Sky is not the limit : Semantic Aware Sky Replacement

[http://www.eecs.harvard.edu/~kalyans/research/skyreplace/SkyReplacement_SIG16.pdf]

DIP Project

Anjali Shenoy (201401114)
Saumya Rawat (201401110)

What is Colour Transfer?

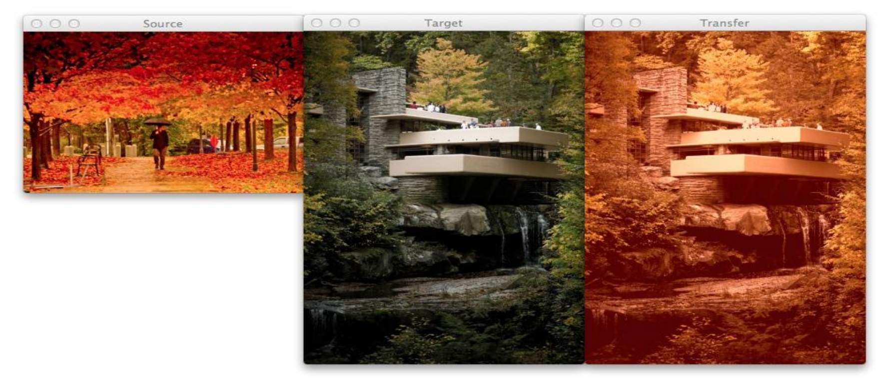

- Changes the appearance of a target image/video according to the colour pattern of a source image.
- Applications include enhancing photos, movie post-production, artistic design, hallucinating a new image of an existing scene at a different time of the day
- The goal of a colour transfer algorithm should be to keep the scene of the target image and accurately apply all the dominant colour styles of the source.

Problem Statement

- Often in photographs of sceneries, skies form an important part of the aesthetic and yet are less interesting due to the time of photographing. To get the best shot professional photographers have to use sophisticated tools with painstaking efforts to get effects that normal users can't achieve.
- In the given paper, an automatic background replacement algorithm is proposed that can generate realistic, artifact-free images with a diverse styles of skies. The algorithm exploits the idea that common objects across images will have similar colour based properties.
- Given an input image, a set of possible skies are proposed for sky replacement based on higher number of matching labels found in the database. The output image consists of a sky chosen from the retrieved options and foreground objects in the modified colour space after applying an appearance transfer method, developed to match statistics locally and semantically.

