

# Urban Perception: Can we understand why a street is safe?

Felipe Moreno-Vera<sup>1</sup>[0000–0002–2477–9624], Bahram Lavi<sup>2</sup>[0000–0002–9226–9116],  
and Jorge Poco<sup>2</sup>[0000–0001–9096–6287]  
felipe.moreno@ucsp.edu.pe, {bahram.lavi, jorge.poco}@fgv.br

<sup>1</sup> Universidad Católica San Pablo, Arequipa, Perú

<sup>2</sup> Fundação Getulio Vargas, Rio de Janeiro, Brazil

**Abstract.** The importance of urban perception computing is relatively growing in machine learning, particularly in related areas to Urban Planning and Urban Computing. This field of study focuses on developing systems to analyze and map discriminant characteristics that might directly impact the city’s perception. In other words, it seeks to identify and extract discriminant components to define the behavior of a city’s perception. This work will perform a street-level analysis to understand safety perception based on the “visual components”. As our result, we present our experimental evaluation regarding the influence and impact of those visual components on the safety criteria and further discuss how to properly choose confidence on safe or unsafe measures concerning the perceptual scores on the city street levels analysis.

**Keywords:** Urban Perception · Urban Computing · Interpretability · LIME · Computer Vision · Perception Computing · Deep Learning · Street-level imagery · Visual processing · Street View · Cityscape · ADE20K · Place Pulse · Perception Learning · segmentation.

## 1 Introduction

“Cities are designed to shape and influence the lives of their inhabitants” [14]. Various studies have shown that the visual appearance of cities plays a key role in human perception that could cause variant reactions (e.g., abnormality) in the city’s environments, such as “The image of the city” [16]. A notable example is the Broken Window Theory [40] which delivers that visual signs of environmental disruption, such as broken windows, abandoned cars, trash, and graffiti, can induce social outcomes like an increase in crime levels. This theory has greatly influenced on public policy makers that lead to aggressive police tactics to control the manifestations of social and physical disorders. For example, in social experiments and studies on the perceived quality of life in the streets of New York, [11] reports the high correlation between graffiti or garbage presence and dangerous places. On the other hand, clean places present high correlation with safety places. Similar results were reported by [31, 38, 16] concluding that in places where “the rules are violated”, none of the rules will be fulfilled in

that place negatively influenced by the environment (e.g. graffiti, garbage). In addition, other studies have shown that the visual aspect of environments of a city affects the psychological state of its inhabitants [14, 10]; Other studies show that the impact of green areas in urban cities has a positives relation to safety perception [39, 30, 13].

In this study, we present a methodology to analyze the influence of objects on a street image by taking into account their corresponding perceptual scores. Furthermore, we investigate machine learning techniques to alleviate the relationship between urban visual components and their perceptions score.

This paper is organized as follows: Section 2 explains the related works; Section 3 introduces our methodology in perception score analysis; Section 4 presents our experiments and discussion about our achieved results by providing some signs over the limitations on this research field; and finally, Section 5 concludes our work.

## 2 Related Works

Previous works have difficulty explaining the direct relation between the visual appearance of a city and its corresponding non-visual attributes. Therefore, these works focused on finding the relation between the data from police records and census statistics (e.g., robbery rate, house prices, population density, graffiti existence) with the the visual appearance of a city area. In the following, we will highlight and discuss them in details.

### 2.1 Urban Perception

Some studies have addressed urban perception analysis by examining different methods in computation and extracting knowledge from various resources (visual and non-visual components) – aiming to seek a proper correlation among them. The work proposed in [7] attempts to address the key question on the appearance of the Paris, "What makes Paris look like Paris?". The work was developed to compare, differentiate, and correlate the visual representation between 12 cities. Similarly, the work in [25] addressed another proposal as "What Makes London Look Beautiful, Quiet, and Happy?", which explores nearly to 700,000 street images through an online web survey. In[3], the work studied the correlation between non-visual-attributes from the city along with its visual appearance using some datasets containing the data from crimes statistics, robbery rate, house pricing rate, population density, graffiti presence, and a perception survey.

In addition, MIT Media Lab releases the PlacePulse dataset [28] which is composed of images over different streets from many capital cities like New York, Boston, Linz, and Salzburg. They also provided the associated perceptual scores for each of the images.

