

Sky pixel detection in outdoor imagery using an adaptive algorithm and machine learning



Kerry A. Nice^{a,d,e,*}, Jasper S. Wijnands^a, Ariane Middel^f, Jingcheng Wang^b, Yiming Qiu^b, Nan Zhao^b, Jason Thompson^a, Gideon D.P.A. Aschwanden^a, Haifeng Zhao^a, Mark Stevenson^{a,c}

^a Transport, Health, and Urban Design Research Hub, Faculty of Architecture, Building, and Planning, University of Melbourne, Australia

^b School of Computing and Information Systems, University of Melbourne, Australia

^c Melbourne School of Engineering; and Melbourne School of Population and Global Health, University of Melbourne, Australia

^d School of Earth, Atmosphere and Environment, Monash University, Australia

^e Cooperative Research Centre for Water Sensitive Cities, Melbourne, Australia

^f School of Arts, Media and Engineering (AME), School of Computing, Informatics, Decision Systems Engineering (CIDSE), Arizona State University, USA

ARTICLE INFO

Keywords:

Sky view factor
Google street view
Machine learning
WUDAPT
Sky pixel detection
Skyfinder

ABSTRACT

Computer vision techniques enable automated detection of sky pixels in outdoor imagery. In urban climate, sky detection is an important first step in gathering information about urban morphology and sky view factors. However, obtaining accurate results remains challenging and becomes even more complex using imagery captured under a variety of lighting and weather conditions.

To address this problem, we present a new sky pixel detection system demonstrated to produce accurate results using a wide range of outdoor imagery types. Images are processed using a selection of mean-shift segmentation, K-means clustering, and Sobel filters to mark sky pixels in the scene. The algorithm for a specific image is chosen by a convolutional neural network, trained with 25,000 images from the Skyfinder data set, reaching 82% accuracy for the top three classes. This selection step allows the sky marking to follow an adaptive process and to use different techniques and parameters to best suit a particular image. An evaluation of fourteen different techniques and parameter sets shows that no single technique can perform with high accuracy across varied Skyfinder and Google Street View data sets. However, by using our adaptive process, large increases in accuracy are observed. The resulting system is shown to perform better than other published techniques.

1. Introduction

Sky pixel detection in images is an ongoing computer vision challenge with a large range of applications such as autonomous vehicle or drone navigation (Shen and Wang, 2013), real-time weather classification (Roser and Moosmann, 2008), image editing (Laffont et al., 2014; Tao et al., 2009), sky replacement (Tsai et al., 2016), and scene parsing (Tighe and Lazebnik, 2010; Hoiem et al., 2005). It is also an important tool in urban climate research. The use of fisheye photography to calculate sky view factors (SVF), the

* Corresponding authors at: Transport, Health, and Urban Design Research Hub, Faculty of Architecture, Building, and Planning, University of Melbourne, Australia.

E-mail address: kerry.nice@unimelb.edu.au (K.A. Nice).

<https://doi.org/10.1016/j.uclim.2019.100572>

Received 18 February 2019; Received in revised form 2 December 2019; Accepted 6 December 2019
2212-0955/ © 2019 Elsevier B.V. All rights reserved.

fraction of sky visible to a point, at individual locations has long been used in urban climate. Numerous techniques exist to process this sort of imagery (Grimmond et al., 2001; Chapman and Thornes, 2004; Ali-Toudert and Mayer, 2007).

In computer vision research, sky detection techniques have largely followed two main paths, either finding the pixels associated with the sky on a pixel by pixel basis or by finding a sky/ground boundary and labelling the sky as everything above that boundary. The first approach focuses on finding individual pixels associated with the sky. Luo and Etz (2002) employed a physics-based approach using the changes of sky colours from zenith to horizon. Gallagher et al. (2004) generated sky pixel probability maps based on colour values and a two-dimensional polynomial for each colour channel. Zafarifar and de With (2007) added texture, gradients, and vertical position to colour values to generate their probability map, however, only blue sky is detected accurately with clouds marked with low probabilities. Schmitt and Priese (2009) added in an analysis of position and shape and were reportedly able to also accurately perform sky detection under cloudy conditions.

Other approaches have focused on finding a sky/ground boundary. Straight lines were first used to define a horizon located via an energy function (Ettinger et al., 2003). Using an improved energy function optimisation and gradient information from the image, Shen and Wang (2013) allowed the horizon line to follow the boundary instead of being restricted to a straight line. Additional variations allowed increasingly difficult sky regions (i.e. regions separated from the main sky region by buildings, flags, or other obstructions) to be detected (Zhijie et al., 2014, 2015). Another approach, not specifically designed for sky pixel identification, attempts to classify around a dozen classes (such as sky, buildings, trees, and cars) in outdoor imagery through semantic segmentation using trained convolutional neural networks (CNN), such as SegNet (Badrinarayanan et al., 2017), and other variations (Holder et al., 2016; Middel et al., 2019).

Comparing existing approaches is difficult, as most studies do not report accuracy metrics, and when they do, they are reported using different metrics and benchmark data sets. Luo and Etz (2002) reports 90.4% correct detection of blue sky pixels and 13% misclassifications. Chapman and Thornes (2004) reports a root-mean-squared error (RMSE) of 0.06 in marking SVF in fisheye images. Schmitt and Priese (2009) reports an accuracy of 0.90 for 80% of their validation images and an accuracy of 0.95 for 75% of the images. Liang et al. (2017), using SegNet, reports an accuracy of 0.96 for sky pixel identification. Shen et al. (2018), using an off-the-shelf version of SegNet, reports an accuracy of 0.83 with lateral views. Finally, Middel et al. (2019) reports an accuracy of 0.95 with lateral views.

An effort is ongoing to provide worldwide databases of standardised urban morphology information (Ching et al., 2018, 2019). Now with the widespread availability of urban imagery, an opportunity exists to expand the range of research that can be conducted without the requirement of manually collecting urban morphology parameters and to accelerate the population of databases such as the World Urban Database and Portal Tool (WUDAPT) (Mills et al., 2015).

Studies have started to utilise automated methods to build SVF data sets from Google Street View (GSV) imagery (Middel et al., 2018; Gong et al., 2018). The results are promising but show some accuracy problems. Urban areas with large numbers of street trees in particular are cited by Gong et al. (2018) as a key source of inaccuracies in their system. In addition, these systems are highly dependent on GSV imagery, which as of 2018 has become more restrictive to license and expensive to obtain. This necessitates the need to expand the type of imagery used, imagery that unlike GSV, might vary more in lighting, weather conditions, camera angles, and aspect ratios.

With these factors in mind, we present a sky pixel detection system that has been tested using various types of outdoor imagery collected under a wide range of lighting and weather conditions, camera angles, and aspect ratios. This system, built on artificial intelligence, is adaptive and uses a range of algorithms and combinations of parameters to locate the sky pixels to ensure the highest accuracy for each individual class of images. We evaluate a number of existing and new techniques for sky pixel classification and demonstrate that this new adaptive system performs with greater accuracy than any of these individual techniques on their own.

2. Methods

Our sky pixel identification system used data from two main sources, the Skyfinder data set (Mihail et al., 2016) and GSV (Google Maps, 2017). Three computer vision techniques and a number of parameter variations were used to process the data (see Table 1). Two were developed/customised as part of this research. One existing method was implemented and also used as a benchmark. The overall process flow is shown in Fig. 1. Finally, a previously published fourth technique (i.e. Sobel/flood-fill) was also used as a benchmark test for our system.

2.1. Data

2.1.1. GSV data

Panoramas for 406 locations in a variety of cities (Adelaide, Brisbane, Paris, Sydney, Tokyo, Perth, and Melbourne) were retrieved using the Google Maps API. Images were retrieved as six 640×640 tiles (one each for up, down, left, right, front, and back directions). The six images were stitched together into a 1280×960 cubic image using Java 8 (Oracle, 2018) and OpenCV (Bradski, 2000). Validation images were created by hand marking sky regions (with blue, RGB 0,0,255) in each image using the GNU Image Manipulation Program (GIMP, 2019). Fig. 2 shows an example of a GSV panorama image and the corresponding hand-marked validation image. This data set was only used as part of the validation data set.

