

Deep Learning the City: Quantifying Urban Perception at a Global Scale

Abhimanyu Dubey¹, Nikhil Naik³(✉), Devi Parikh², Ramesh Raskar³,
and César A. Hidalgo³

¹ Indian Institute of Technology, Delhi, India
abhimanyu1401@gmail.com

² Virginia Tech, Blacksburg, USA
parikh@vt.edu

³ MIT Media Lab, Cambridge, USA
{naik,raskar,hidalgo}@mit.edu

Abstract. Computer vision methods that quantify the perception of urban environment are increasingly being used to study the relationship between a city’s physical appearance and the behavior and health of its residents. Yet, the throughput of current methods is too limited to quantify the perception of cities across the world. To tackle this challenge, we introduce a new crowdsourced dataset containing 110,988 images from 56 cities, and 1,170,000 pairwise comparisons provided by 81,630 online volunteers along six perceptual attributes: safe, lively, boring, wealthy, depressing, and beautiful. Using this data, we train a Siamese-like convolutional neural architecture, which learns from a joint classification and ranking loss, to predict human judgments of pairwise image comparisons. Our results show that crowdsourcing combined with neural networks can produce urban perception data at the global scale.

Keywords: Perception · Attributes · Street view · Crowdsourcing

1 Introduction

We shape our buildings, and thereafter our buildings shape us. – Winston Churchill.

These famous remarks reflect the widely-held belief among policymakers, urban planners and social scientists that the physical appearance of cities, and its perception, impacts the behavior and health of their residents. Based on this idea, major policy initiatives—such as the New York City “Quality of Life Program”—have been launched across the world to improve the appearance of cities. Social scientists have either predicted or found evidence for the impact of the perceived unsafety and disorderliness of cities on criminal behavior [1,2], education [3],

Electronic supplementary material The online version of this chapter (doi:[10.1007/978-3-319-46448-0_12](https://doi.org/10.1007/978-3-319-46448-0_12)) contains supplementary material, which is available to authorized users.

health [4], and mobility [5], among others. However, these studies have been limited to a few neighborhoods, or a handful of cities at most, due to a lack of quantified data on the perception of cities. Historically, social scientists have collected this data using field surveys [6]. In the past decade, a new source of data on urban appearance has emerged, in the form of “Street View” imagery. Street View has enabled researchers to conduct virtual audits of urban appearance, with the help of trained experts [7, 8] or crowdsourcing [9, 10].

However, field surveys, virtual audits and crowdsourced studies lack both the resolution and the scale to fully utilize the global corpus of Street View imagery. For instance, New York City alone has roughly one million street blocks, which makes generating an exhaustive city-wide dataset of urban appearance a daunting task. Naturally, generating urban appearance data through human efforts for hundreds of cities across the world, at several time points, and across different attributes (e.g., cleanliness, safety, beauty), remains impractical. The solution to this problem is to develop computer vision algorithms—trained with human-labeled data—that conduct automated surveys of the built environment at street-level resolution and global scale.

A notable example of this approach is Streetscore by Naik et al. [11]—a computer vision algorithm trained using Place Pulse 1.0 [9], a crowdsourced game. In Place Pulse 1.0, users are asked to select one of the two Street View images in response to question “Which place looks safer?”, “Which place looks more unique?”, and “Which places looks more upper class?”. This survey collected a total of 200,000 pairwise comparisons across the three attributes for 4,109 images from New York, Boston, Linz, and Salzburg. Naik et al. converted the pairwise comparisons for perceived safety to ranked scores and trained a regression algorithm using generic image features to predict the ranked score for perceived safety (also see the work by Ordonez and Berg [12] and Porzi et al. [13]). Streetscore was employed to automatically generate a dataset of urban appearance covering 21 U.S. cities [14], which has been used to identify the impact of historic preservation districts on urban appearance [15], for quantifying urban change using time-series street-level imagery [16], and to determine the effects of urban design on perceived safety [17].