This work was born from the attempt to relate people’s perception of a street through an online survey. This dataset conducted new studies for the problems like urban mapping [23] which performs as a classification/regression task to

compare the performance of features extractors like Gist, SIFT+Fisher Vectors, and DeCAF [8]. In [20], a StreetScore approach proposed to compare a set of low-level features such as GIST, Geometric Probability Map, Text on Histograms, Color Histograms, Geometric Color Histograms, HOG 2x2, Dense SIFT, LBP, Sparse SIFT histograms, and SSIM features extractors doing a similar research on urban perception analysis. Following their methodology, a similar study was performed over the city of Bogotá, Colombia [1].

In summary, these works have difficulty in extracting information about the natural image because they use traditional image representations including Hog+Color descriptor, Locality-Sensitive Hashing, Gist, HOG+color [3], SIFT Fisher Vectors, DeCAF features [23], geometric classification map, color Histograms, HOG2x2, and Dense SIFT [20]. Besides, to train those features non-linear methods are used like SVM [4], Linear Regression [23], SVR [20], RankingSVM [24], Multi Task Learning [17], Transfer Learning based models and pre-trained networks in [21, 9, 41, 13, 18, 12, 1].

## 2.2 Model Interpretation

Model interpretation methods allow us to get insights and understand the behavior of the learning model in its training phase. In the line with those methods, there are several works whose purpose is to understand and explain predictions. Previous works such as LIME [26], SHAP [15], and Anchor [27] explain a model based on their local and global level feature components. Other approach based on gradient attribution methods used to generate feature maps of an input to provide a visual idea about the explanation like Saliency Maps [35], Gradient [34], Integrated Gradients [37], DeepLIFT [33], Grad-CAM [32], Guided Back Propagation [36], Guided gradCAM[32], and SmoothGrad [2]. These methods can ease and assist us in explaining simple/complex models aiming to identify the dependence of variables and determine whether one of them can be isolated or not; to ascertain which one has a better representation for prediction depending on the input sample.

In this work, our primary goal is to understand the behavior of urban perception. First, we extract the objects segmented from Place Pulse images denominating our visual components. Then, we train those components, considering them as a feature vector using the standard classifiers: SVC model with RBF and Linear kernels, Logistic Regression, and Ridge Classifier. Finally, we aim to understand the impact of the visual component representations by adopting the LIME (black-box) method to analyze the behavior of the predictions.

## 3 Methodology

In this section, we describe our methodology in urban perception analysis. Our main goal is to analyze the key components that mainly affect a security perception, such as safety. To this end, we first explain our utilized urban datasets, mainly they contain the perceptual scores for the safety criteria. Then we

explain a possible solution for learning the key components to understand the importance of their presence in the images, which they have assigned as safe or unsafe.

### 3.1 Datasets and Data Pre-Processing

PlacePulse has two versions. The first one is Placepulse 1.0 which composed by a set of street views images and provides their corresponding perceptual scores [28]. At the end of 2013, Place Pulse 1.0 was organized with a total of 73,806 comparisons of 4,109 images from 4 cities: New York City (including Manhattan and parts of Queens, Brooklyn and The Bronx), Boston (including parts of Cambridge), Linz and Salzburg of two countries (US and Austria) and three types of comparisons: *safe*, *wealth*, y *unique*. This dataset has been pre-processed for quick use, containing information on the position of each image (latitude and longitude), perception score for each category, an image identifier and the city to which said image belongs.

The second dataset is PlacePulse 2.0 [9] that contains a set of comparisons between image pairs, and include the latitude and longitude points for each. In addition, each comparison has the respective winner (or draw). In 2016, Place Pulse 2.0 already contained around 1.22 million comparisons of 111,390 images of 56 cities in 28 countries across the 5 continents and six types of comparisons: *safe*, *wealth*, *depress*, *beautiful*, *boring*, and *lively*. This dataset contain 8 columns: image ID (left and right), latitude and longitude (of each image), the result of the comparison, and the respective evaluated category.