2.1.2. Skyfinder data

This data set was built from 90,000 long-term timelapse images from 53 outdoor webcams over a variety of lighting and weather

Table 1
Techniques and variations of each algorithm used in this study for sky pixel detection.

Sobel		Implementation of Wang et al. (2015)'s Sobel operator/hybrid probability model		
Variations	Probability threshold			
Sobel_50	0.5			
Sobel_60	0.6			
Sobel_70	0.7			
Sobel_80	0.8			
Sobel_90	0.9			
Sobel_95	0.95			

Mean shift Variations		Algorithm developed by the authors based on mean shift segmentation		
	Spatial radius (pixels)	Range radius (pixels)	Min. density (pixels)	
Mean_7_8_300	7	8	300	
Mean_3_6_100	3	6	100	
Mean_5_7_210	5	7	210	
Mean_7_6_100	7	6	100	

K-means Variations		Algorithm developed by the authors based on K-means clustering and HSL colour filtering					
	Clusters	Skyreq	H _{high}	H _{low}	L _{lightness}	L _{grey}	S _{grey}
K-mean_12	12	0.7	0.75	0.3	0.95	0.75	0.2
K-mean_6	6	0.6	0.75	0.3	0.95	0.75	0.2
K-mean_14	14	0.4	0.75	0.3	0.95	0.65	0.2

Sobel/flood-fill		Middel et al. (2018)'s Sobel operator/flood-fill algorithm used as benchmark		
------------------	--	--	--	--

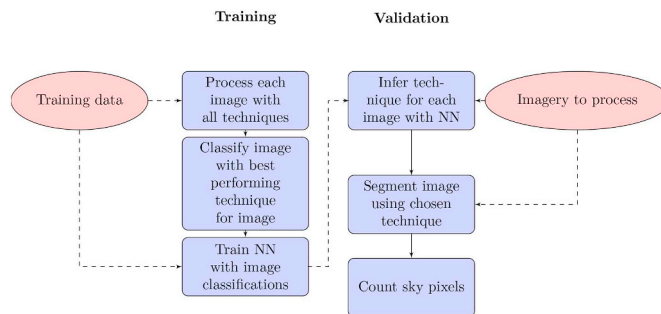


Fig. 1. Process flow of training and validation steps.

conditions. Images are of a wide range of sizes and aspect ratios, including 640×489 , 857×665 , 960×600 , 1280×720 , and 1280×960 . For each location, a binary sky mask was created by Mihail et al. (2016) for validation purposes. An example image and mask (with sky marked with white, RGB 255,255,255) are shown in Fig. 2. All of these images are available from the Skyfinder website (Mihail, 2015).

In this study, we selected 38,115 images from 40 locations. Night-time and images with heavy fog were removed as these are conditions unlikely to be encountered in imagery used to calculate SVF. The data set was split into two data sets, 28,586 for neural network training and 9529 for validation.

2.2. Techniques and parameters

2.2.1. Wang et al. (2015) Sobel operator/hybrid probability model

An implementation of the sky detection algorithm presented in Wang et al. (2015) was implemented using OpenCV and Java 8. This method proceeds by calculating grey-scale gradient images using x- and y-directional Sobel operators to estimate sky colour. An optimised objective function attempts to find the best sky/ground boundary in the gradient image using the covariance matrices of a first calculation of sky and ground regions. Using this best sky boundary, probability models are created from i) the centre and standard deviations of the colours, ii) the gradient values, and iii) the vertical position of each pixel (vertically higher pixels are more likely to be sky). An overall probability model, ranking each pixel's probability (0 to 1) to be sky, is generated from these three




Fig. 2. a) Original GSV panorama image and b) hand-marked validation image (sky marked with blue, RGB 0,0,255). c) Original Skyfinder image and d) Skyfinder validation mask (sky marked with white, RGB 255,255,255). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

probability models. Wang et al. (2015) reports an error average of 0.051 and standard deviation of 0.058 in their evaluation using human-labelled images.

Wang et al. (2015) did not recommend a probability threshold, so a number of thresholds were tested (0.5, 0.6, 0.7, 0.8, 0.9, and 0.95) and given the designations of Sobel_50, Sobel_60, Sobel_70, Sobel_80, Sobel_90, and Sobel_95 respectively (Table 1). The algorithm was applied to each image and pixels that exceeded the chosen threshold were marked as sky pixels (using blue, RGB 0,0,255). Results from our implementation are shown in Fig. 3. Note, two of the best performing variations are also used for a benchmark comparison in Section 3.3.1.

2.2.2. Mean shift segmentation algorithm

Mean shift is an algorithm often used for image segmentation (Comaniciu and Meer, 1997, 2002). Image segmentation involves decomposing images into homogeneous contiguous regions of pixels of similar colours or grey levels. Mean shift iteratively picks search windows (spatial and range) of a certain radius in an initial location in an image, computes a mean shift vector, and translates the search window by that amount until convergence (Comaniciu and Meer, 1997). Segmentation results are highly dependent on input parameters for the algorithm, which include the spatial and colour range radius of the search window and minimum density (the minimum number of pixels to constitute a region). The mean shift used in this project is based on a Java port by Pangburn (2002) of the C++ based EDISON vision toolkit (Christoudias et al., 2002).

Four different variations of the input parameters were used, determined experimentally through a sensitivity test to work across the widest variety of images. A comparison of a GSV and a Skyfinder location with different mean shift variations and marked results is shown in Fig. 4. The technique designations and parameters are detailed in Table 1. Mean shift is applied to each image with the chosen set of parameters and pixels of the most common colour (in the top half of the segmented image) are marked as sky.

To illustrate the effects of different mean shift variations can have in a variety of images, under varying sky conditions, results from two additional locations are shown in Fig. 5. To shift the entire sky to a single colour, images with patchy multi-coloured clouds are more accurately segmented when the radius and density parameters are increased (e.g. in column b compared to column a in Fig. 5). However, at other locations, this can have the effect of creating false positives, for example, where the building in the centre left background is increasingly segmented into the sky in column d in Fig. 5 compared to column c.

2.2.3. K-means clustering and HSL colour filtering

A third sky segmentation technique was designed using K-means clustering and hue, saturation, and lightness (HSL) colour filtering. K-means clustering iteratively splits an image into K number of clusters, terminating when a specified criteria is met (i.e. maximum iterations and/or desired accuracy). The K-means clustering was performed using the K-means method from the OpenCV library. Three different input parameter settings were used, determined experimentally through a sensitivity test to work on a wide variety of images. The technique designations and parameters are detailed in Table 1. K-mean_6 was more accurate with cloudy skies. K-mean_14 handles sky scenes broken up by tree canopies. The last variation, K-mean_12, generally performed with low accuracy for most images but in a small number of cases (when the sky is mostly obscured by a solid object, building or bridge but not trees)




Fig. 3. Results of Sobel operator/hybrid probability model, showing a) Sobel operator gradient image, b) resulting probability predictions, c) Sobel_50, d) Sobel_60, e) Sobel_70, f) Sobel_80, g) Sobel_90, and h) Sobel_95 for a GSV image. Subfigures i) to p) same as a) to h) but with Skyfinder image.

performed better than all the other techniques and variations.

After K-means clustering, cluster regions were processed based on HSL values. The following conditions (for H , hue, S , saturation, and L , lightness) must be met to add a colour region to a list of possible sky regions:

$$H_{low} < H < H_{high}$$

$$\cup L > L_{lightness}$$

$$\cup L > L_{grey} \cap S < S_{grey}$$

Of these possible sky clusters, only clusters with a number of pixels greater than the *Skyreq* threshold (percent of all pixels) in the image were finally marked as sky regions. Example results are shown in Fig. 6.

