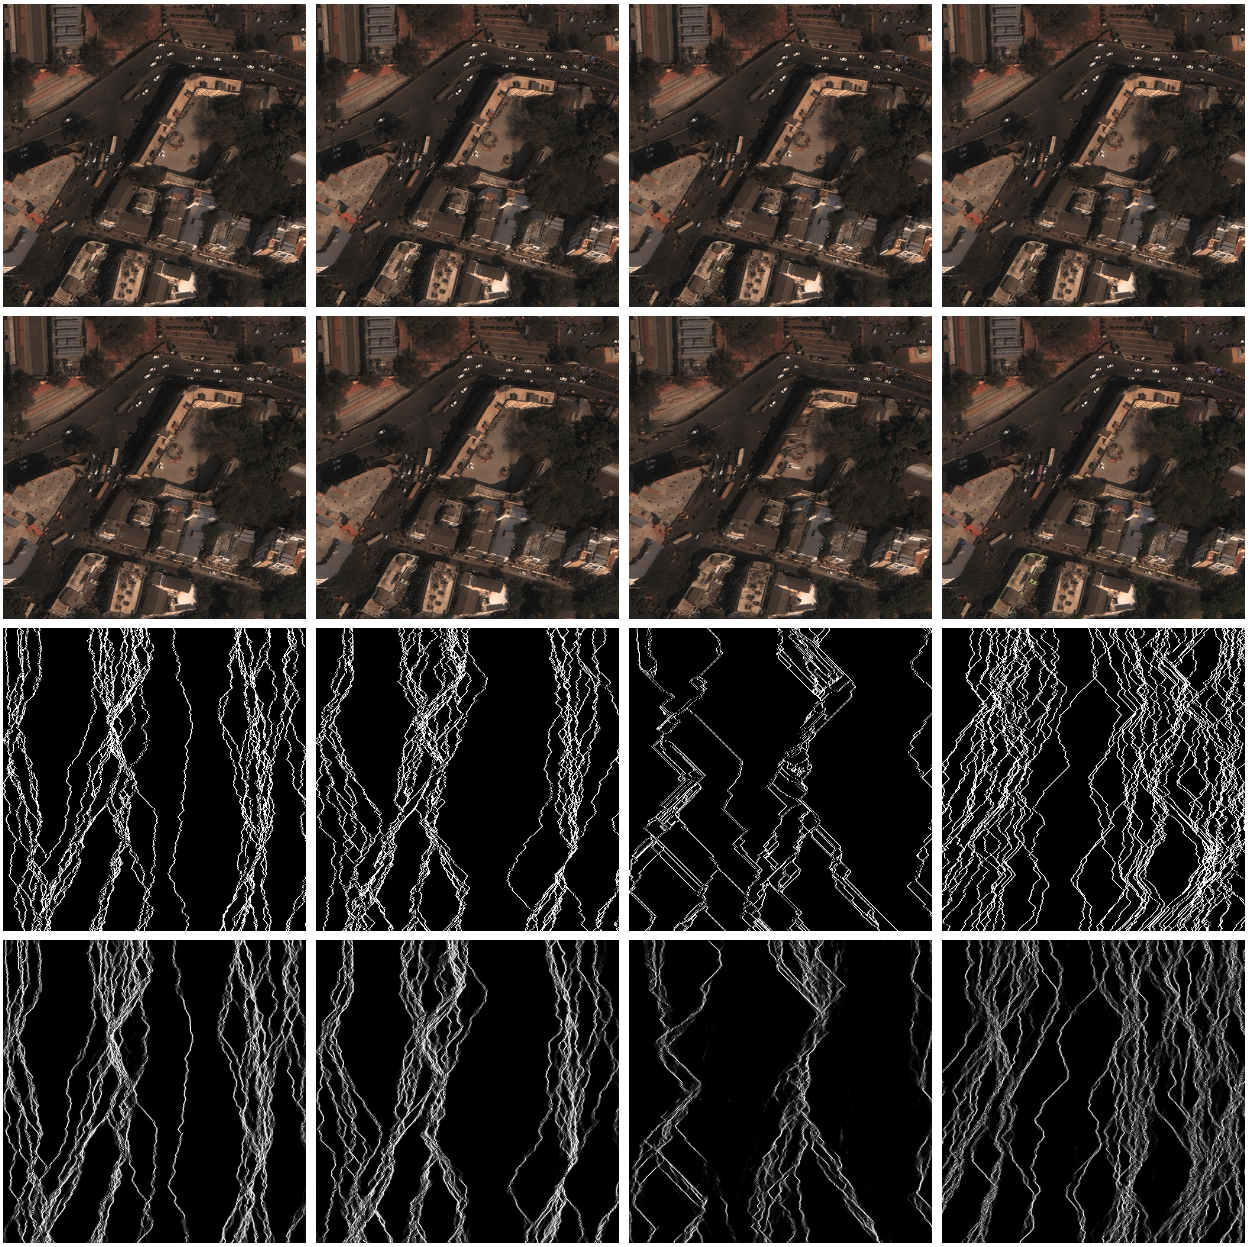

Figure 3. Seam Carving Ground Truth and Prediction Comparison using Various Seam Carving Algorithms. From left to right: Forward Energy, Backward Energy, Saliency Maps, Seam Merging. From top to bottom: Original image, seam carved image, ground truth seam removal mask, predicted seam removal mask