Place Pulse 1.0				
City	# images	<i>safe mean</i>	<i>wealth mean</i>	<i>unique mean</i>
Linz	650	4.85	5.01	4.83
Boston	1237	4.93	4.97	4.76
New York	1705	4.47	4.31	4.46
Salzburg	544	4.75	4.89	5.04
Total	4136			

**Table 1.** Data summary about Place Pulse 1.0 and their respective category mean.

We perform the method proposed by [29] to pre-process all comparisons in the dataset: for each compared image  $i$  with other images  $j$  many times in different categories, we define the intensity of perception of any image  $i$  as the percentage of times that the image was selected. Besides, the intensity of  $j$  affects  $i$  intensity. Due to this, we define the positive rate  $W_i$  (1) and the negative rate  $L_i$  (2) of an image  $i$  corresponding to a specific category:

$$W_i = \frac{w_i}{w_i + d_i + l_i} \tag{1}$$

$$L_i = \frac{l_i}{w_i + d_i + l_i} \quad (2)$$

Where,  $w_i$  is the number of wins,  $l_i$  number of loses, and  $d_i$  draws; From the equations 1 and 2 we can calculate the perceptual score associated for each an image  $i$  called  $Q$ -score with notation  $q_{i,k}$  in a category  $k$ :

$$q_{i,k} = \frac{10}{3} \left( W_{i,k} + \frac{1}{n_{i,k}^w} \left( \sum_{j_1} W_{j_1,k} \right) - \frac{1}{n_{i,k}^l} \left( \sum_{j_2} L_{j_2,k} \right) + 1 \right) \quad (3)$$

The Equation 3 is the perceptual score of the image  $i$  to be ranked, where  $j$  is an image compared to  $i$ ,  $n_i^w$  is equal to the total number of images  $i$  beat and  $n_i^l$  is equal to the total number of images to which  $i$  lost. Besides,  $j_1$  is the set of images that loses against the image  $i$  and  $j_2$  is the set of images that wins against the image  $i$ . Finally,  $Q$  is normalized to fit the range 0 to 10; this scale is a standard measurement whereby one can evaluate the perceptions [22]. In this score, 10 represents the highest possible score for a given question. For example, if an image receives a calculated score of 0 for the question “Which place looks safer?” indicating that specific image is perceived as the least safe image in the dataset.

Place Pulse 2.0		
Continent	# cities	# images
America	22	50,028
Europe	22	38,747
Asia	7	11,417
Oceania	2	6,097
Africa	3	5,101
Total	56	111,390

(a)

Place Pulse 2.0		
Category	# comparisons	mean
<i>Safety</i>	368,926	5.188
<i>Lively</i>	267,292	5.085
<i>Beautiful</i>	175,361	4.920
<i>Wealthy</i>	152,241	4.890
<i>Depressing</i>	132,467	4.816
<i>Boring</i>	127,362	4.810
Total	1,223,649	

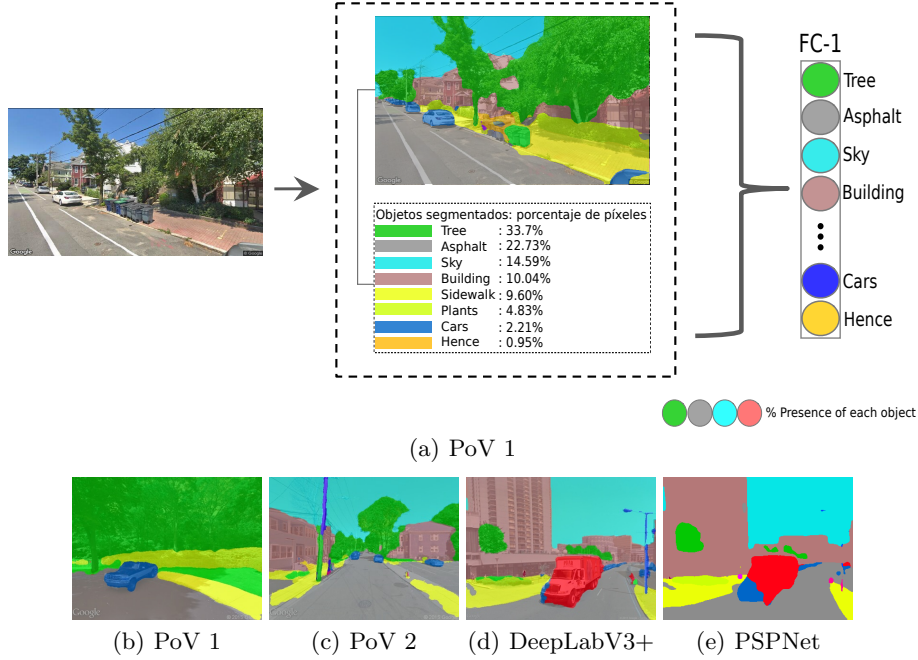
(b)

**Table 2.** Statistics obtained after process all comparisons from Place Pulse, containing information about images per city in each continent and the mean score for each requested category.

