Sky is Not the Limit: Semantic-Aware Sky Replacement

Yi-Hsuan Tsai[1*] Xiaohui Shen[2] Zhe Lin[2] Kalyan Sunkavalli[2] Ming-Hsuan Yang[1]
[1]University of California, Merced [2]Adobe Research

Analysis of sky segmentation results. In Table 1 and Figure 1, we show the importance of our sky refinement step and analyze the effect of each term of the online models. In Table 1, using online models for fine segmentation consistently outperforms the coarse level (FCN), especially in terms of the boundary accuracy (BPR). For the fine level, we show that IOU ratio is higher when combining three cues (color, texture and location terms). While only using the color cue obtains better result on BPR, it often produces noisy segments as shown in Figure 1. In practice, considering both the quality of entire sky segmentation and boundary accuracy, we use the model combining all the cues.

In Figure 1, FCN generates results with noisy boundaries that include foreground scenes. The online models using only one or two cues produce noisy segments that do not belong to sky regions. With the guide of the location term U_f and using all the three cues, the segmentation results are complete and accurate, especially around the boundaries. More results of our fine-level sky segmentation are shown in Figure 2.

Table 1: *Sky segmentation results with the comparison of coarse and fine level segmentation. The overlap ratio (IOU) and boundary precision-recall (BPR) are computed for evaluation.*

Methods	FCN	U_c	$U_c + U_g$	$U_c + U_g + U_f$
IOU	87.6	87.8	88.4	**88.7**
BPR	0.639	**0.869**	0.829	0.839

CCT transfer function. We use a continuous transfer function for histogram matching:

$$ct^I(x) = g(ct^R(x)) = \frac{\arctan(\frac{\alpha}{\delta}) + \arctan(\frac{ct^R(x)-\alpha}{\delta})}{\arctan(\frac{\alpha}{\delta}) + \arctan(\frac{1-\alpha}{\delta})}, \quad (1)$$

where ct^I and ct^R are the color temperature of I^{fg} after adjustment and R^{fg}, respectively. In addition, α and δ are two parameters of the mapping function which can be estimated by minimizing a cost function with the CCT histogram CT (we use 32 bins):

$$\begin{aligned}(\hat{\alpha}, \hat{\delta}) &= \arg\min_{\alpha,\delta} \sum_k (g(CT^I(k)) - \tilde{CT}(k))^2, \\ \tilde{CT}(k) &= CT^I(k) + (CT^R(k) - CT^I(k))\frac{\tau}{t} \\ &\text{s.t. } t = \min(\tau, \left|CT^R - CT^I\right|_\infty) \end{aligned} \quad (2)$$

where $CT^I(k)$ and $CT^R(k)$ present the k^{th} element of the input and reference histograms, and $\tilde{CT}$ is an intermediate histogram that normalizes the difference between these two histograms bounded by τ (we use 0.1 for all the results). Similar to the luminance transfer mentioned in the manuscript, we introduce a weighted regularization $\beta \cdot \tau$ that accounts for the sky similarity, where β has the same definition as in the manuscript. We then optimize (2) using parameter sweeping in a branch-and-bound scheme.

*e-mail:ytsai2@ucmerced.edu

More comparisons of transfer methods. We present more visual comparisons of different transfer methods in Figure 3, including the SkyFinder [Tao et al. 2009], our global transfer method without semantic cues and our semantic-aware transfer method. Note that, we implement the SkyFinder method which utilizes the transfer technique [Reinhard et al. 2001] in the Lab color space.

More results of automatic sky replacement. We show our final sky replacement results from Figure 4 to Figure 9 which have been used for our three user study tasks, and additional results in Figure 10 and Figure 11. For each test image, we use the skies from the top five reference images retrieved by our semantic search method, and generate sky replacement results.

References

REINHARD, E., ASHIKHMIN, M., GOOCH, B., AND SHIRLEY, P. 2001. Color transfer between images. *IEEE Comp. Graph. Appl. 21*, 5, 34–41.

TAO, L., YUAN, L., AND SUN, J. 2009. Skyfinder: Attribute-based sky image search. *ACM Trans. Graph. (proc. SIGGRAPH) 28*, 3.

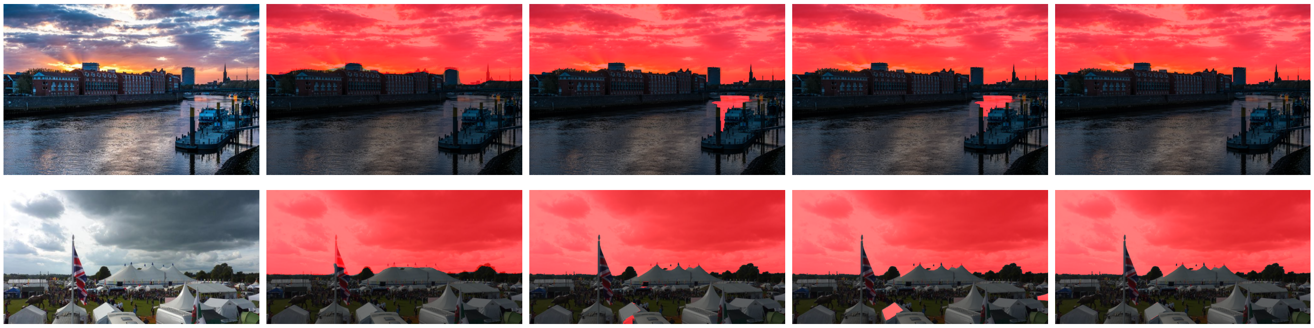

(a) Input image (b) FCN (c) U_c (d) $U_c + U_g$ (e) $U_c + U_g + U_f$

Figure 1: *Two examples of the sky segmentation results. Given the input image (a), we compare the results of coarse-level FCN segmentation (b), and the fine-level segmentation via online models from (c) to (e), where the segmentation is illustrated as the red mask.* *Photo credits: Andreas Dantzand muffinn.*