Pipeline

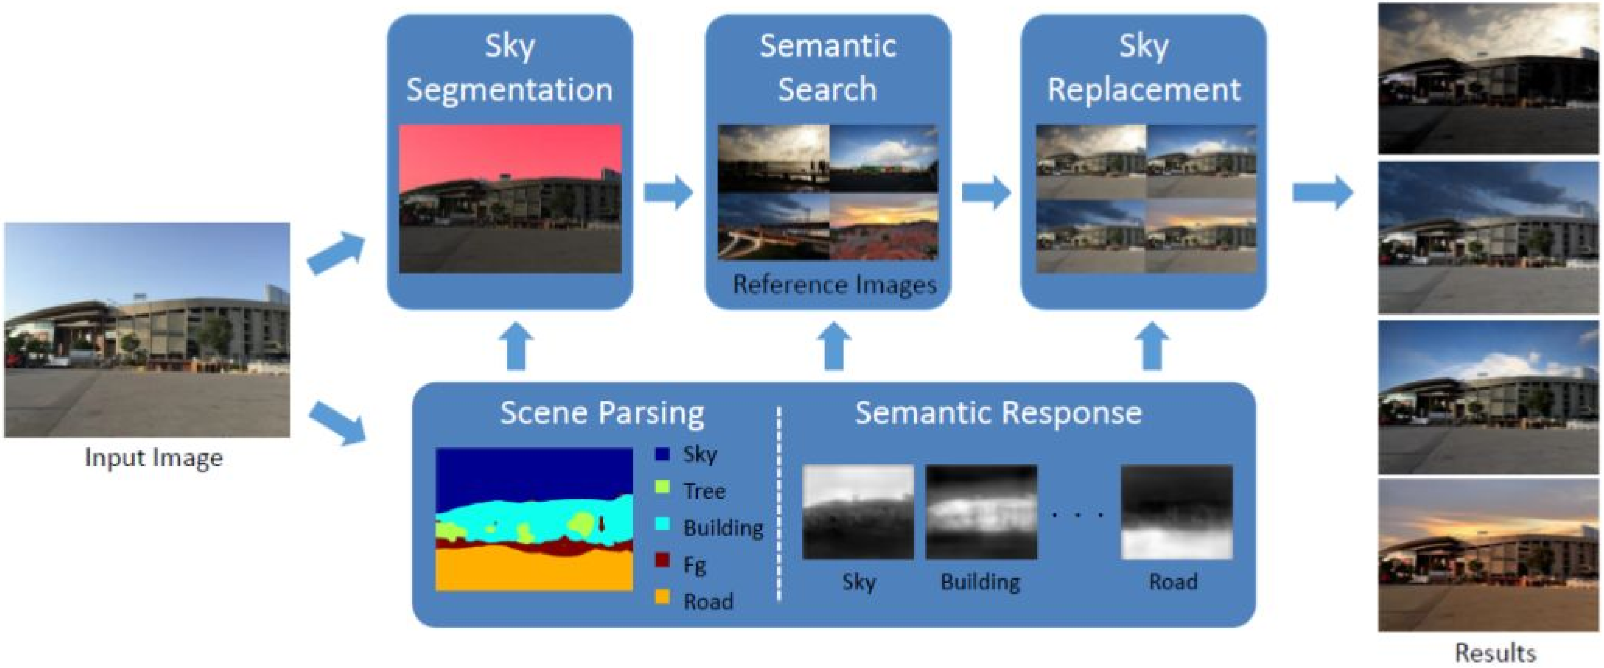

Input Image
Sky Segmentation
Semantic Search
Reference Images
Sky Replacement
Results
Scene Parsing
Sky
Tree
Building
Fg
Road
Semantic Response
Sky
Building
Road

Sky segmentation

Input Image → Fully Convolutional Neural network → Pixel-wise probability map of images→ Labelled image

To refine this labelled output, the paper proposes a graph cut based energy minimization.

$$E(X) = \lambda_1 \sum_i U_c(x_i) + \lambda_2 \sum_i U_t(x_i) + \lambda_3 \sum_i U_f(x_i) + \lambda_4 \sum_{(i,j)\in\mathcal{E}} V(x_i, x_j),$$

U_c -> color potential
U_t -> texture potential
for the sky and non-sky label costs (obtained from the learned online classifier)
U_f -> FCN output
V -> the pairwise potential for smoothness

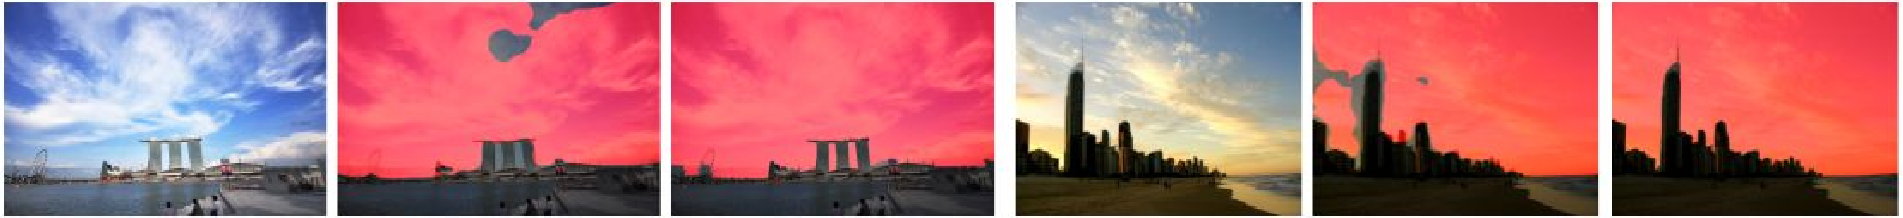

(a) Input image (b) FCN result (c) Fine segmentation (d) Input image (e) FCN result (f) Fine segmentation

Sky Search & Retrieval

- Descriptor computation :

FCN output of Image → Spatially pooling method, 9 grids of image $h^j = \frac{1}{m}\sum_{i=1}^{m} f_i^j$

+

Global histogram of image

- Pairwise distance calculation and retrieve top 4 closest skies.
- Another constraint on the retrieval is the aspect ratio and resolution of the pair, reference and input, to ensure that the two images are similar in terms of areas of sky.