### 3.2 Visual Components Extraction

In this work, we will focus on Boston city for our experiments. We use two segmentation Network: (i) PSPNet [42] and (ii) DeepLabV3+ [5]. We define as “visual components” the object pixel presence extracted from (i) and (ii) per image of the city. As we show in Figure 1, our main idea is to transform the percentage of objects present in each image. As we know, different images could present different % pixel ratio segmented (see Figure 1 (a), (b), (c), and (d), we

perform a Standardization of all features obtained by segmentation. We perform both extractions to compare these methods using the PSPNet as the baseline.



**Fig. 1.** (a) Input image and output features (based on the object pixel ratio segmented). (a) y (b): Different Point of View of the images, each object presence will depend on the image evaluated. (c) y (d): Different pixel ratios extracted by DeepLabV3+ and PSP-Net. As you can see in (d) DeepLabV3+ detects the street light, but failed to detect by PSPNet.

Both object segmentation extractors were trained in ADE20K [43], using networks like ResNet101 and Xception as backbone, respectively. We prefer to use the ADE20K dataset pre-trained weights instead of Pascal, COCO, or CityScapes [6], due to the number of classes between the datasets. ADE20K presents 150 classes and a hierarchical tree of indoor-outdoor classes, CityScapes provides 50 classes (most of them contained in ADE20K), COCO 91 classes, and Pascal-VOC 20 classes.

Next, after extracting our features, we train them using a Support Vector Classifier with 10 KFold cross-validations varying our regularization parameter  $l_2$ . To perform our classification task to predict if a street image is safe or not safe. In order to classify, we need to select subsets from each city dataset. To do this, we define a parameter called  $\delta$  with a value between 0.05 - 0.5. This delta will create a subset using the binary labels  $y_{i,k} \in \{1, -1\}$  for both training and testing as:

$$y_{i,k} = \begin{cases} 1 & \text{if } (q_{i,k}) \text{ in the top } \delta\% \\ -1 & \text{if } (q_{i,k}) \text{ in the bottom } \delta\% \end{cases} \quad (4)$$


Since we know from previous works results [23, 28, 19], we focus our study on the worst case reported, corresponding to  $\delta = 0.5$  using all labels divided into safe and not safe.

## 4 Experiments and Discussions

This work presents a methodology to learn and explain which features have more impact on the prediction of safe or not safe categories. We perform our experiments in Boston city (1327 images) from Place Pulse 2.0 dataset. We focus only in the safety perception due to the larger number of image compared in that category per city (see Table 2 (b) ). We extract our visual components using both methods as mentioned above (PSP-Net, and DeepLabV3+). In order to visualize our distribution of visual components, we perform a mini-process in our pixel ratio extracted for each feature  $X_{i,k}$ :

$$X_{i,k} = \begin{cases} 1 & \text{if object } (k) \text{ is present in image } X_i \\ 0 & \text{if object } (k) \text{ is not present in image } X_i \end{cases} \quad (5)$$

Then, sum by column and divided by the total number of images, we got an object distribution presence in the whole Boston city for each method (see Figure 2). Then, we train both visual components extracted by both methods (PSP-Net and DeepLabV3), we can see that both features yielded a poor performance on the dataset (see Table 3).



**Fig. 2.** (a) DeepLabV3+ object distribution presence. (b) PSPNet Object distribution presence. As we can see, there are objects detected by only DeepLabV3 method.

### Model Explanation

In this work, we want to understand why our street images are predicted as “safe” or “not safe”. To do this, we use the black-box model explainer LIME: Local Interpretable Model-agnostic technique. LIME explains a black-box model by simulating local candidates close to the original prediction. By using these prediction outcomes, LIME generates a random distribution set of possible predictions based on  $L_2$  distance called “local fidelity” taken as a reference to the original prediction.

First, based on the results of training using different models like Logistic Regression, Ridge Classifier, SVC with kernel Linear and RBF presented in the Table 3. Then, we choose the Logistic Regression model to analyze the feature importance and influence. We started analyzing the object presence in both subsets divided by safe and not safe categories.