2.2.4. Middel et al. (2018) Sobel operator/flood-fill algorithm

For benchmark comparisons, we used an algorithm developed by Middel et al. (2018). The process is based on a Sobel filter (Sobel and Feldman, 1968) and flood-fill algorithm (Laungrunthip et al., 2008; Middel et al., 2017). The method was designed to calculate SVF from GSV image cubes that were projected into upwards facing fisheye views. Note, this algorithm also rescales the imagery to 512×512 . All 38,521 images in the combined training and validation data set were processed with this algorithm (sky pixels marked with white, RGB 255,255,255), compared to validation images, and results saved for a comparison with our results. The results were kept separate from the other 13 techniques and were not included in the NN training process (see Section 2.3.2). Also, in our benchmark comparison, we used this system to process a varied outdoor imagery data set, not the fisheye imagery (cropped below the horizon) the algorithm was originally designed to process. Examples are shown in Fig. 7.

2.3. Neural network

2.3.1. Inception v3

The Microsoft Cognitive Toolkit (CNTK) (Yu et al., 2015; Agarwal et al., 2016), with the Inception v3 network (Szegedy et al., 2016), was used in this project to route images through our adaptive algorithm. This neural network (NN) is a widely used model for image classification across a large variety of fields (Xia et al., 2017; Hassannejad et al., 2016). The original Inception model (Szegedy et al., 2015) is a deep convolutional neural network (CNN) built using convolutions, average pooling, max pooling, concatenations, dropouts,




Fig. 4. Comparison outputs of intermediate mean shift segmentation algorithm processing steps using varying parameters, showing columns a) Mean_3_6_100, b) Mean_7_6_100, c) Mean_5_7_210, d) Mean_7_8_300 and intermediate mean shifted (GSV row 1, Skyfinder row 3) and the final marked images (GSV row 2, Skyfinder row 4).




Fig. 5. Examples of false negatives and false positives with different mean shift variations at two additional locations of mean shift segmentation parameters. Intermediate mean shifted (row 1) and the final marked images (row 2) of Mean_7_6_100 (columns a and c) and Mean_5_7_210 (columns b and d).

and fully connected layers. A reduction of parameters and increased performance over other competing architectures was achieved by the use of Inception modules, combining convolution operations with varying filter sizes. Inception v3 deepens the network (i.e., 42 layers, see Fig. 8) without a large increase in computational requirements by factorising convolution operations with large filter sizes into smaller stacked layers.

2.3.2. Neural network training

The Skyfinder data set of 38,115 images was split into two data sets for training (75%) and validation (25%). All training and validation images (which consisted of images of a wide variety of sizes and aspect ratios) were preprocessed and rescaled using ImageJ (Rueden et al., 2017) to Inception v3's native input size, 299×299 , prior to training and validation.

The network was calibrated using supervised learning, trained with a list of images assigned to categories, in our case, which one of the 13 sky detection technique variations performed most accurately for each image. The training process was run until the model




Fig. 6. Results of K-means clustering and HSL colour filtering (K-mean_6), showing GSV (top row) and Skyfinder (bottom) and K-means clustered image (left), HSL intermediate image (middle), and final marked image (right).

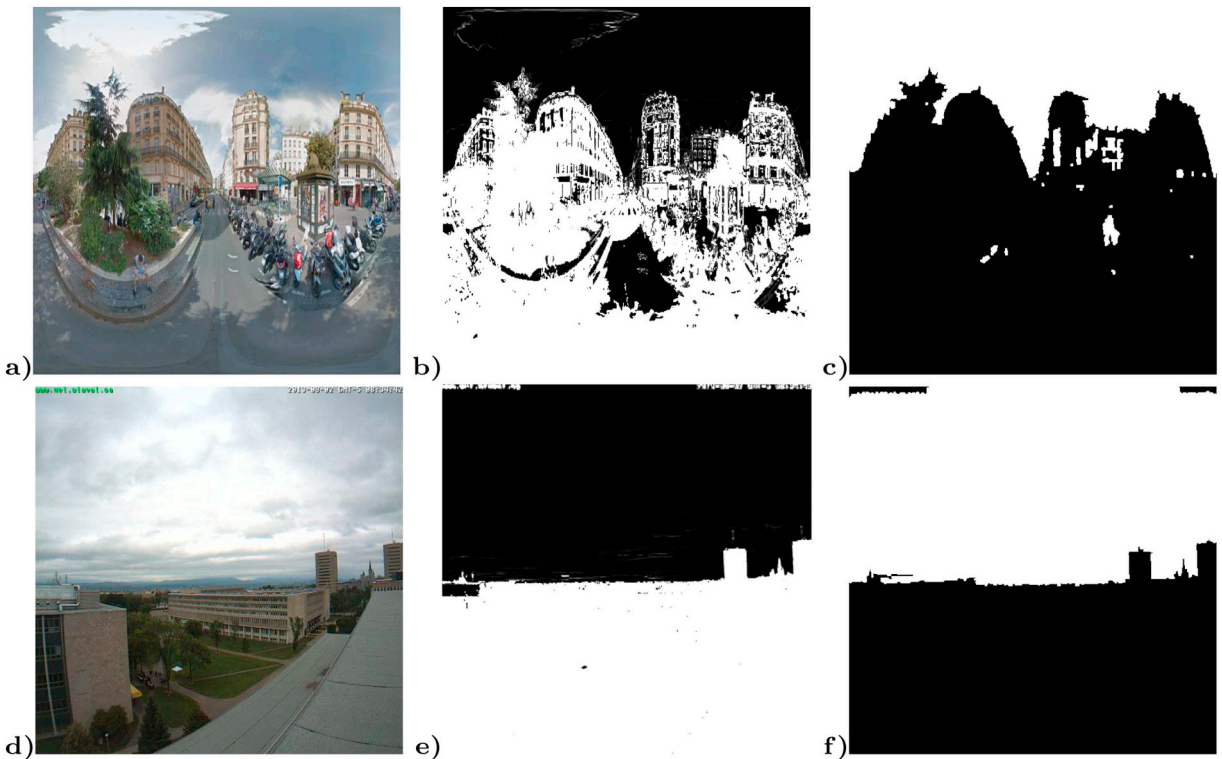


Fig. 7. Results of Sobel/flood-fill combination, showing a) original GSV image (rescaled to 512×512), b) intermediate Sobel image, and c) final marked sky image. Subfigures d) to f) same as a) to c) but using Skyfinder imagery.

reached peak accuracy (convergence) at recognising these classifications from images. Note, none of the GSV imagery was used in the NN training process.

2.3.3. Neural network inference

Using the trained model, inferences were performed using the images from the validation data set (the 25% of images from Section 2.3.2) as well as the 406 GSV panoramas. The techniques and parameters picked by the NN as the most appropriate for that image were used to mark the sky pixels. The marked sky pixels were then compared to the ground truth to assess accuracy.

3. Results

Three sets of results are reported in this section. The first is a comparison of all the techniques and parameters run individually




Fig. 8. Inception v3 architecture (Google, 2019).

against the Skyfinder and GSV data sets (a total of 38,521 images). The second presents the results of 9636 validation images using our adaptive process flow with the techniques and parameters chosen by the trained NN. The third presents a comparison to two benchmark models: a) the Wang et al. (2015) Sobel operator/hybrid probability model (from which we chose the two best performing variations) and b) the Middel et al. (2018) Sobel operator/flood-fill algorithm.