Yet the Streetscore algorithm is not unboundedly scalable. Streetscore was trained using a dataset containing a few thousand images from New York and Boston, so it cannot accurately measure the perceived safety of images from cities outside of the Northeast and Midwest of United States, which may have different architecture styles and urban planning constructs. This limits our ability to generate a truly **global dataset of urban appearance**. Streetscore was also trained using a dataset with a relatively dense set of preferences (each image was involved in roughly 30 pairwise comparisons). But collecting such a dense set of preferences with crowdsourcing is challenging for a study that involves hundreds of thousands of images from several cities, and multiple attributes. So scaling up the computational methods to map urban appearance from the regional scale, to the global scale, requires methods that can be trained on larger and sparser datasets—which contain a large, visually diverse set of images with relatively few comparisons among them.

With the motivation of developing a global dataset of urban appearance, in this paper, we introduce a new crowdsourced dataset of urban appearance and a computer vision technique to rank street-level images for urban appearance in this paper. Our dataset, which we call the **Place Pulse 2.0** dataset, contains 1.17 million pairwise comparisons for 110,988 images from 56 cities from 28 countries across 6 continents, scored by 81,630 online volunteers, along six perceptual dimensions: safe, lively, boring, wealthy, depressing, and beautiful. We use the Place Pulse 2.0 (PP 2.0) dataset to train convolutional neural network models which are able to predict the pairwise comparisons for perceptual attributes by taking an image pair as input. We propose two related network architectures: (i) the Streetscore-CNN (**SS-CNN** for short) and (ii) the Ranking SS-CNN (**RSS-CNN**). The SS-CNN consists of two disjoint identical sets of layers with tied weights, followed by a fusion sub-network, which minimizes the classification loss on pairwise comparison prediction. The RSS-CNN includes an additional ranking sub-network, which tries to simultaneously minimize the loss on both pairwise classification and ordinal ranking over the dataset. The SS-CNN architecture—fine-tuned with the PP 2.0 dataset—significantly outperforms the same network architecture with pre-trained AlexNet [18], PlacesNet [19], or VGGNet [20] weights. RSS-CNN shows better prediction performance than SS-CNN, owing to end-to-end learning based on both classification and ranking loss. Moreover, our CNN architecture obtains much better performance over a geographically disparate test set when trained with PP 2.0, in comparison to PP 1.0, due to the larger size and visual diversity (110,988 images from 56 cities, versus 4,109 images from 4 cities).

We find that networks trained to predict one visual attribute (e.g., *Safe*), are fairly accurate in the prediction of other visual attributes (e.g., *Lively*, *Beautiful*, etc.). We also use a trained network to predict the perceived safety of streetscapes from 6 new cities from 6 continents, that were not part of the training set. Finally, we hope that this work and our publicly released dataset will enable further progress on global studies of the social and economic effects of architectural and urban planning choices.

2 Related Work

Our paper speaks to four different strands of the academic literature: (1) predicting perceptual responses to images, (2) using urban imagery to understand cities, (3) understanding the connection between urban appearance and socioeconomic outcomes, and (4) generating image rankings and comparisons.

There is a growing body of literature on **predicting the perceptual responses to images**, such as aesthetics [21], memorability [22], interestingness [23], and virality [24]. In particular, our work is related to the literature on predicting the perception of street-level imagery. Naik et al. [11] use generic image features and support vector regression to develop Streetscore, an algorithm that predicts the perceived safety of street-level images from United States, using training data from the Place Pulse 1.0 dataset [9]. Ordonez and Berg [12] use

the Place Pulse 1.0 dataset and report similar results for prediction of perceived safety, wealth, and uniqueness using Fisher vectors and DeCAF features [25]. Porzi et al. [13] identify the mid-level visual elements [26] that contribute to the perception of safety in the Place Pulse 1.0 dataset.