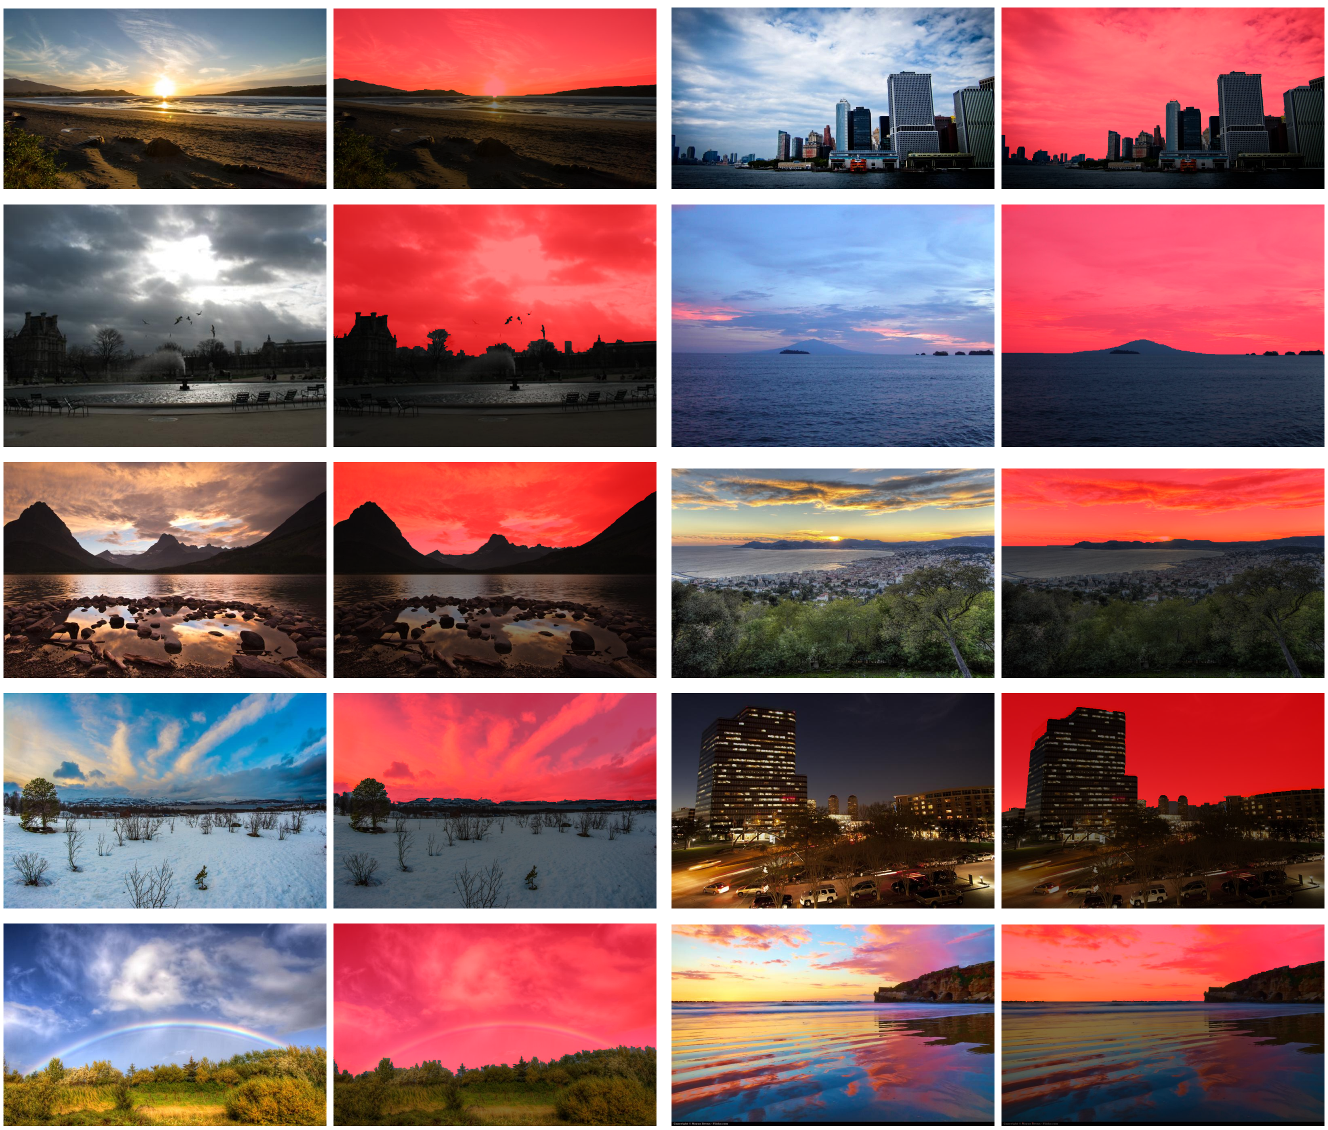

(a) Input image (b) Fine segmentation (c) Input image (d) Fine segmentation

Figure 2: *Examples of our fine-level sky segmentation results, where the segmentation is illustrated as the red mask.* *Photo credits: Tom Hall, Cédric Boismain, Jeff P, Vander Muniz, Axel Kristinsson, Nan Palmero, Hans Kylberg, ludovick, josephdepalma and Moyan Brenn.*