3.1. Results from all techniques

All the technique and parameter variations were used to process the two data sets of the 38,115 Skyfinder and 406 GSV images. Fig. 9a and b shows a comparison of RMSE statistics for the two data sets. Plots of a number of the better performing techniques are

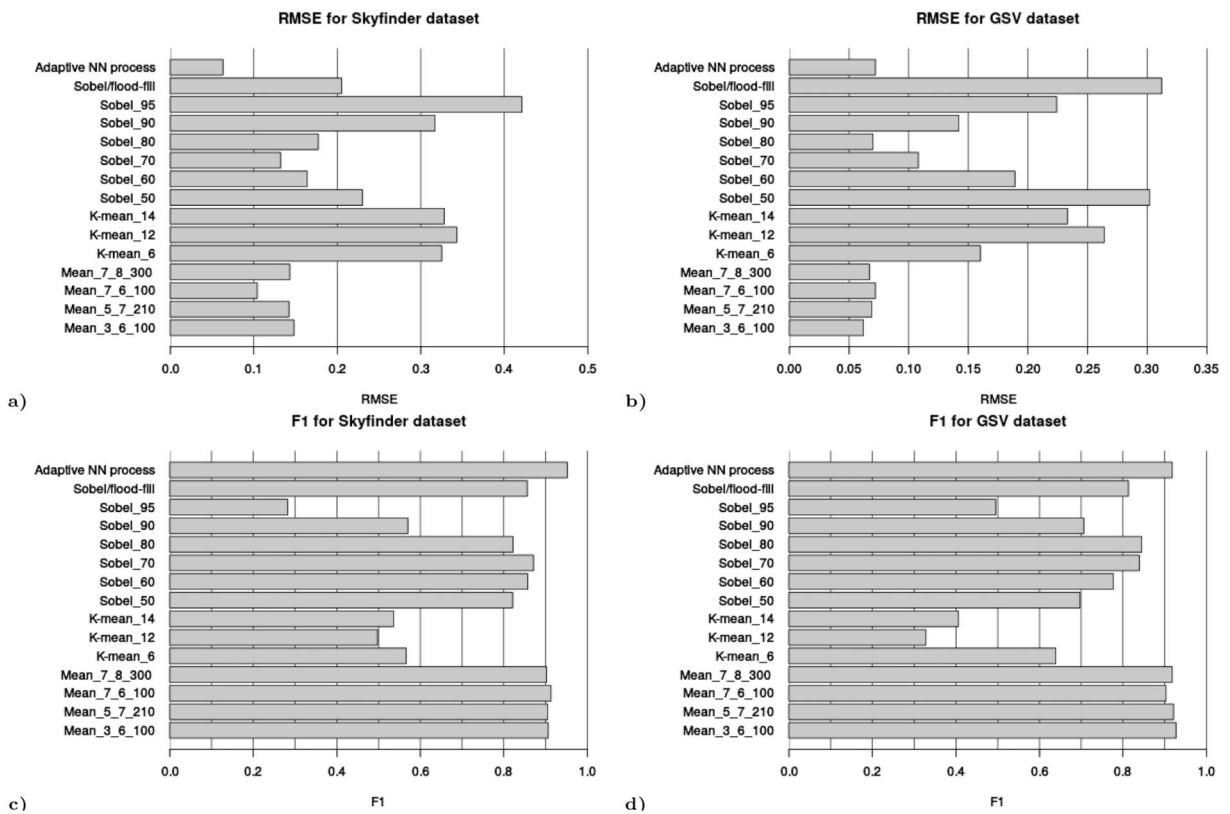


Fig. 9. Evaluation of all techniques and parameters showing statistics for RMSE (lower values are more accurate) for a) Skyfinder data set and b) GSV data set. Evaluation of all techniques and parameters showing F1 statistics (higher values are more accurate) for the 9636 image validation data set, split by c) Skyfinder and d) GSV images.

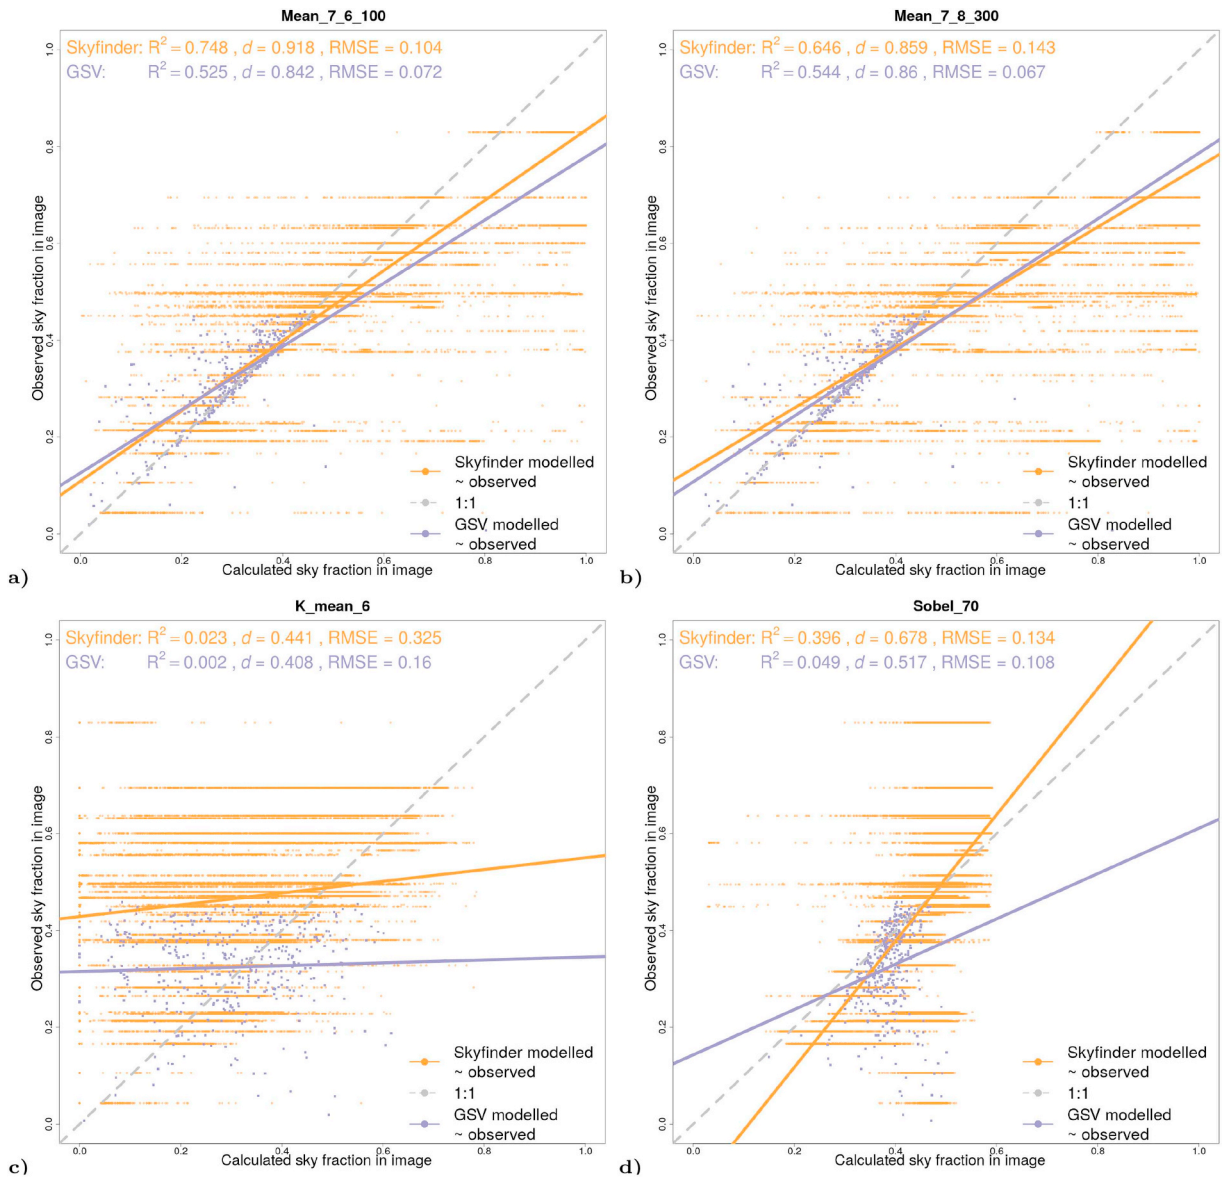


Fig. 10. Observed vs. calculated sky pixels using the a) Mean_7_6_100, b) Mean_7_8_300, c) K-mean_6, and d) Sobel_70 techniques on the 38,521 image combined data set (406 GSV and 38,115 Skyfinder images).

presented in Fig. 10. Note, the strong horizontal lines in these figures are due to the nature of the Skyfinder data set that contains large groups of the same scenes (with the same percentage of sky) under different lighting and weather conditions, often resulting in a wide range of predicted results. A complete summary of Willmott (1981)'s index of agreement, *d*, R², and RMSE statistics for evaluations against the two data sets are presented in Appendix Table 2. F1 statistics for each technique against the 9636 validation images are presented in Fig. 9c and d. A complete summary of precision, recall and F1 statistics are shown in Appendix Table 3.