Sky replacement

maximum rectangle enclosed in the reference sky region

(resized) ↓

the largest rectangular convex hull enclosing the sky region in the input image (A)

The sky pixels in the input image are given values from A ⟶ colour transfer.

Reference

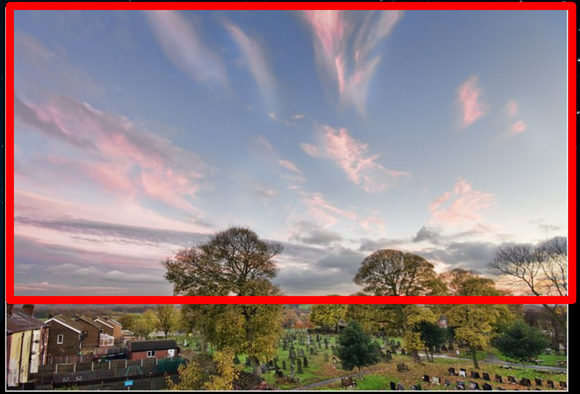

Input

Semantic aware transfer

- Compute luminance difference in sky and use this for regularisation of luminance across labels of input and reference simultaneously.
- For chrominance transfer, map the Gaussian statistics of Reference Image's chrominance to Input Image's chrominance across labels.

Special case:

If no label found in reference image, for a label input image, deviating from the paper, we resort to finding the regions neighbouring the label in Input image and try and see if this exists in reference too.