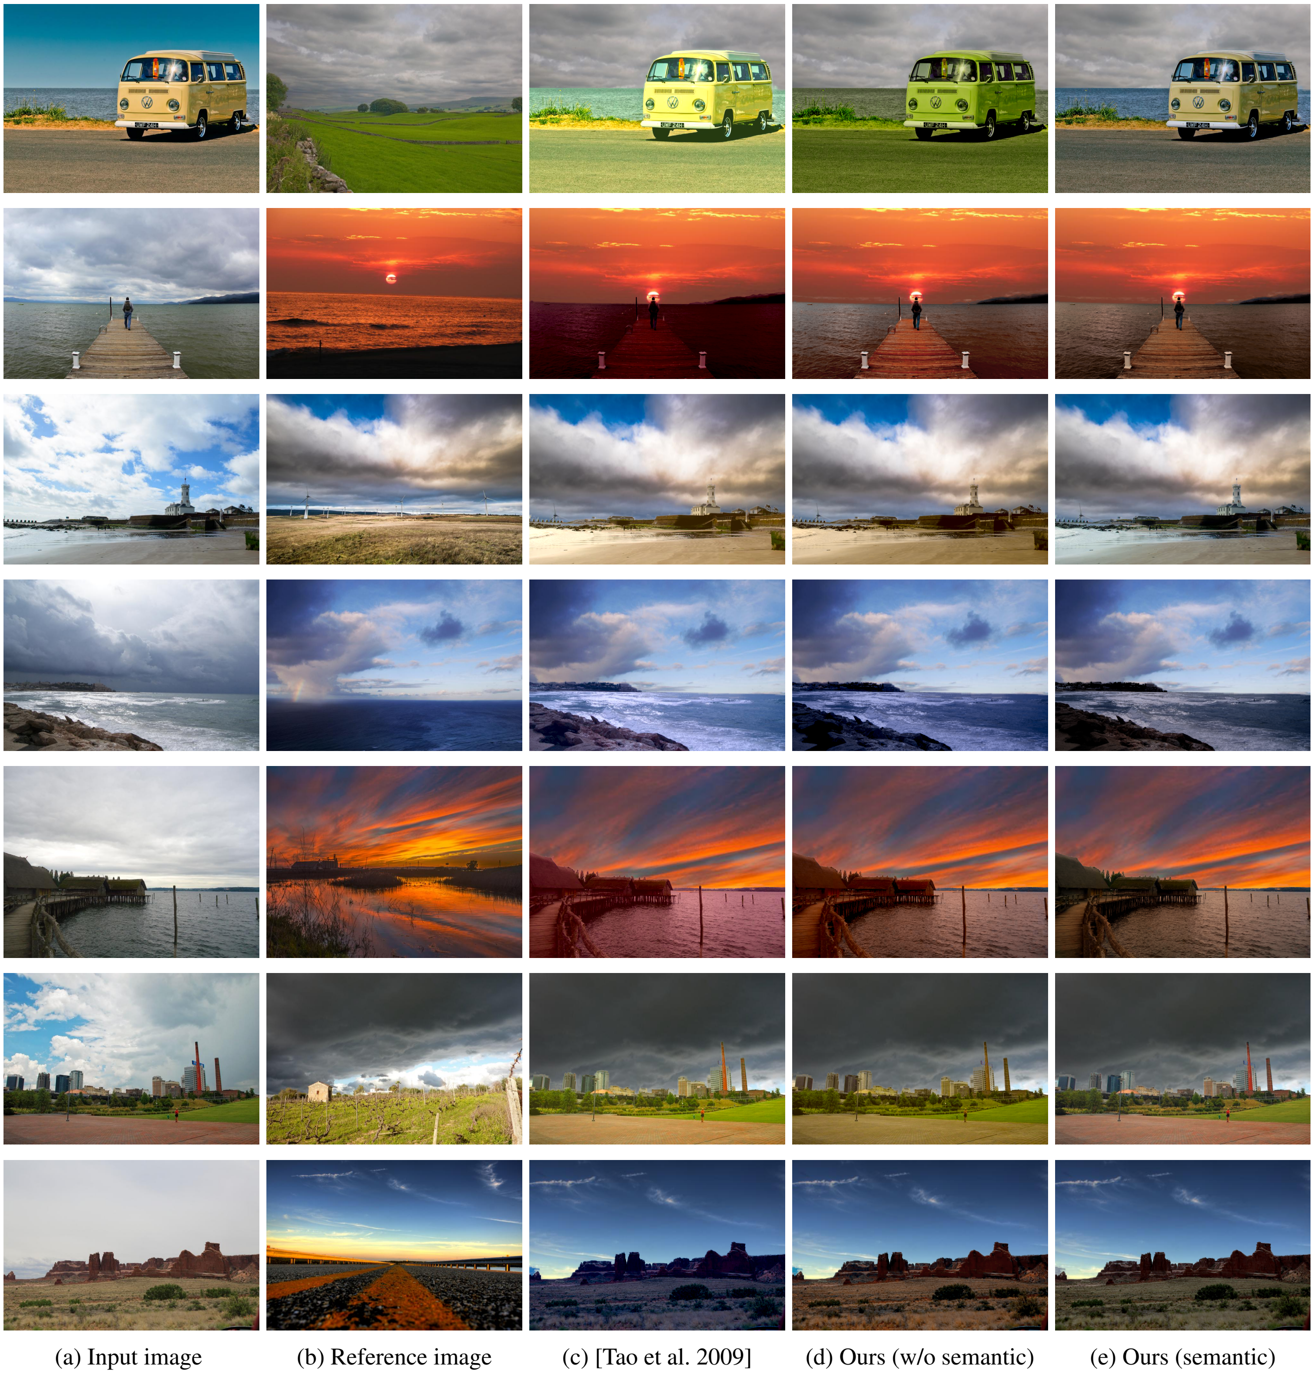

Figure 3: *Comparison of different sky transfer methods. Given the input image (a) and reference image (b), we show the result of the transfer method proposed in SkyFinder [Tao et al. 2009] (c), our method without using semantic matching (d) and our semantic transfer approach (e). Photo credits: Beverley Goodwin, daveynin, gordon.milligan, Cédric Boismain, Sílvia Martín, Max Wolfe and sugarbear96.*