3.2. Results from neural network classified techniques

Fig. 11a presents a theoretical best case. If the NN was 100% accurate in picking the best technique from the 13 possible combinations based on its training, a RMSE of 0.026 and 0.020 (and *d* index of agreement of 0.994 and 0.988) is possible against the validation data set of 9529 Skyfinder images and 406 GSV images respectively. Samples of the imagery used in training for selected classifications are shown in Fig. 12. As can be seen in this figure, there are no strong visual themes in each of the classifications (i.e. all very cloudy, clear blue sky, or multi-coloured sky), however the NN is able to pick up on more subtle features not readily visible to the eye.

The NN was trained for 250 epochs on an Nvidia GeForce GTX 1080 GPU, requiring about 12 h. The NN reached a peak accuracy

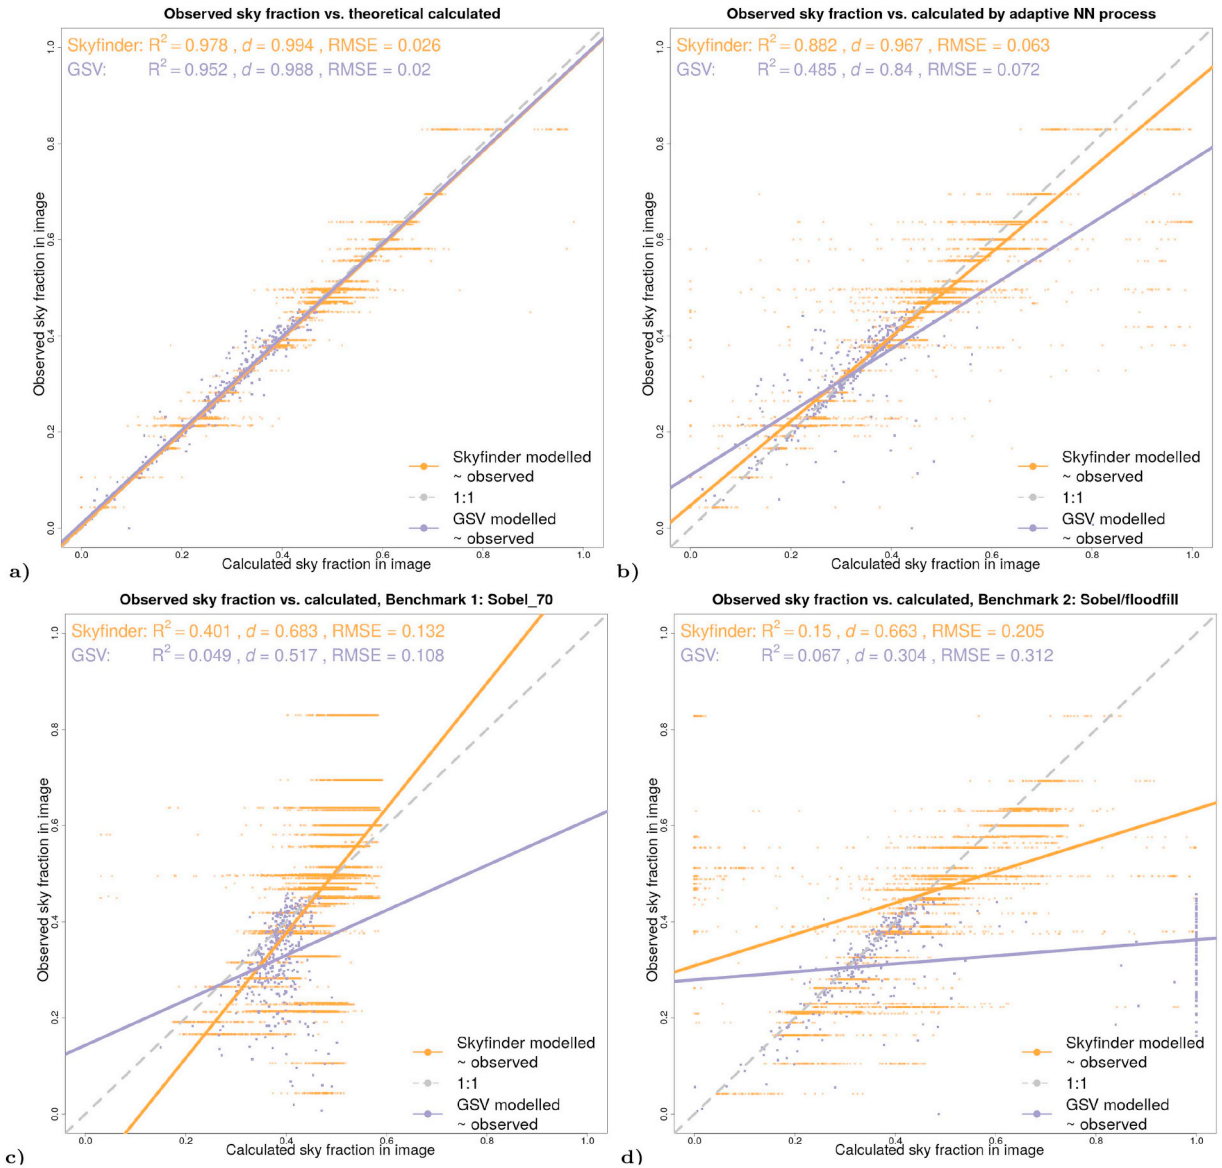


Fig. 11. a) Theoretical best case results if the NN is 100% accurate for the 9636 validation images. Other subfigures show results of b) our adaptive NN process, c) Benchmark 1: Sobel_70, and d) Benchmark 2: Sobel/flood-fill combination.

rate of 52.6% in choosing the optimal algorithm from the 13 options. However, as some techniques did only slightly better than others, the error introduced by picking the second or third best algorithm is generally limited. Breaking down the accuracy with this in mind, the NN picked the best method 52.6% of the time, the second best 18.9%, and the third best 10.5% (for a total of the three of 82.0%).

Fig. 11b shows the overall accuracy of the NN chosen pathway process flow against the 9636 validation images, for the Skyfinder and GSV images (respectively) from the validation data set with a RMSE of 0.063 and 0.072, R^2 of 0.882 and 0.485, and d of 0.967 and 0.840. The accuracy of the NN has impacted the overall accuracy of the system (i.e. not reaching the theoretical accuracy of 0.026 or 0.020 RMSE), but the results on out-of-sample images are still very good. In addition, precision, recall, and F1-scores for the Skyfinder and GSV validation data sets show good results (see Appendix Table 3). Precision is 0.946 and 0.933, recall is 0.965 and 0.918, while the F1-scores are 0.952 and 0.918 (all respectively).

3.3. Benchmark results

3.3.1. Results from the Wang et al. (2015) Sobel operator/hybrid probability model

To compare the results from our adaptive process (shown in Section 3.2), we picked the Sobel variations that performed best with




Fig. 12. Selected imagery used for NN training for classifications a) Mean_7_8_300, b) Mean_7_6_100, c) K-mean_6, and d) Sobel_70.

the two types of imagery in the validation data set. Sobel_70 (shown in Fig. 11c) performed best with the Skyfinder imagery, with a RMSE of 0.132, R^2 of 0.401, and d of 0.683. The best performing variation with the GSV imagery, Sobel_80 (figure not presented), resulted in a RMSE of 0.07, R^2 of 0.433 and d of 0.655. Precision, recall, and F1-scores statistics (see Appendix Table 3) show that Sobel_70 has a lower precision than Sobel_80 with the Skyfinder validation images (at 0.869 vs. 0.920) while Sobel_70 has better recall scores than Sobel_80 (0.914 vs. 0.781). Similar patterns are seen with the GSV imagery. Overall, the F1-scores are in the range of 0.82 to 0.87 for both techniques and data sets.