This new body of literature that utilizes **urban imagery to understand cities** has been enabled by new sources of data from both commercial providers (e.g., Google Street View) and photo-sharing websites (e.g., Flickr). These data sources have enabled applications for computer vision techniques in the fields of architecture, urban planning, urban economics and sociology. Doersch et al. [26] identify geographically distinctive visual elements from Street View data. Lee et al. [27] extend this work in the temporal domain by identifying architectural elements which are distinctive to specific historic periods. Arietta et al. [28] and Glaeser et al. [29] develop regression models based on Street View imagery to predict socioeconomic indicators. Zhou et al. [30] develop a unique city identity based on a high-level set of attributes derived from Flickr images. Khosla et al. [31] use Street View data and crowdsourcing to demonstrate that both humans and computers can navigate an unknown urban environment to locate businesses.

Our research also speaks to the more traditional stream of literature studying the connection between **urban appearance and socioeconomic outcomes** of urban residents, especially health and criminal behavioral. Researchers have studied the connection between the perception of unsafety and alcoholism [32], obesity [33], and the spread of STDs [4]. The influential “Broken Windows Theory (BWT)” [1] hypothesizes that criminal activity is more likely to occur in places that appear disorderly and visually unsafe. There has been a vigorous debate among scholars on BWT, who have found evidence in support [2, 34] and against the theory [35, 36]. Once again, this is another area where methods to quantify urban appearance may illuminate important questions.

Finally, our work is related to literature on **ranking and comparing images** based on both semantic and subjective attributes, or generating metrics for image comparisons. The concept of “relative attributes” [37]—ranking object/scene types according to different attributes—has been shown to be useful for applications such as image classification [38] and guided image search [39]. Kiapour et al. [40] rank images based on clothing styles using annotations collected from an online game, and generic image features. Zhu et al. [41] rank facial images for attractiveness, for generating better portrait images. Wang et al. [42] introduce a deep ranking method for image similarity metric computation. Zagoruyko and Komodakis [43] develop a Siamese architecture for computing image patch similarity for applications like wide-baseline stereo. Work on image perception summarized earlier [11–13] also ranks street-level images based on perceptual metrics.

In this paper, we contribute to these literatures by introducing a CNN-based technique to predict human judgments on urban appearance, using a global crowdsourced dataset.



(a) Snapshot of the game

(b) Distribution of #comparisons

Fig. 1. Using a crowdsourced online game (a), we collect 1.1 million pairwise comparisons on urban appearance from 81,630 volunteers. The distribution of number of pairwise comparisons contributed by players is shown in (b).

3 The Place Pulse 2.0 Dataset

Our first goal is to collect a crowdsourced dataset of perceptual attributes for street-level images. To create this dataset, we chose Google Street View images from 56 major cities from 28 countries spread across all six inhabited continents. We obtained the latitude-longitude values for locations in these cities using a uniform grid [44] of points calculated on top of polygons of city boundaries. We queried the Google Street View Image API¹ using the latitude-longitude values, and obtained a total of 110,988 images captured between years 2007 and 2012.

Following Salesses et al. [9], we created a web-interface (Fig. 1-(a)) for collecting pairwise comparisons from users. Studies have shown that gathering relative comparisons is a more efficient and accurate way of obtaining human rankings as compared to obtaining numerical scores from each user [45, 46]. In our implementation, we showed users a randomly-chosen pair of images side by side, and asked them to choose one in response to one of the six questions, preselected by the user. The questions were: “Which place looks safer?”, “Which place looks livelier?”, “Which place looks more boring?”, “Which place looks wealthier?”, “Which place looks more depressing?”, and “Which place looks more beautiful?”.