(a) Input image

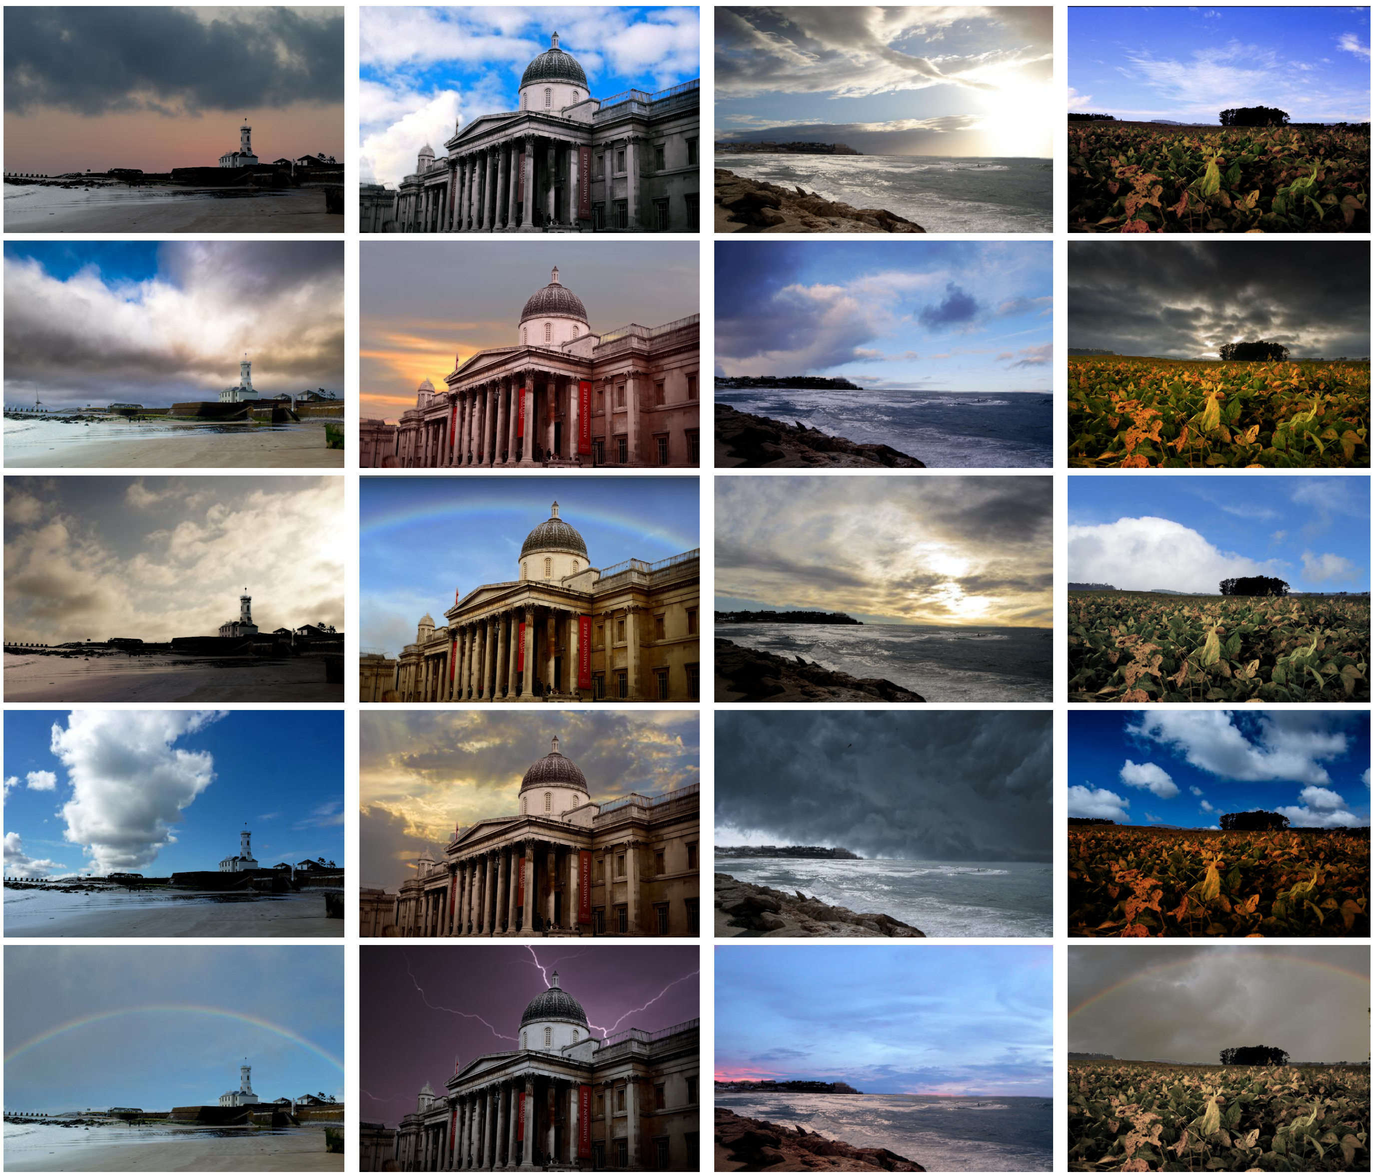

(b) Our results

Figure 4: *Sky replacement results automatically generated by our method. Given an input image (a), we show the top five results (b) with a set of diverse skies. Photo credits: gordon.milligan, Shimelle Laine, Cédric Boismain and Nicholas A. Tonelli.*