3.3.2. Results from the Middel et al. (2018) Sobel operator/flood-fill algorithm

In the evaluation of the Sobel/flood-fill algorithm, results from the Skyfinder and GSV 9636 validation images are shown in Fig. 11d. This algorithm yields a RMSE of 0.205, R^2 of 0.150 and d of 0.663 against the Skyfinder portion of validation data set and a RMSE of 0.312, R^2 of 0.067, and d of 0.304 against with the GSV portion of the validation data set. The results from the evaluation of GSV imagery showed a number of images miss-marked as 100% sky, inflating the error rate for this data set. A similar problem was seen with the Skyfinder images, several of them were miss-marked as 0% sky. This is also reflected in the precision, recall, and F1-scores (see Appendix Table 3). Precision scores are 0.840 and 0.761 (for Skyfinder and GSV) are low while the recall scores are much higher (0.900 and 0.948), resulting in overall F1-scores of 0.856 and 0.813.

4. Discussion

The results from Section 3.1 show that no single technique and parameter combination performs sky pixel identification with high accuracy across the data sets used by this project. These data sets contain a wide variety of outdoor scenes with various lighting and weather conditions (as can be seen in some of the sample images in Fig. 12), challenging many of the techniques. It was expected that algorithms would perform better with GSV imagery, due to their regularity. These images were captured with the same type of equipment, using the same camera angles (horizon at 50% image height), under clear sky or partly cloudy conditions. The results show that almost all of the variations perform better with the GSV data than the Skyfinder data. Some of the variations even approach the accuracy of our system with the GSV data, for example Sobel_80.

However, the Skyfinder data set challenged all of the variations with the Mean shift and Sobel based methods achieving no better than 0.1 to 0.2 RMSE. However, for some individual images, even the poorest performing techniques (i.e. K-means variations) excelled compared to all of the other techniques. Also, some techniques perform poorly for certain images. In Fig. 10d, the results for the Sobel_70 method show wide variations in R^2 between the Skyfinder and GSV data sets, while the RMSE values are roughly similar. In the case of Sobel_70, sky fractions for images with low sky fractions (a small number of images in the data set) are systematically overestimated but this has a significant impact on the GSV R^2 value. Both of these cases validate the need for an adaptive process that can respond to the specific challenges each image presents to deliver overall better results than any single algorithm. Further, this allows certain techniques to be avoided in the cases that they will perform poorly.

Having a broad range of technique variations was important for the overall accuracy. Experimentation was performed to reduce the number of classifications to possibly increase the accuracy of the NN. However, in removing some of the worst performing methods (many of the K-means variations), the overall accuracy degraded. While some of the variations had very low accuracy overall, they were the best choice for some images. Having those available more than compensated for the NN not always picking the most accurate variations for each image. For example, reducing the choices to the best three performing classes (Mean_7_6_100, K-mean_6, and Sobel_70) showed a reduction in RMSE to 0.039 over the best accuracy of the 13 approach combination (RMSE of 0.026 and 0.020 for Skyfinder and GSV respectively). This attempt to increase the accuracy of the NN predictions highlights a limitation of this study. 53% accuracy in picking between 13 approaches shows that there is room for improvement with this part of the method.

With a larger training data set, a lower NN error rate might be achieved, improving performance towards the theoretical RMSE of 0.026 for sky pixel identification. New algorithms and variations of existing algorithms can be added to the system to handle new imagery with greater accuracy, enabled by our flexible framework, as future research. In addition, a NN could be used in a more systematic manner and explore predicting technique parameters directly.

One difficulty in this study was comparing different sky pixel detection schemes in an objective manner. As noted in the introduction, most studies either do not provide metrics, or provide different types of metrics. Also, with our results evaluating the Skyfinder and GSV data sets showing some large differences in performance, a lack of standardised benchmarks makes comparisons less meaningful. We attempted to overcome these difficulties by implementing some existing methods and including them in our evaluation against common data sets. In addition, the needs of the eventual application should be kept in mind. As precision and recall scores often varied widely for each algorithm, the impact of either higher false positives or false negatives should guide algorithm choice or at least be considered in the results. For example, an algorithm with higher false positives will overestimate sky pixels, leading to a higher SVF estimate and possibly higher maximum temperatures in urban canyon modelling.

5. Conclusion

In conclusion, we present a system of sky pixel identification that shows high accuracy rates with varied and challenging outdoor imagery. This system sits between algorithms that can be quickly set up and run but are not as accurate with challenging data sets, e.g. [Middel et al. \(2018\)](#), and more complex systems such as [Gong et al. \(2018\)](#) that require a more extensive deep learning algorithm. Our adaptive system uses the best elements of each in pursuit of the most accurate results possible.

In comparison to published methods, our adaptive process performs well. Our accuracy against the Skyfinder images of RMSE of 0.063 compares well to the RMSE of 0.205 for the [Middel et al. \(2018\)](#) Sobel/flood-fill algorithm and the best performing [Wang et al. \(2015\)](#) Sobel variations, Sobel_70 and Sobel_80, achieving an RMSE of 0.132 and 0.177. Similarly, our adaptive process also performs well on the GSV validation images with a RMSE of 0.072 compared to 0.07 and 0.312 for Sobel_80 and Sobel/flood-fill respectively. Finally, our adaptive process performs with the best precision, recall, and F1-scores for the Skyfinder data set compared to any of the evaluated techniques.

In the Zenodo data set corresponding to this article (see Section 6), we provide our trained NN model (and configuration files), which can be used to infer the best algorithm for any type of outdoor imagery, as well as all the training and validation imagery used in this study. This system can then be used out of the box. Using our system, it will be possible to populate databases (such as WUDAPT) of urban morphology information. This data set also provides a standardised data set to reproduce our results and allow benchmark comparisons with other sky pixel detection systems.

6. Code availability and licensing

Code and data are available from the corresponding author on request. Also, data and code are available at <https://doi.org/10.5281/zenodo.2562396> and <https://bitbucket.org/politemadness/skypixeldetection> ([Nice, 2019](#)) and are distributed under the Creative Commons Attribution-NonCommercial-ShareAlike 4.0 Generic (CC BY-NC-SA 4.0) license.

Declaration of Competing Interest

None.

Acknowledgements

The support of the Commonwealth of Australia through the Cooperative Research Centre program is acknowledged. At Monash University, Kerry Nice was funded by the Cooperative Research Centre for Water Sensitive Cities, an Australian Government initiative. At the University of Melbourne, Kerry Nice was funded by the Transport, Health, and Urban Design (THUD) research hub and a Graham Treloar Fellowship for Early Career Researchers.

Appendix A. Appendix

A.1. Additional data tables

The following tables provide the underlying data for [Fig. 9](#).

Table 2

Evaluation of all techniques and parameters for the full Skyfinder and GSV data sets showing statistics for R^2 , RMSE, and d index of agreement.

Designation	Skyfinder			GSV		
	R^2	RMSE	d	R^2	RMSE	D
Adaptive NN process	0.882	0.063	0.967	0.485	0.072	0.840
Sobel/flood-fill	0.124	0.211	0.651	0.067	0.312	0.304
Sobel_95	0	0.421	0.372	0.053	0.224	0.388
Sobel_90	0.041	0.317	0.443	0.214	0.142	0.516
Sobel_80	0.255	0.177	0.561	0.433	0.070	0.652
Sobel_70	0.396	0.134	0.678	0.049	0.108	0.517
Sobel_60	0.288	0.164	0.659	0.005	0.189	0.387
Sobel_50	0.104	0.230	0.550	0.026	0.302	0.287
K-mean_14	0.029	0.328	0.439	0.023	0.233	0.307
K-mean_12	0.028	0.343	0.427	0.017	0.264	0.299
K-mean_6	0.023	0.325	0.441	0.002	0.160	0.408
Mean_7_8_300	0.646	0.143	0.859	0.544	0.067	0.860
Mean_7_6_100	0.748	0.104	0.918	0.525	0.072	0.842
Mean_5_7_210	0.659	0.142	0.862	0.526	0.069	0.852
Mean_3_6_100	0.658	0.148	0.854	0.603	0.062	0.866

Table 3

Evaluation of all techniques and parameters for the 9636 image validation data set, split by Skyfinder and GSV images, showing statistics for precision, recall, and F1 scores.