We generated traffic on our website primarily from organic media sources and by using Facebook advertisements targeted to English-speaking users who are interested in online games, architecture, cities, sociology, and urban planning. We collected a total of 1,169,078 pairwise comparisons from 81,630 online users between May 2013 and February 2016. The online users provided 16.6 comparisons on average. 6,118 users provided a single comparison each, while 30 users provided more than 1,000 comparisons (Fig. 1-(b)). The maximum number of comparisons provided by a single user was 7,168. We obtained the highest responses (370,134) for the question “Which place looks safer?”, and the lowest responses (111,184) for the question “Which place looks more boring?”. We attracted users from 162 countries (based on data from web analytics).

¹ <https://developers.google.com/maps/documentation/streetview/>.

Our user base contained a good mix of residents of both developed and developing countries. The top five countries of origin for these users were United States (31.4%), India (22.4%), United Kingdom (5.8%), Brazil (4.6%), and Canada (3.6%). It is worth noting that the Place Pulse 1.0 study found that individual preferences for urban appearance were not driven by participants’ age, gender, or location [9], indicating that there is no significant cultural bias in the dataset. Place Pulse 1.0 also found high inter-user reproducibility and high transitivity in people’s perception of urban appearance, which is indicative of consistency in data collected for this task. With that established, we did not collect demographic information from users for our much larger PP 2.0 dataset, but we did use the exact same data collection interface and user recruitment strategy as PP 1.0. Table 1 summarizes the key facts about the Global Perception dataset.

Table 1. The place pulse 2.0 dataset at a glance

(a) Statistics on Images			(b) Statistics on Pairwise Comparisons (PC)		
Continent	#Cities	#Images	Question	#PC	#Per-image PC
Asia	7	11,342	Safe	370,134	7.67
Africa	3	5,069	Lively	268,494	5.52
Australia	2	6,082	Beautiful	166,823	3.46
Europe	22	38,636	Wealthy	137,688	2.87
North America	15	33,691	Depressing	114,755	2.47
South America	7	16,168	Boring	111,184	2.40
Total	56	110,988	Total	1,169,078	16.73

4 Learning from the Place Pulse 2.0 Dataset

We now describe how we use the Place Pulse 2.0 dataset to train a neural network model to predict pairwise comparisons. Collecting pairwise comparisons has been the method of choice for learning subjective visual attributes such as style, perception, and taste. Examples include learning clothing styles [40], urban appearance [11], emotive responses to GIFs [47], or affective responses to paintings [48]. All these efforts use a two-step process for learning these subjective visual attributes—image ranking, followed by image classification/regression based on the visual attribute. In the first step, these methods [11, 40, 47, 48] convert the pairwise comparisons to ranked scores for images using the Microsoft TrueSkill [49] algorithm. TrueSkill is a Bayesian ranking method, which generates a ranked score for a player (in this case, an image) in a two-player game by iteratively updating the ranked score of players after every contest (in this case, a human-contributed pairwise comparison). Note that this approach for producing image rankings does not take image features into account. In the next step, the ranked scores, along with image features are used to train classification or regression algorithms, to predict the score of a previously unseen image.

However, this two-step process has a few limitations. First, for larger datasets, the number of crowdsourced pairwise comparisons required becomes quite large. TrueSkill needs 24 to 36 comparisons per image for obtaining stable rankings [49]. Therefore, we would require **~ 1.2 to 1.9 million comparisons** per question, to obtain stable TrueSkill scores for 110,988 images in the Place Pulse 2.0 dataset. This number is hard to achieve, even with the impressive number of users attracted by the Place Pulse game. Indeed, we are able to collect only 3.35 comparisons per image per question on average, after 33 months of data collection. Second, this two-step process ignores the visual content of images in the ranking process. We believe it is better to use visual content in the image ranking stage itself by learning to predict pairwise comparisons directly, which is similar in spirit to learning ranking functions for semantic attributes from image data [37] (also see Porzi et al. [13] for additional discussion on ranking versus regression). To address both problems, we propose to predict pairwise comparisons by training a neural network directly from image pairs and their crowdsourced comparisons from the Global Perception dataset. We describe the problem formulation and our neural network model next.