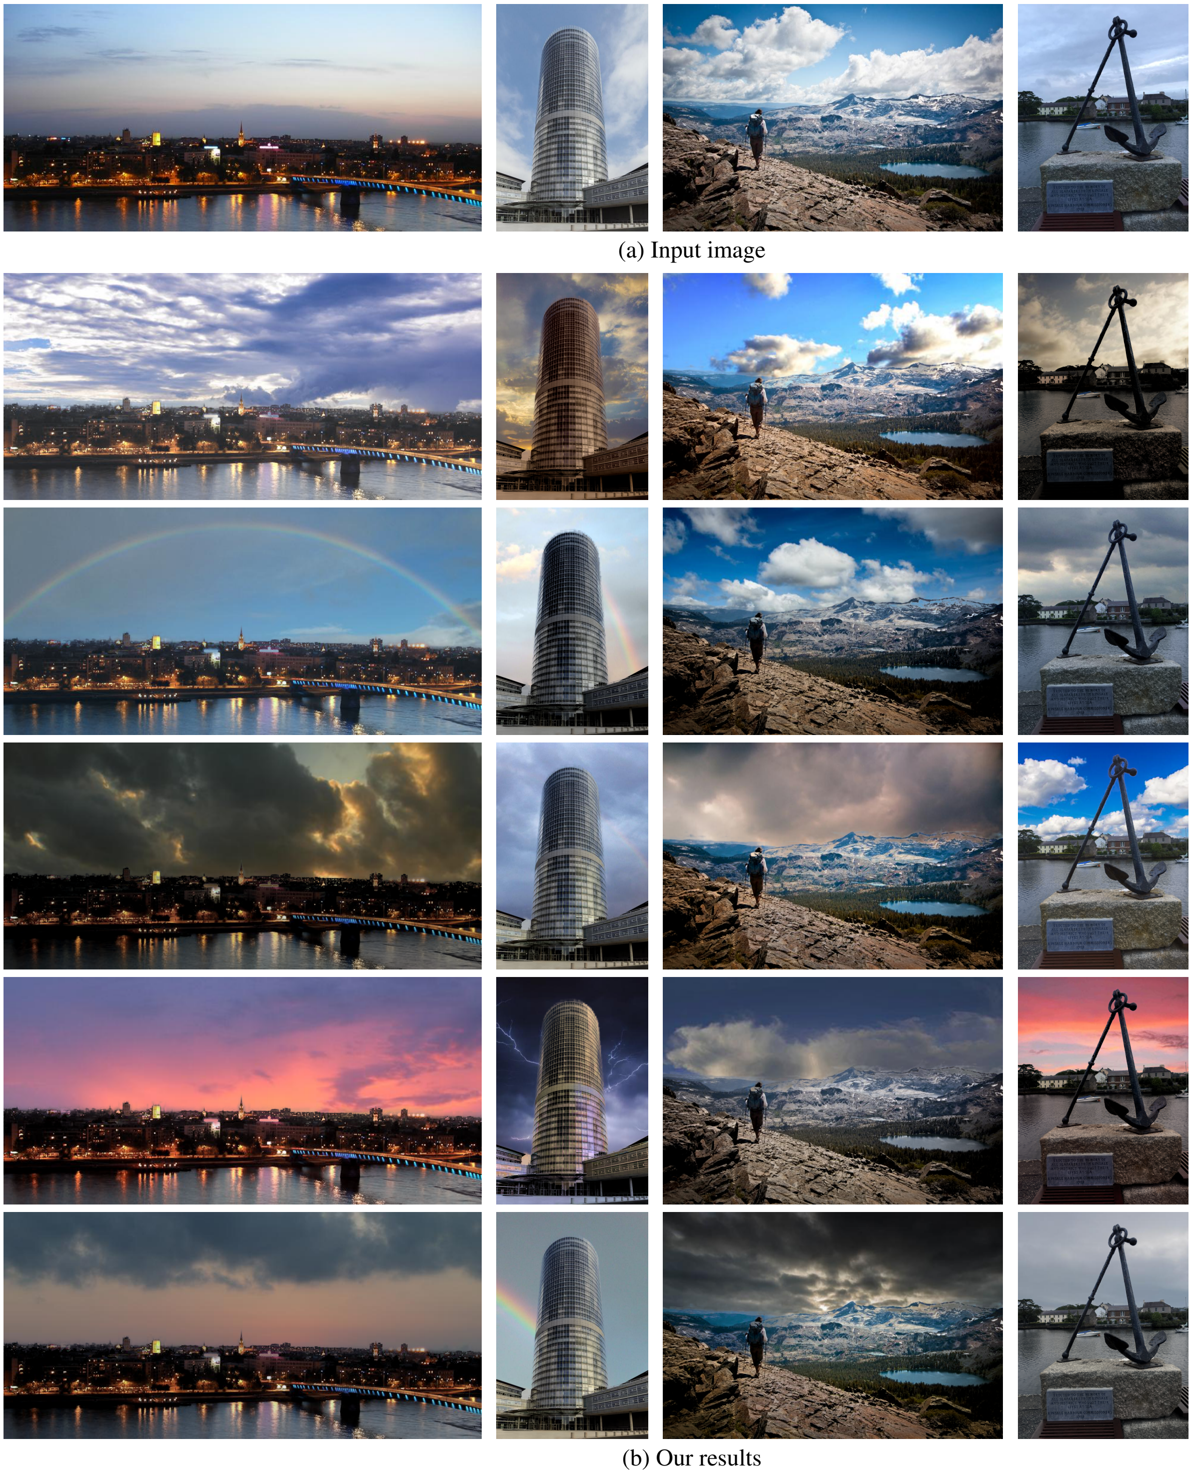

Figure 5: *Sky replacement results automatically generated by our method. Given an input image (a), we show the top five results (b) with a set of diverse skies.* *Photo credits: Ivan Aleksic, David Schiersner, Jeff P and Scott Cohen.*