Features	Metric	RBF-SVC	Linear SVC	Logistic Regression	Ridge
PSPNet	<i>AUC</i>	0.46465	0.47036	0.465	0.48551
	<i>ACC</i>	0.44516	0.48065	0.47097	0.48387
	<i>F1</i>	0.42282	0.5752	0.50602	0.47712
DeepLabV3	<i>AUC</i>	<b>0.55553</b>	0.51255	0.51895	<b>0.56066</b>
	<i>ACC</i>	0.5129	0.50323	<b>0.52258</b>	0.51935
	<i>F1</i>	0.47018	<b>0.59043</b>	0.53459	0.52396

**Table 3.** We report Test classification for  $\delta = 0,5$  (Worst case) for each feature extractor method. We note that DeepLab and PSPNet have a poor performance in this task, but DeepLab present a better learning process.

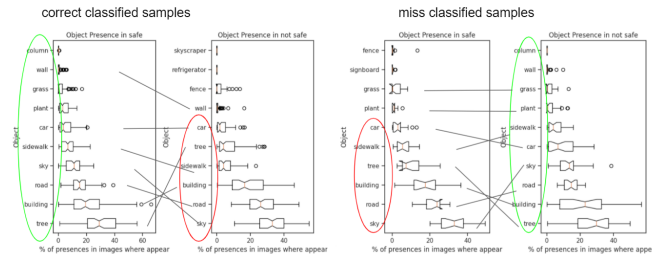
Second, we analyze the Feature Presence in whole images and the respective subsets corresponding to safety and not safety. At this step, we divided in two sub-subsets: miss-classified and correct classified subsets (see Figure 3). Then, we perform the following Feature Importance, Permutation Importance in the whole dataset and the subsets divided by safe and not safe categories. In Figure 3-(a), we show the global and divided object presence, we can see that the first four objects in safety are the reverse of not safety. In Figure 3-(b), we note that object presence in miss-classified samples correspond more to the opposite class (e.g., miss-classified as safe that has presented like not safe), these regions are highlighted in red and green: red correspond to not safe class, which is present in miss-classified safe images, the same happens with miss-classified not safety highlighted in green. In Figure 3 (c) we present the object influence, this report obtained by the methods mentioned above. The results show the Influence of permute features or visual components and the relevance in the predictions, as the component ‘tree’. The feature of ‘tree’ component has higher permutation importance in not safe class.

Our last step is to use the model explainer LIME to determine the object importance of predictions. This step shows the main results of LIME. The most






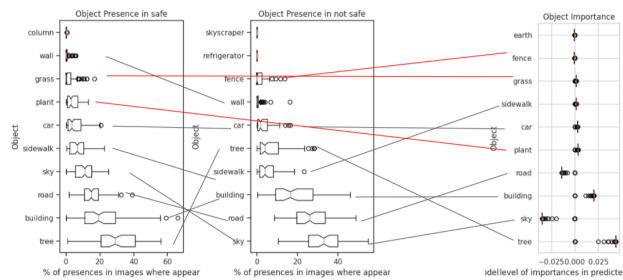
(a) Presence: General Overview



(b) Comparison classified



(c) Permutations: Object Influence



(d) LIME: Object Importance

**Fig. 3.** (a) General object presence in the whole dataset divided by safe and not safe categories. (b) sub-subsets of correct/miss classified. (c) Object Influence after permutations on features. (d) Object importance calculated by LIME, red lines mean the match between presence in safety/not safety and the importance in predictions.

highlighted ones are tree, sky, building, and road which are the most discriminant features to the predictions method in comparison with cars and sidewalk. This can be explained due to the high frequently appearance of these objects in all the images. Besides, objects like grass, plants, earth, and fence are very related to a particular category. In safety class, we have grass and plants, while in not safe class, it observes more with earth and fence components.

**Limitations** : We found three main limitations in this work. The first one is about the Place Pulse dataset that constructed using an online survey. Each volunteer chose between two images that are the most "safe" depending on their biased personal perception criteria. The second limitation is the small number of sample images per city. Comparing with other dataset with millions of samples, in total is not above of 100,000 that yields the method to have week performance when a lower number of data samples are available. The last limitation is the impracticality of creating a general city perceptual predictor due to the significant difference between cities and their unique visual appearance.