Problem Formulation: The Place Pulse 2.0 dataset consists of a set of m images $I = \{x_i\}_{i=1}^m \in \mathbb{R}^n$ in pixel-space and a set of N image comparison triplets $P = \{(i_k, j_k, y_k)\}_{k=1}^N$, $i, j \in \{1, \dots, m\}$, $y \in \{+1, -1\}$, which specify a pairwise comparison between the i th and the j th image in the set. $y = +1$ denotes a win for image i , and $y = -1$ denotes a win for image j . Our goal is to learn a ranking function $f_r(x)$ on the raw image pixels such that we satisfy the maximum number of constraints

$$y \cdot (f_r(x_i) - f_r(x_j)) > 0 \quad \forall (i, j, y) \in P \quad (1)$$

over the dataset. We aim to approximate a solution for this NP-hard problem [50] using a ranking approach, motivated by the direct adaptation of the RankSVM [50] formulation by Parikh and Grauman [37].

As the first step towards solving this problem, we transform the ranking task to a classification task. Specifically, our goal is to design a function which given an image pair, extracts low-level and mid-level features for each image as well as higher-level features discriminating the pair of images, and then predicts a winner. We next describe a convolutional neural network architecture which learns such a function.

4.1 Streetscore-CNN

We design the Streetscore-CNN (**SS-CNN**) for predicting the winner in a pairwise comparison task, by taking an image pair as input (Fig. 2). SS-CNN consists of two disjoint identical sets of layers with tied weights for feature extraction (similar to a Siamese network [51]). These feature extractor layers are concatenated and followed by a *fusion* sub-network, which consists of a set of convolutional layers culminating in a fully-connected layer with softmax loss used to

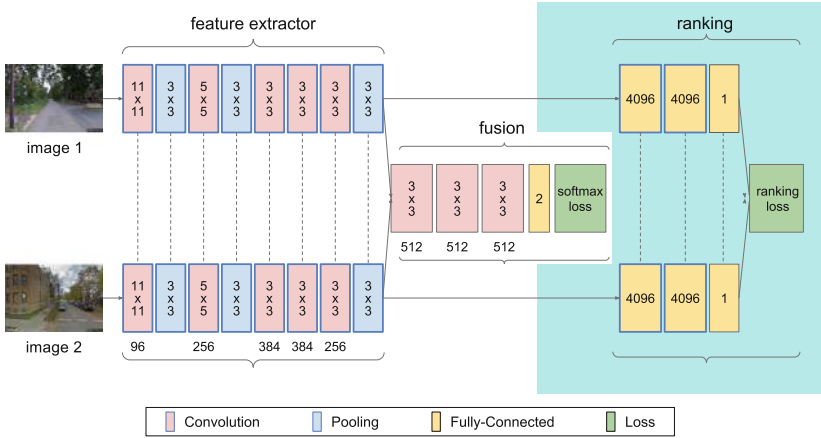


Fig. 2. We introduce two networks architectures, based on the Siamese model, for predicting pairwise comparisons of urban appearance. The basic model (SS-CNN) is trained with softmax loss in the fusion layer. We also introduce ranking loss layers to train the Ranking SS-CNN (additional layers shown in light blue background). While we experiment with AlexNet, PlacesNet, and VGGNet, this figure shows the AlexNet configuration. (Color figure online)

train the network. The fusion sub-network was inspired by the *temporal fusion* architecture [52] used to learn temporal features from video frames. The temporal fusion architecture learns convolutional filters by combining information from different activations in time. We employ a similar tactic to learn discriminative filters from pairwise image comparisons. We train SS-CNN for binary classification using the standard softmax or classification loss (L_c) with stochastic gradient descent. Since we perform classification between two categories (left image, right image), the softmax loss is specified as

$$L_c = \sum_{(i,j,y) \in P} \sum_k^K -\mathbb{1}[y = k] \log(g_k(x_i, x_j)) \quad (2)$$

where $K = 2$ and g is the softmax of final layer activations.