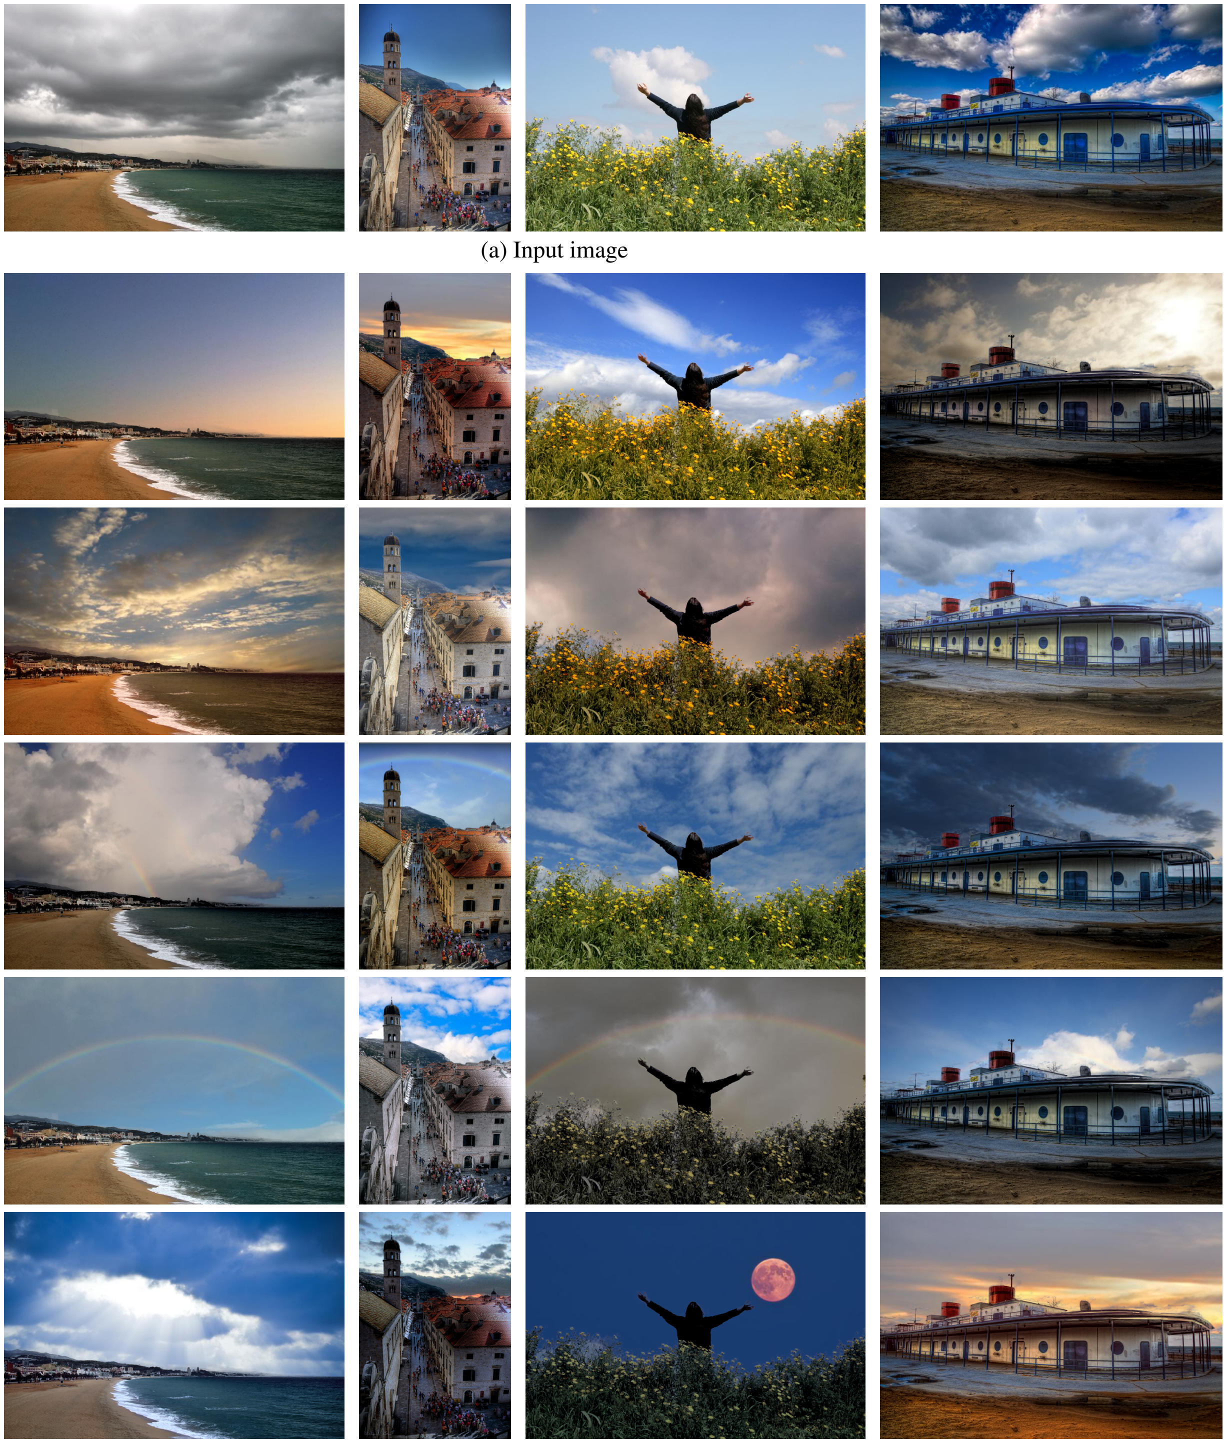

(a) Input image

(b) Our results

Figure 6: *Sky replacement results automatically generated by our method. Given an input image (a), we show the top five results (b) with a set of diverse skies.* *Photo credits: Jorge Franganillo, Michael Caven, Vic and Rick Seidel.*

(a) Input image

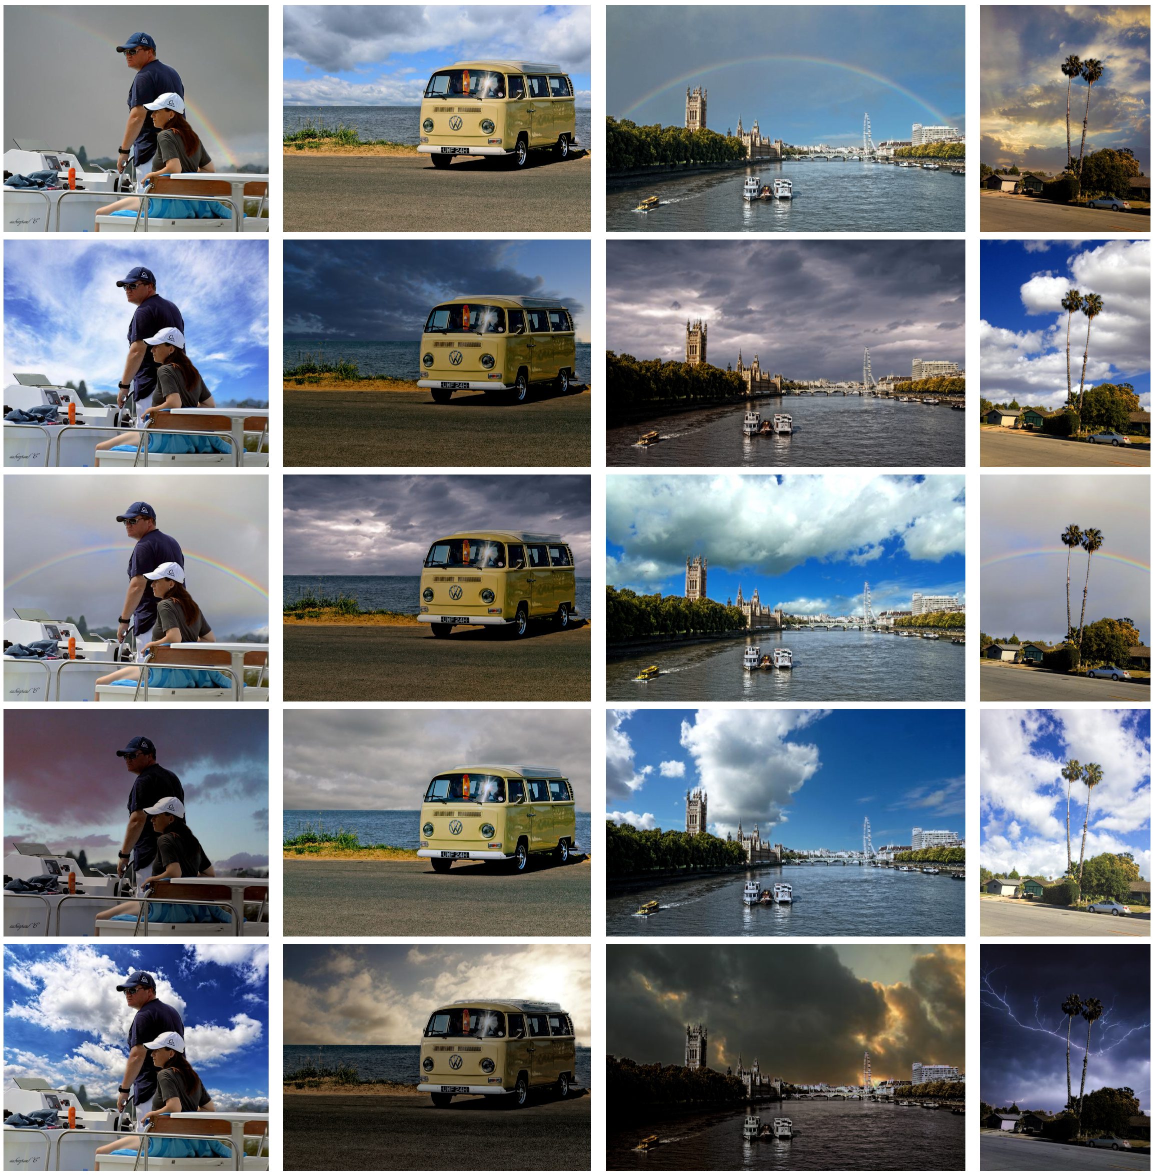

(b) Our results

Figure 7: *Sky replacement results automatically generated by our method. Given an input image (a), we show the top five results (b) with a set of diverse skies.* *Photo credits: sabin paul croce, Beverley Goodwin, Berit Watkin and Scott Cohen.*