Designation	Skyfinder			GSV		
	Precision	Recall	F1	Precision	Recall	F1
Adaptive NN process	0.946	0.965	0.952	0.933	0.918	0.918
Sobel/flood-fill	0.840	0.900	0.856	0.761	0.948	0.813
Sobel_95	0.921	0.181	0.282	0.931	0.349	0.495
Sobel_90	0.936	0.436	0.570	0.916	0.586	0.706
Sobel_80	0.920	0.781	0.822	0.866	0.839	0.845
Sobel_70	0.869	0.914	0.871	0.788	0.925	0.839
Sobel_60	0.807	0.967	0.857	0.675	0.959	0.777
Sobel_50	0.742	0.987	0.821	0.565	0.973	0.697
K-mean_14	0.708	0.484	0.536	0.640	0.349	0.406
K-mean_12	0.677	0.441	0.497	0.547	0.272	0.328
K-mean_6	0.658	0.538	0.566	0.694	0.663	0.639
Mean_7_8_300	0.879	0.944	0.902	0.941	0.909	0.918
Mean_7_6_100	0.912	0.930	0.913	0.951	0.876	0.903
Mean_5_7_210	0.879	0.949	0.905	0.946	0.911	0.921
Mean_3_6_100	0.876	0.956	0.906	0.952	0.915	0.927

References

- Agarwal, A., Akchurin, E., Basoglu, C., Chen, G., Cyphers, S., Droppo, J., Eversole, A., Guenter, B., Hillebrand, M., Huang, X., Huang, Z., Ivanov, V., Kamenev, A., Kranen, P., Kuchaiev, O., Manousek, W., May, A., Mitra, B., Nano, O., Navarro, G., Orlov, A., Padmilac, M., Parthasarathi, H., Peng, B., Reznichenko, A., Seide, F., Seltzer, M.L., Slaney, M., Stolcke, A., Wang, H., Yao, K., Yu, D., Zhang, Y., Zweig, G., 2016. An Introduction to Computational Networks and the Computational Network Toolkit. MSR-TR-2014-112, Microsoft.
- Ali-Toudert, F., Mayer, H., 2007. Effects of asymmetry, galleries, overhanging façades and vegetation on thermal comfort in urban street canyons. *Solar Energy* 81 (6), 742–754. <http://linkinghub.elsevier.com/retrieve/pii/S0038092X06002623>.
- Badrinarayanan, V., Kendall, A., Cipolla, R., 2017. SegNet: a deep convolutional encoder-decoder architecture for image segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.* 39 (12), 2481–2495.
- Bradski, G., 2000. The OpenCV Library. Dr. Dobb's Journal of Software Tools. URL: <http://www.drdoobs.com/open-source/the-opencv-library/184404319>.
- Chapman, L., Thornes, J.E., 2004. Real-time sky-view factor calculation and approximation. *J. Atmos. Ocean. Technol.* 21, 730–741.
- Ching, J., Mills, G., Bechtel, B., See, L., Feddema, J., Wang, X., Ren, C., Brousse, O., Martilli, A., Neophytou, M., Mouzourides, P., Stewart, I., Hanna, A., Ng, E., Foley, M., Alexander, P., Aliaga, D., Niyogi, D., Shreevastava, A., Bhalachandran, P., Masson, V., Hidalgo, J., Fung, J., Andrade, M., Baklanov, A., Dai, W., Milcinski, G., Demuzere, M., Brunzell, N., Pesaresi, M., Miao, S., Mu, Q., Chen, F., Theeuwes, N., 2018. World Urban Database and Access Portal Tools (WUDAPT), an Urban Weather, Climate and Environmental Modeling Infrastructure for the Anthropocene. *Bulletin of the American Meteorological Society* URL: <https://journals.ametsoc.org/doi/abs/10.1175/BAMS-D-16-0236.1>.
- Ching, J., Aliaga, D., Mills, G., Masson, V., See, L., Neophytou, M., Middel, A., Baklanov, A., Ren, C., Ng, E., Fung, J., Wong, M., Huang, Y., Martilli, A., Brousse, O., Stewart, I., Zhang, X., Shehata, A., Miao, S., Wang, X., Wang, W., Yamagata, Y., Duarte, D., Li, Y., Feddema, J., Bechtel, B., Hidalgo, J., Roustani, Y., Kim, Y., Simon, H., Kropp, T., Bruse, M., Lindberg, F., Grimmond, S., Demuzere, M., Chen, F., Li, C., Gonzales-cruz, J., Bornstein, B., He, Q., Hanna, A., Erell, E., Tapper, N., Mall, R.K., Niyogi, D., 2019. Pathway using WUDAPT's digital Synthetic City tool towards generating urban canopy parameters for multi-scale urban atmospheric modeling. *Urban Climate* 28 100459. <https://www.sciencedirect.com/science/article/abs/pii/S2212095519300975>.
- Christoudias, C.M., Georgescu, B., Meer, P., 2002. Synergism in low level vision. In: *Proceedings Of The Sixteenth International Conference On Pattern Recognition*