## 5 Conclusions

In this work, we propose a methodology that allows us to understand the behavior of the urban safety perception on street view images. To this end, we pre-processed the dataset Place Pulse 2.0, analyzing the 110 thousand images obtained by comparisons and calculated their corresponding perception scores in six different categories. In our study, we focused on Boston city with its safety scores. We investigated and analyzed which visual components are impacting positively and negatively in the predictions. To understand the predictions, we used LIME to determine the importance of feature components. From the result, we conclude that our model is capable to predict the safety perception from street view images. In addition, we showed the correlation between higher safety perception with the presence of trees or green areas, skies, and roads.

## References

1. Acosta, S.F., Camargo, J.E.: Predicting city safety perception based on visual image content. In: Iberoamerican Congress on Pattern Recognition. pp. 177–185. Springer (2018)
2. Ancona, M., Ceolini, E., Öztireli, C., Gross, M.: A unified view of gradient-based attribution methods for deep neural networks. ETH Zurich (2017)
3. Arietta, S.M., Efros, A.A., Ramamoorthi, R., Agrawala, M.: City forensics: Using visual elements to predict non-visual city attributes. *IEEE transactions on visualization and computer graphics* **20**(12), 2624–2633 (2014)
4. Boser, B.E., Guyon, I.M., Vapnik, V.N.: A training algorithm for optimal margin classifiers. In: Proceedings of the fifth annual workshop on Computational learning theory. pp. 144–152. ACM (1992)

5. Chen, L.C., Zhu, Y., Papandreou, G., Schroff, F., Adam, H.: Encoder-decoder with atrous separable convolution for semantic image segmentation. In: Proceedings of the European conference on computer vision (ECCV). pp. 801–818 (2018)
6. Cordts, M., Omran, M., Ramos, S., Rehfeld, T., Enzweiler, M., Benenson, R., Franke, U., Roth, S., Schiele, B.: The cityscapes dataset for semantic urban scene understanding. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 3213–3223 (2016)
7. Doersch, C., Singh, S., Gupta, A., Sivic, J., Efros, A.: What makes paris look like paris? (2012)
8. Donahue, J., Jia, Y., Vinyals, O., Hoffman, J., Zhang, N., Tzeng, E., Darrell, T.: Decaf: A deep convolutional activation feature for generic visual recognition. In: International conference on machine learning. pp. 647–655 (2014)
9. Dubey, A., Naik, N., Parikh, D., Raskar, R., Hidalgo, C.A.: Deep learning the city : Quantifying urban perception at A global scale. CoRR (2016)
10. Kaplan, R., Kaplan, S.: The experience of nature: A psychological perspective. Cambridge university press (1989)
11. Keizer, K., Lindenberg, S., Steg, L.: The spreading of disorder. *Science* **322**(5908), 1681–1685 (2008)
12. León-Vera, L., Moreno-Vera, F.: Car monitoring system in apartments’ garages by small autonomous car using deep learning. In: Annual International Symposium on Information Management and Big Data. pp. 174–181. Springer, Springer International Publishing (2018)
13. Li, X., Zhang, C., Li, W., Ricard, R., Meng, Q., Zhang, W.: Assessing street-level urban greenery using google street view and a modified green view index. *Urban Forestry & Urban Greening* **14**(3), 675–685 (2015)
14. Lindal, P.J., Hartig, T.: Architectural variation, building height, and the restorative quality of urban residential streetscapes. *Journal of Environmental Psychology* **33**, 26–36 (2013)
15. Lundberg, S.M., Lee, S.I.: A unified approach to interpreting model predictions. In: Guyon, I., Luxburg, U.V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., Garnett, R. (eds.) *Advances in Neural Information Processing Systems* 30, pp. 4765–4774. Curran Associates, Inc. (2017)
16. Lynch, K.: Reconsidering the image of the city. In: *Cities of the Mind*, pp. 151–161. Springer (1984)
17. Min, W., Mei, S., Liu, L., Wang, Y., Jiang, S.: Multi-task deep relative attribute learning for visual urban perception. *IEEE Transactions on Image Processing* **29**, 657–669 (2019)
18. Moreno-Vera, F.: Performing deep recurrent double q-learning for atari games. In: 2019 IEEE Latin American Conference on Computational Intelligence (LA-CCI). pp. 1–4 (2019). <https://doi.org/10.1109/LA-CCI47412.2019.9036763>
19. Moreno-Vera, F.: Understanding safety based on urban perception. In: *International Conference on Intelligent Computing*. pp. 54–64. Springer (2021)
20. Naik, N., Philipoom, J., Raskar, R., Hidalgo, C.: StreetScore: predicting the perceived safety of one million streetscapes. 2014 IEEE Conference on Computer Vision and Pattern Recognition Workshops (2014)
21. Naik, N., Raskar, R., Hidalgo, C.A.: Cities are physical too: Using computer vision to measure the quality and impact of urban appearance. *American Economic Review* **106**(5), 128–32 (2016)
22. Nasar, J.L.: *The evaluative image of the city* (1998)
23. Odonez, V., Berg, T.L.: Learning high-level judgments of urban perception. *European Conference on Computer Vision (ECCV)* (2014)