(a) Input image

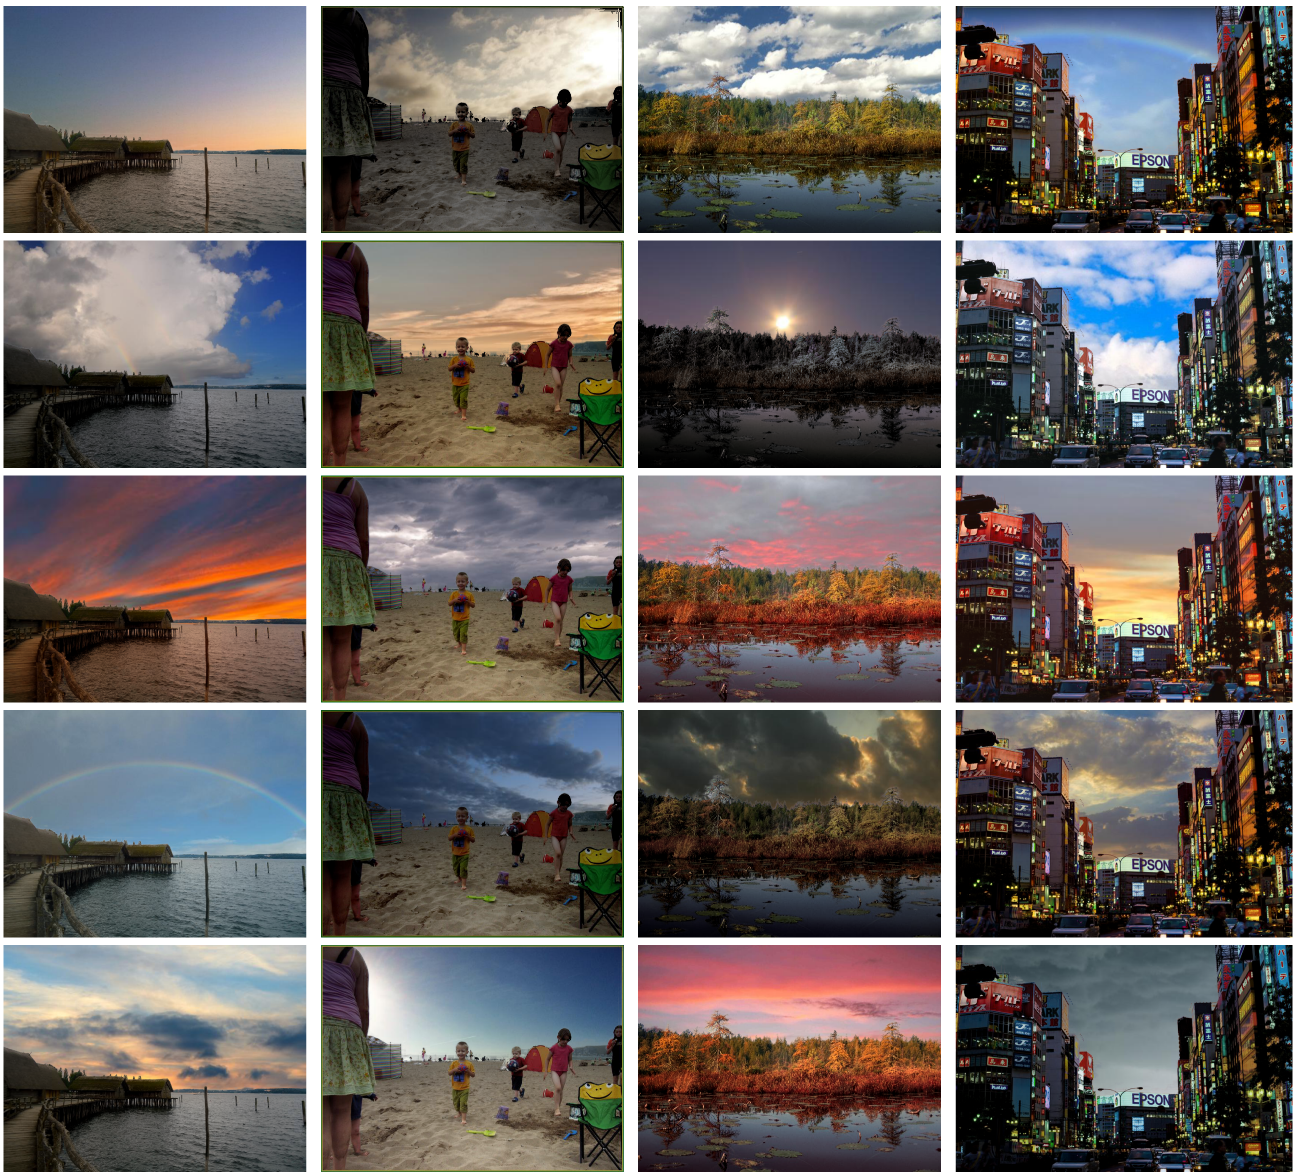

(b) Our results

Figure 8: *Sky replacement results automatically generated by our method. Given an input image (a), we show the top five results (b) with a set of diverse skies. Photo credits: Sílvia Martín, Graham Richardson, Andy Arthur and Joi Ito.*

(a) Input image

(b) Our results

Figure 9: *Sky replacement results automatically generated by our method. Given an input image (a), we show the top five results (b) with a set of diverse skies. Photo credits: Eoin O'Mahony, Paula J Andrews, sugarbear96 and Scott Cohen.*