4.2 Ranking Streetscore-CNN

While the SS-CNN architecture learns to predict pairwise comparisons from two images, training with logistic loss does not account for the ordinal ranking over all the images in the dataset. Moreover, training for only binary classification may not be sufficient to train such complex networks to understand the fine-grained differences between image pairs [42]. Therefore, we explicitly incorporate the ranking function $f_r(x)$ (Eq. 1) in the end-to-end learning process, we modify this basic SS-CNN architecture by attaching a ranking sub-network, consisting

of fully-connected weight-tied layers (Fig. 2, in light blue). We call this network the Ranking SS-CNN (**RSS-CNN**). The RSS-CNN learns an additional set of weights—in comparison to SS-CNN—for minimizing a ranking loss,

$$L_r = \sum_{(i,j,y) \in P} (\max(0, y \cdot (f_r(x_j) - f_r(x_i)))^2. \quad (3)$$

The ranking loss (L_r) is designed to penalize the network to satisfy the constraints of our ranking problem—which is identical to the the loss function of the RankSVM [50,53] formulation. To train RSS-CNN, we minimize the loss function (L), which is a weighted combination of the classification (or softmax) loss (L_c), and the ranking loss (L_r), in the form $L = L_c(P) + \lambda L_r(P)$. We set the hyper-parameter λ using a grid-search to maximize the classification accuracy on the validation set.

5 Experiments and Results

After defining SS-CNN and RSS-CNN, we evaluate their performance in Sects. 5.1 and 5.2, using the 370, 134 pairwise comparisons collected for the question “Which place looks safer?”, since this question has the highest number of responses. Results for other attributes are described in Sect. 5.3.

Implementation Details. For all experiments, we split the set of triplets (P) for a given question randomly in the ratio 65–5–30 for training, validation and testing. We conducted experiments using the latest stable implementation of the Caffe library [54]. For both SS-CNN and RSS-CNN, we initialized the feature extractor layers using the pre-trained model weights of the following networks using their publicly available Caffe models² (one at a time): (i) the **AlexNet** image classification model [18], (ii) the **VGGNet** [20] 19-layer image classification model, and (iii) the **PlacesNet** [19] scene classification model. The weights for layers in fusion and ranking sub-networks were initialized from a zero-mean Gaussian distribution with standard deviation 0.01, following [18].

We trained the models on a single NVIDIA GeForce Titan X GPU. The momentum was set to 0.9. The initial learning rate was set to 0.001. When the validation error stopped improving with current learning rate, we reduced it by a factor of 10, repeating this process a maximum of four times (following [18]). The networks were trained to 100,000–150,000 iterations, stopping when the validation error stopped improving even after decreasing the learning rate. We will publicly release the models and the dataset upon acceptance.

5.1 Predicting Pairwise Comparisons

SS-CNN: We experiment with SS-CNN initialized using AlexNet, PlacesNet, and VGGNet, and evaluated their performance using three methods described below.

² <https://github.com/BVLC/caffe/wiki/Model-Zoo>.

Table 2. Pairwise comparison prediction accuracy

(a) SS-CNN				(b) RSS-CNN	
Network	Ranking Method			Model	Prediction Acc.
	Softmax	TrueSkill	RankSVM		
AlexNet	53.0%	55.7%	58.4%	AlexNet	64.1%
SS-CNN (AlexNet)	60.3%	62.6%	65.5%	PlacesNet	68.8%
PlacesNet	56.4%	58.8%	61.6%	VGGNet	73.5%
SS-CNN (PlacesNet)	62.2%	64.7%	68.1%		
VGGNet	60.9%	62.7%	63.5%		
SS-CNN (VGGNet)	65.3%	67.8%	72.4%		