- (Icpr 2002) Held In Quebec City, Canada, August 11-15, 2002. IEEE Press, New York, pp. 150–155. NY. No. 2. <https://ieeexplore.ieee.org/document/1047421>.
- Comaniciu, D., Meer, P., 1997. Robust analysis of feature spaces: color image segmentation. In: Proceedings Of The 1997 Conference On Computer Vision And Pattern Recognition (Cvpr '97) Held In San Juan, Puerto Rico, June 17-19, 1997. IEEE Press, New York, pp. 750–755. <https://ieeexplore.ieee.org/document/609410>.
- Comaniciu, D., Meer, P., 2002. Mean shift: a robust approach toward feature space analysis. *IEEE Trans. Pattern Anal. Mach. Intell.* 24 (5), 603–619. <https://ieeexplore.ieee.org/document/1000236/>.
- Ettinger, S.M., Nechyba, M.C., Waszak, M., Ifju, P.G., 2003. Vision-guided flight stability and control for micro air vehicles. *Adv. Robot.* 17 (7), 617–640.
- Gallagher, A.C., Luo, J., Hao, W., 2004. Improved blue sky detection using polynomial model fit. In: 2004 International Conference on Image Processing, 2004. ICIP '04, pp. 2367–2370.
- GIMP, 2019. GNU Image Manipulation Program. <https://www.gimp.org/>.
- Gong, F.-Y., Zeng, Z.-C., Zhang, F., Li, X., Ng, E., Norford, L.K., 2018. Mapping sky, tree, and building view factors of street canyons in a high-density urban environment. *Build. Environ.* 134, 155–167. <https://doi.org/10.1016/j.buildenv.2018.02.042>.
- Grimmond, C., Potter, S., Zutter, H., Souch, C., 2001. Rapid methods to estimate sky-view factors applied to urban areas. *Int. J. Climatol.* 21 (7), 903–913. <http://doi.wiley.com/10.1002/joc.659>.
- Hassannejad, H., Matrella, G., Ciampolini, P., Mordonini, M., Cagnoni, S., 2016. Food image recognition using very deep convolutional networks. In: *MADiMa'16*, October 16 2016. Amsterdam, The Netherlands, pp. 41–49.
- Hoiem, D., Efron, A.A., Hebert, M., 2005. Geometric context from a single image. In: Tenth IEEE International Conference on Computer Vision (ICCV'05) Volume 1, URL. <https://ieeexplore.ieee.org/document/1541316>.
- Holder, C.J., Breckon, T.P., Wei, X., 2016. From on-road to off: transfer learning within a deep convolutional neural network for segmentation and classification of off-road scenes. In: European Conference on Computer Vision-ECCV 2016: Computer Vision ECCV 2016 Workshops, pp. 149–162.
- Google, 2019. Advanced Guide to Inception v3 on Cloud TPU. <https://cloud.google.com/tpu/docs/inception-v3-advanced>.
- Oracle, 2018. Java Platform Standard Edition 8 Documentation. <https://docs.oracle.com/javase/8/docs/index.html>.
- Laffont, P.-Y., Ren, Z., Tao, X., Qian, C., Hays, J., 2014. Transient attributes for high-level understanding and editing of outdoor scenes. *ACM Transactions On Graphics* 33 (4) URL. <https://dl.acm.org/citation.cfm?id=2601101>.
- Laungrunthip, N., McKinnon, A.E., Churcher, C.D., Unsworth, K., 2008. Edge-based detection of sky regions in images for solar exposure prediction. In: 2008 23rd International Conference Image and Vision Computing New Zealand, IVCNZ.
- Liang, J., Gong, J., Sun, J., Zhou, J., Li, W., Li, Y., Liu, J., Shen, S., 2017. Automatic sky view factor estimation from street view photographs—a big data approach. *Remote Sensing* 9 (5), 411. URL. <http://www.mdpi.com/2072-4292/9/5/411>.
- Luo, J., Eitz, S.P., 2002. A physical model-based approach to detecting sky in photographic images. *IEEE Trans. Image Process.* 11 (3), 201–212.
- Maps, Google, 2017. Google Street View API. Available from: <https://developers.google.com/maps/documentation/streetview/>, Accessed date: 15 June 2017.
- Middel, A., Lukaszczuk, J., Maciejewski, R., 2017. Sky view factors from synthetic fisheye photos for thermal comfort Routing-A case study in Phoenix, Arizona. *Urban Planning* 2 (1), 19. <http://www.cogitatiopress.com/urbanplanning/article/view/855>.
- Middel, A., Lukaszczuk, J., Maciejewski, R., Demuzere, M., Roth, M., 2018. Sky view factor footprints for urban climate modeling. *Urban Climate* 25, 120–134. <https://www.sciencedirect.com/science/article/pii/S2212095518301883>.
- Middel, A., Lukaszczuk, J., Zakrzewski, S., Arnold, M., Maciejewski, R., 2019. Urban form and composition of street canyons: a human-centric big data and deep learning approach. *Landscape and Urban Planning* 183, 122–132. <https://doi.org/10.1016/j.landurbplan.2018.12.001>.
- Mihail, R.P., Workman, S., Bessinger, Z., Jacobs, N., 2015. Skyfinder Dataset. 2016 In: Sky Segmentation in the Wild: An Empirical Study. 2016 IEEE Winter Conference on Applications of Computer Vision, WACV 2016. <https://ieeexplore.ieee.org/document/7477637/>.
- Mihail, R.P., Workman, S., Bessinger, Z., Jacobs, N., 2016. Sky segmentation in the wild: An empirical study. 2016 IEEE Winter Conference on Applications of Computer Vision, WACV 2016. <https://ieeexplore.ieee.org/document/7477637/>.
- Mills, K., Bechtel, B., Ching, J., See, L., Feddema, J., Foley, M., Alexander, P., Connor, M. O., 2015. An Introduction to the WUDAPT project. In: ICUC9 - 9th International Conference on Urban Climate Jointly with 12th Symposium on the Urban Environment. URL http://www.wudapt.org/wp-content/uploads/2015/05/Mills_et_al_ICUC9.Pdf.
- Nice, K.A., 2019. Code Repository For 'Sky Pixel Detection In Outdoor Imagery Using An Adaptive Algorithm And Machine Learning'. <https://bitbucket.org/politemadness/skypixeldetection>.
- Pangburn, B.E., 2002. Experience-Based Language Acquisition: A Computational Model of Human Language Acquisition. Phd. Louisiana State University.
- Roser, M., Moosmann, F., 2008. Classification of weather situations on single color images. In: 2008 IEEE Intelligent Vehicles Symposium, pp. 798–803. <https://ieeexplore.ieee.org/abstract/document/4621205>.
- Rueden, C.T., Schindelin, J., Hiner, M.C., DeZonia, B.E., Walter, A.E., Arena, E.T., Eliceiri, K.W., 2017. ImageJ2: ImageJ for the next generation of scientific image data. *BMC Bioinformatics* 18 (529), 1–26. URL. <https://bmcbioinformatics.biomedcentral.com/articles/10.1186/s12859-017-1934-z>.
- Schmitt, F., Priese, L., 2009. Sky detection in CSC-segmented color images. In: VISAPP 2009 - Proceedings of the Fourth International Conference on Computer Vision Theory and Applications, Lisboa, Portugal, February 5-8, 2009 - Volume 2.
- Shen, Y., Wang, Q., 2013. Sky region detection in a single image for autonomous ground robot navigation. *Int. J. Adv. Robot. Syst.* 10, 1–13.
- Shen, Q., Zeng, W., Ye, Y., Stefan, M., Schubiger, S., Burkhard, R., Qu, H., 2018. StreetVizor: visual exploration of human-scale urban forms based on street views. *IEEE Trans. Vis. Comput. Graph.* 24 (1), 1004–1013. <https://ieeexplore.ieee.org/abstract/document/8017655>.
- Sobel, I., Feldman, G., 1968. A 3x3 isotropic gradient operator for image processing. In: Stanford Artificial Intelligence Project (SAIL).
- Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., Rabinovich, A., 2015. Going deeper with convolutions. In: Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition 07-12-June, pp. 1–9.
- Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., Wojna, Z., 2016. Rethinking the Inception Architecture for Computer Vision. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR).
- Tao, L., Yuan, L., Sun, J., 2009. SkyFinder: attribute-based sky image search. *ACM Trans. Graph., Siggraph* 28, 68.
- Tighe, J., Lazebnik, S., 2010. SuperParsing: scalable nonparametric image parsing with Superpixels. In: European Conference on Computer Vision, ECCV 2010: Computer Vision ECCV 2010, pp. 352–365.
- Tsai, Y.-H., Shen, X., Lin, Z., Sunkavalli, K., Yang, M.-H., 2016. Sky is not the limit: semantic-aware sky replacement. *ACM Trans. Graph.* 35 (4), 149. URL. <https://dl.acm.org/citation.cfm?id=2925942>.
- Wang, C.-W., Ding, J.-J., Chen, P.-J., 2015. An efficient sky detection algorithm based on hybrid probability model. In: Proceedings of APSIPA Annual Summit and Conference 2015, pp. 919–922.
- Willmott, C.J., 1981. On the validation of models. *Phys. Geogr.* 2 (2), 184–194.
- Xia, X., Xu, C., Nan, B., 2017. Inception-v3 for flower classification. In: 2017 2nd International Conference on Image, Vision and Computing (ICIVC). IEEE, pp. 783–787.
- Yu, D., Eversole, A., Seltzer, M.L., Yao, K., Huang, Z., Guenter, B., Kuchaiev, O., Zhang, Y., Seide, F., Wang, H., Droppo, J., Zweig, G., Rossbach, C., Currey, J., Gao, J., May, A., Peng, B., Stolcke, A., Slaney, M., 2015. An Introduction To Computational Networks And The Computational Network Toolkit. Microsoft Technical Report MSR-TR-2014-112. Tech. rep., Microsoft.
- Zafarifar, B., de With, P.H.N., 2007. Blue sky detection for content-based television picture quality enhancement. In: 2007 Digest of Technical Papers International Conference on Consumer Electronics.
- Zhijie, Z., Qian, W., Huadong, S., Xuesong, J., 2014. An effective method for sky region detection. *Adv. Sci. Technol. Lett.* 77, 96–100.
- Zhijie, Z., Qian, W., Huadong, S., Xuesong, J., Qin, T., Xiaoying, S., 2015. A novel sky region detection algorithm based on border points. *Int. J. Signal Proc., Image Proc. Pattern Recognit.* 8 (3), 281–290.