(a) Input image

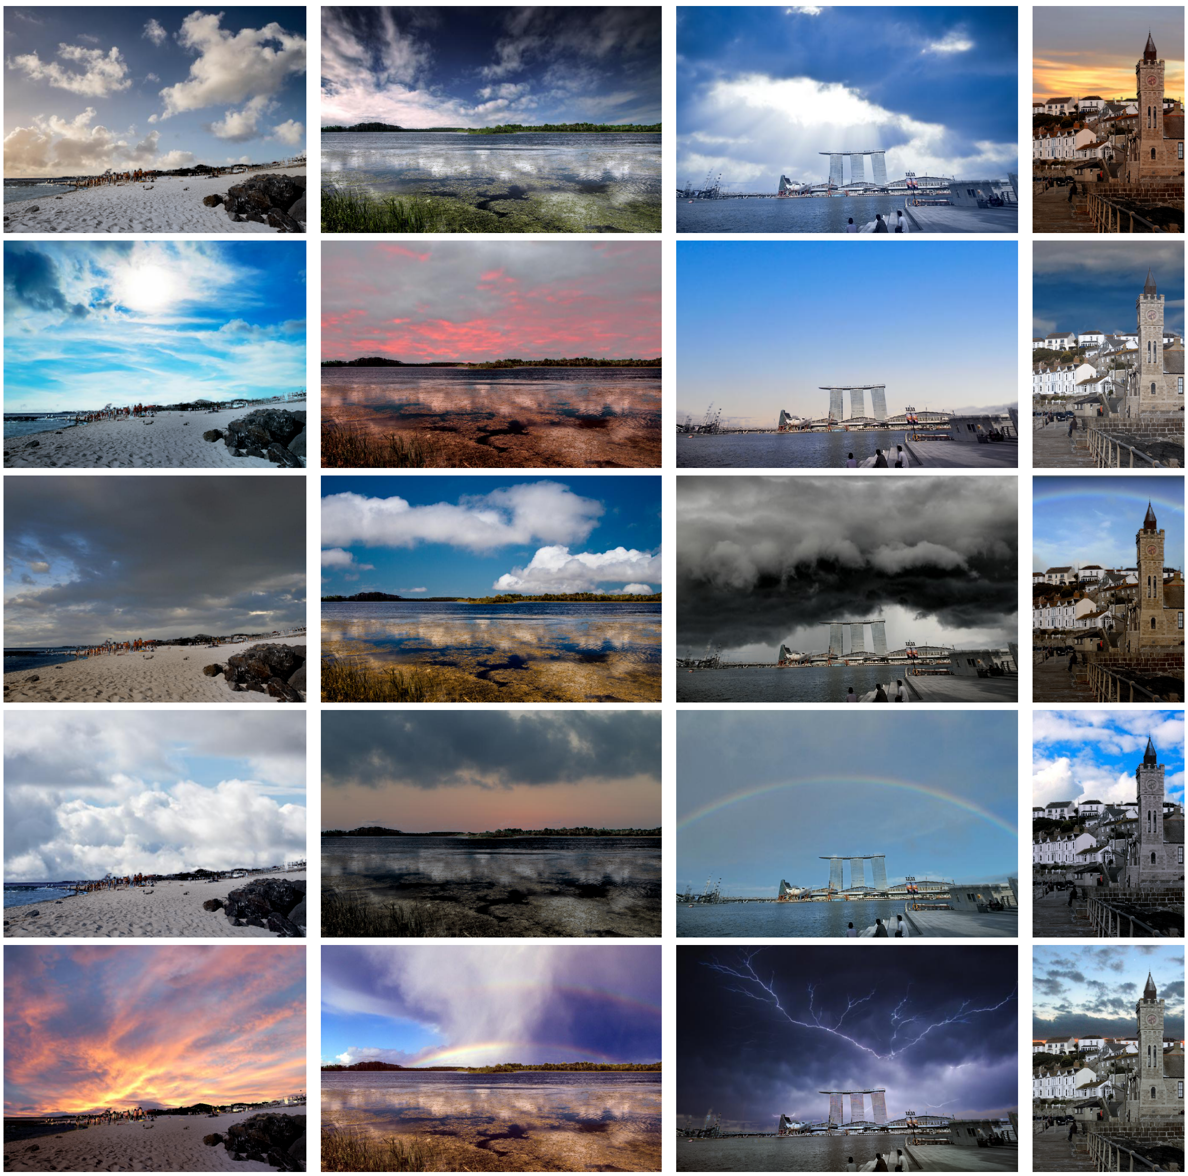

(b) Our results

Figure 10: *Sky replacement results automatically generated by our method. Given an input image (a), we show the top five results (b) with a set of diverse skies. Photo credits: Kim MyoungSung, Ricardo's Photography, motiqua and umbrellahead56.*

(a) Input image

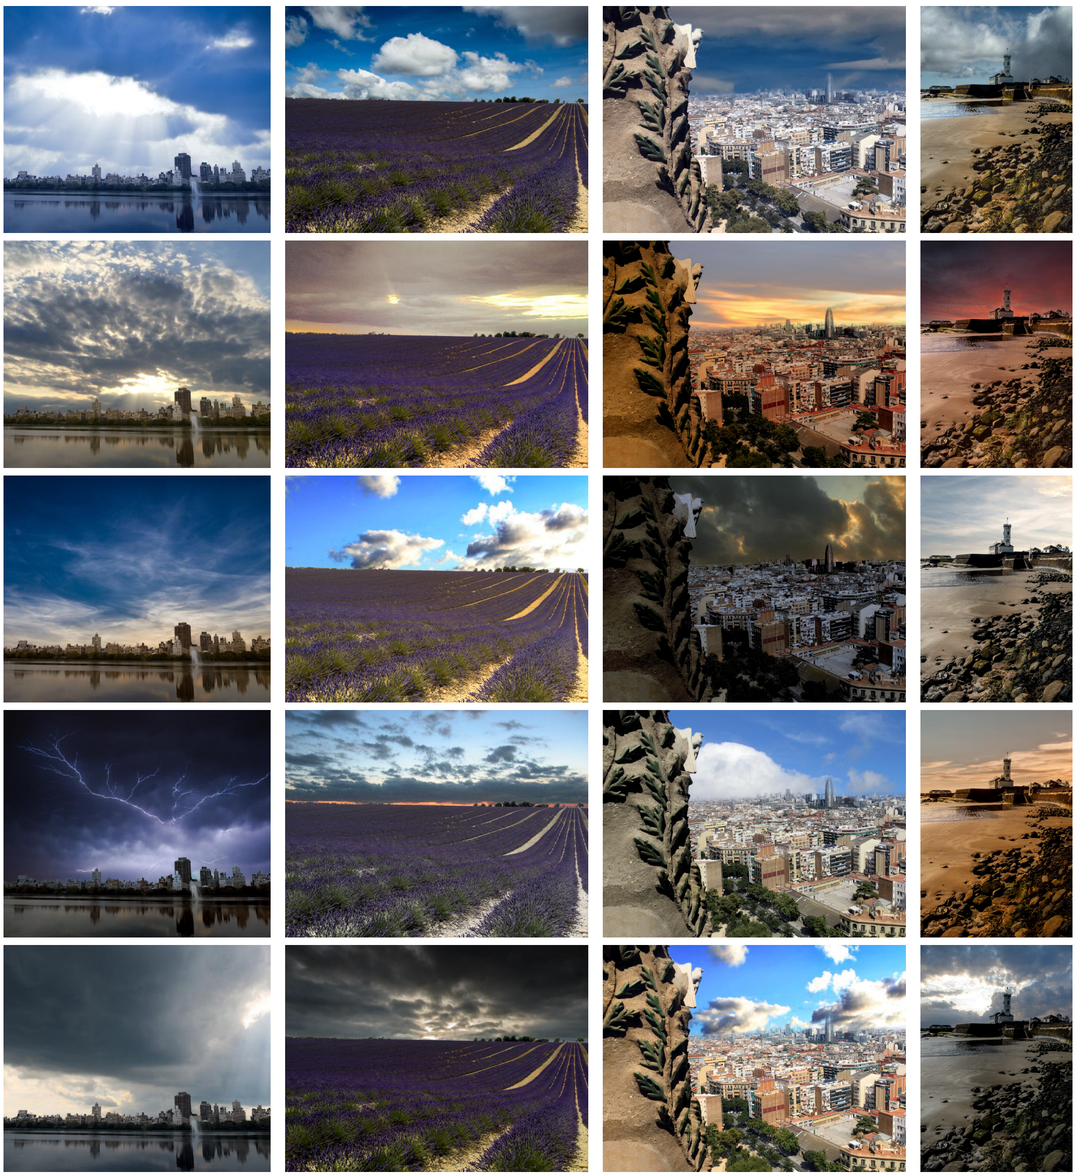

(b) Our results

Figure 11: *Sky replacement results automatically generated by our method. Given an input image (a), we show the top five results (b) with a set of diverse skies. Photo credits: Guillermo Palacios, Scott Cohen and gordon.milligan.*