- Softmax:** We calculate the binary prediction accuracy of the softmax output for prediction of pairwise comparisons.
- TrueSkill:** We generate 30 “synthetic” pairwise comparisons per image using the network, by feeding random image pairs, and calculate the TrueSkill score for each image with these comparisons. We compare TrueSkill scores of the two images in a pair, to predict the winning image for each pair in the test set, and measure the binary prediction accuracy. We use this method since TrueSkill is able to generate stable scores for images, which allows us to reduce the noise in independent binary predictions on image pairs.
- RankSVM:** We feed a combined feature representation of the image pair obtained from the final convolution layer of SS-CNN to a RankSVM [50] (using the LIBLINEAR [55] implementation), and learn a ranking function. We then use the ranking scores for images in the test set to decide the winner from test image pairs, and calculate the binary prediction accuracy.

We evaluate the accuracy for all three networks with (i) original (pre-trained) weights, and (ii) weights fine-tuned with the Place Pulse 2.0 dataset. Table 2-(a) shows that, in all cases, the binary prediction accuracy increases significantly—6.5% on average—across all experiments. The gain in performance can be attributed to both, end-to-end learning of the pairwise classification task and the size and diversity of the Place Pulse 2.0 dataset. SS-CNN (VGGNet), the deepest architecture, obtains the best performance over all three methods. We also observe that RankSVM consistently outperforms TrueSkill, which in turn, outperforms softmax. This makes sense, since TrueSkill is not designed to maximize prediction accuracy for pairwise comparisons, but rather to generate stable ranked scores from pairwise comparisons. In contrast, the RankSVM loss function explicitly tries to minimize misclassification in pairwise comparisons.

RSS-CNN: We test the performance of the RSS-CNN architecture with AlexNet, PlacesNet, and VGGNet. Since we explicitly learn a ranking function $f_r(x)$ in the case of RSS-CNN, we compare the ranking function outputs for both images in a test pair to decide which image wins, and calculate the binary prediction accuracy. Table 2-(b) summarizes the results for the three models.

Table 3. Comparing place pulse 1.0 and place pulse 2.0 datasets

Ranking SS-CNN	AlexNet	PlacesNet	VGGNet
Place Pulse 1.0 (PP 1.0)	59.8 %	60.9 %	64.1 %
Place Pulse 2.0 (same #comparisons as PP 1.0)	61.9 %	66.2 %	64.2 %
Place Pulse 2.0 (all comparisons)	64.1 %	68.8 %	73.5 %

The Ranking SS-CNN (VGGNet) obtains the highest accuracy for pairwise comparison prediction (73.5%). Since the RSS-CNN performs end-to-end learning based on both the classification and ranking loss, it significantly outperforms the SS-CNN trained with only classification loss (Table 2-(a), column 1). The RSS-CNN also does better than the combination of SS-CNN and RankSVM (Table 2-(a), column 3) in most cases. We also find that RSS-CNN learns better with more data, and continues to do so, whereas the SS-CNN architecture plateaus after encountering approximately 60% of the training data. See supplementary material for additional analysis on data size and performance.

5.2 Comparing Place Pulse 1.0 and Place Pulse 2.0 Datasets

The Place Pulse 2.0 (PP 2.0) dataset has significantly higher visual diversity (56 cities from 28 countries) as compared to the Place Pulse 1.0 (PP 1.0) dataset (4 cities from 2 countries). It also contains significantly more training data. For the visual attribute of *Safety*, the PP 2.0 dataset contains 370,134 comparisons for 110,988 images, while the PP 1.0 dataset contains 73,806 comparisons for 4,109 images. We are interested in studying the gain in performance obtained by this increased visual diversity and size. So we compare the binary prediction accuracy on PP 2.0 data, of an RSS-CNN trained with the three network architectures using (i) all 73,806 comparisons from PP 1.0, (ii) 73,806 comparisons randomly chosen from PP 2.0 (the same amount of data as PP 1.0, but an increase in visual diversity), and (iii) 240,587 comparisons from PP 2.0 (the entire training set) (an increase in both the amount and the visual diversity of data). Comparing experiments (i) and (ii) (Table 3), we find that increasing visual diversity improves the accuracy for all three networks, for the same amount of data. The gain in performance is least for VGGNet, which is the deepest network, and hence needs larger amount of data to train. Finally, training with the entire PP 2.0 dataset (experiment (iii)) improves accuracy by an average of 7.2% as compared to training with the PP 1.0.