Results

Input

Selected

Retrieved Skies

Outputs

Input

Retrieval of semantically similar images

+

=

+

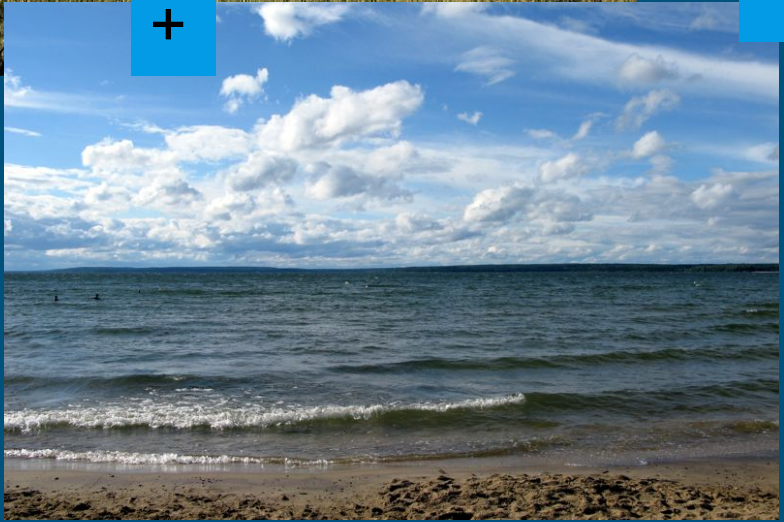

=

+

=

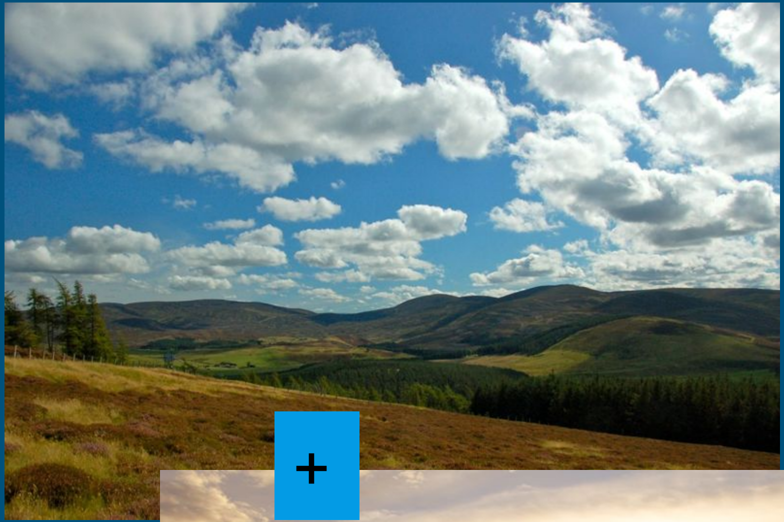

+

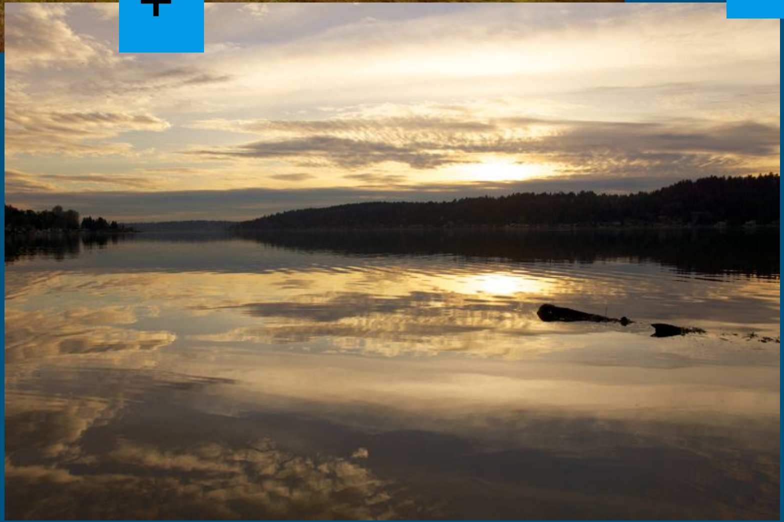

=

When few labels match

Limitations of the given approach

- Does not change the colour for reflected surface - absence of strong light source.
- It is less effective for images with strong directional lighting or high-level cues like shadow directions and reflections.

Challenges

FCN output gave 95% accuracies, so there were some cases where the labelings were mismatched. Eg. labeling a sky reflection on the sea as sky. This error poses a problem in both sky replacement as well as colour transfer.

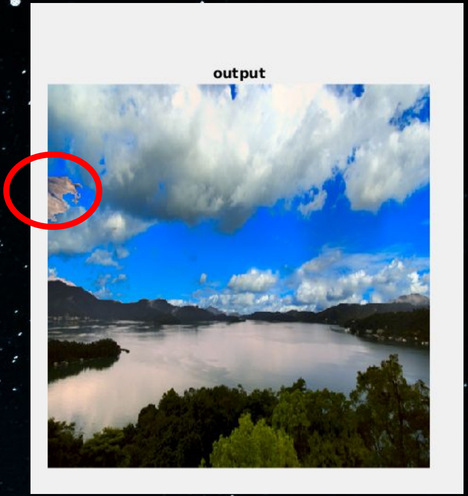

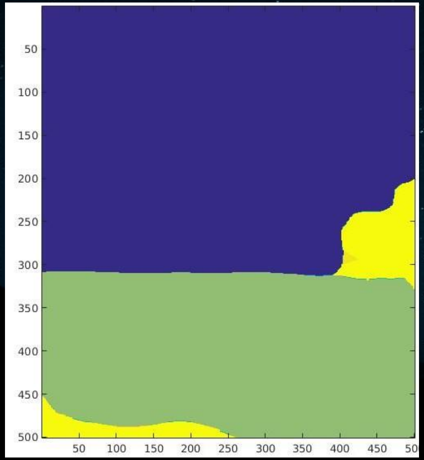

Challenges

- Computing the descriptors and FCN response map for 414 images was computationally heavy.
- The dataset used by the paper is not diverse enough to correctly support the retrieval of images.
- This kind of color transfer is less effective for images with strong directional lighting or high-level cues like shadow directions and reflections.

Timeline

Understand and modify the FCN and form a relevant response map

Getting a set of proposed skies according to the response map

Perform colour transfer of matched labels

Perform realistic composting on unmatched labels

Work distribution

Saumya

Understand and modify the FCN and compute normalized response maps
Compute descriptors for all images
Retrieve a set of proposed skies according to the response map and descriptors
Replace skies by manipulating sky areas and compositing

Anjali

Perform colour transfer of matched labels on the output of the FCN (While that is being worked on : http://dags.stanford.edu/projects/scenedataset.html)

Perform realistic composting on unmatched labels

Thank you!