24. Porzi, L., Rota Bulò, S., Lepri, B., Ricci, E.: Predicting and understanding urban perception with convolutional neural networks (10 2015)
25. Quercia, D., O’Hare, N.K., Cramer, H.: Aesthetic capital: what makes london look beautiful, quiet, and happy? In: Proceedings of the 17th ACM conference on Computer supported cooperative work & social computing (2014)
26. Ribeiro, M.T., Singh, S., Guestrin, C.: Why should i trust you?: Explaining the predictions of any classifier. In: Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining. pp. 1135–1144. ACM (2016)
27. Ribeiro, M.T., Singh, S., Guestrin, C.: Anchors: High-precision model-agnostic explanations. In: Thirty-Second AAAI Conference on Artificial Intelligence (2018)
28. Salesses, M.P.: Place Pulse: Measuring the collaborative image of the city. Ph.D. thesis, Massachusetts Institute of Technology (2012)
29. Salesses, P., Schechtner, K., Hidalgo, C.A.: The collaborative image of the city: Mapping the inequality of urban perception. PLOS ONE (2013)
30. Sampson, R.J., Morenoff, J.D., Gannon-Rowley, T.: Assessing “neighborhood effects”: Social processes and new directions in research. *Annual review of sociology* **28**(1), 443–478 (2002)
31. Schroeder, H.W., Anderson, L.M.: Perception of personal safety in urban recreation sites. *Journal of leisure research* **16**(2), 178–194 (1984)
32. Selvaraju, R.R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., Batra, D.: Grad-cam: Visual explanations from deep networks via gradient-based localization. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 618–626 (2017)
33. Shrikumar, A., Greenside, P., Kundaje, A.: Learning important features through propagating activation differences. arXiv preprint arXiv:1704.02685 (2017)
34. Shrikumar, A., Greenside, P., Shcherbina, A., Kundaje, A.: Not just a black box: Learning important features through propagating activation differences (2016)
35. Simonyan, K., Vedaldi, A., Zisserman, A.: Deep inside convolutional networks: Visualising image classification models and saliency maps (2013)
36. Springenberg, J.T., Dosovitskiy, A., Brox, T., Riedmiller, M.: Striving for simplicity: The all convolutional net. arXiv preprint arXiv:1412.6806 (2014)
37. Sundararajan, M., Taly, A., Yan, Q.: Axiomatic attribution for deep networks (2017)
38. Tokuda, E.K., Silva, C.T., Jr., R.M.C.: Quantifying the presence of graffiti in urban environments. CoRR **abs/1904.04336** (2019), <http://arxiv.org/abs/1904.04336>
39. Ulrich, R.S.: Visual landscapes and psychological well-being. *Landscape research* **4**(1), 17–23 (1979)
40. Wilson, J.Q., Kelling, G.L.: Broken windows. *Atlantic monthly* **249**(3), 29–38 (1982)
41. Zhang, F., Zhou, B., Liu, L., Liu, Y., Fung, H.H., Lin, H., Ratti, C.: Measuring human perceptions of a large-scale urban region using machine learning. *Landscape and Urban Planning* **180**, 148–160 (2018)
42. Zhao, H., Shi, J., Qi, X., Wang, X., Jia, J.: Pyramid scene parsing network. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2881–2890 (2017)
43. Zhou, B., Zhao, H., Puig, X., Xiao, T., Fidler, S., Barriuso, A., Torralba, A.: Semantic understanding of scenes through the ade20k dataset. *International Journal on Computer Vision* (2018)