We also conduct the reverse experiment to measure the performance of the PP 2.0 dataset on PP 1.0. We calculate the five-fold cross-validation accuracy (following [13]) for pairwise comparison prediction for the *Safety* attribute using a RankSVM trained with features of image pairs from the PP 1.0 dataset. We experiment with two different features, extracted, respectively, from (i) the SS-CNN (VGGNet) trained with PP 2.0 data and (ii) the SS-CNN (VGGNet) trained with PP 2.0 data and fine-tuned further with PP 1.0 data. Experiments

Table 4. Prediction performance across attributes

Train	Test					
	Safe	Lively	Beautiful	Wealthy	Boring	Depressing
Safe	73.5 %	67.7 %	66.3 %	60.3 %	47.2 %	42.3 %
Lively	63.8 %	70.3 %	65.8 %	61.3 %	58.9 %	53.7 %
Beautiful	61.2 %	67.1 %	70.2 %	53.5 %	50.2 %	51.4 %
Wealthy	60.7 %	54.6 %	52.7 %	65.7 %	52.8 %	55.9 %
Boring	48.6 %	55.6 %	52.3 %	53.1 %	66.1 %	59.8 %
Depressing	54.5 %	54.2 %	43.2 %	49.7 %	57.2 %	62.8 %

(i) and (ii) yield an accuracy of 81.6% and 81.1% respectively. The previous best result reported for the pairwise comparison prediction task [13] on the PP 1.0 dataset is 70.2%, albeit from a model trained with PP 1.0 data alone. Note that our models are too deep to be trained with only PP 1.0 data.

Comparison with Generic Image Features: Prior work [11–13] has found that generic image features do well on the Place Pulse 1.0 dataset, for predicting both ranked scores and pairwise comparisons. Based on this literature, we extract three best performing features—GIST [56], Texton Histograms [57], and CIELab Color Histograms [58]—from images in the PP 2.0 dataset. We find that the pairwise prediction accuracy of a RankSVM trained with feature vector consisting of these features is 56.7% on the PP 2.0 dataset, significantly lower than all variations of SS-CNN. Our best performing model RSS-CNN (VGGNet) has an accuracy of 73.5%.

5.3 Predicting Different Perceptual Attributes

Our dataset contains a total of six perceptual attributes—*Safe*, *Lively*, *Beautiful*, *Wealthy*, *Boring*, and *Depressing*. We now evaluate the prediction performance of RSS-CNN on these six attributes. Specifically, we train the RSS-CNN (VGGNet) network for each attribute, and measure its performance using binary prediction accuracy. Table 4 shows that the *in-attribute* prediction performance is roughly proportional to the number of comparisons available for training, with the best prediction performance for *Safe*, and the worst performance for *Depressing*. We also evaluate the performance of the network trained to predict one perceptual attribute in predicting the pairwise comparisons for the other three attributes (*cross-attribute* prediction). The *Safe* network shows strong performance in prediction of *Lively*, *Beautiful*, and *Wealthy* attributes, which is indicative of the high correlation between different perceptual attributes.

A model trained to predict pairwise comparisons can be used to generate “synthetic” comparisons by taking random image pairs as input. A large number of comparisons can be then fed to ranking algorithms (like TrueSkill) to



Fig. 3. Example results from the Place Pulse 2.0 dataset, containing images ranked based on pairwise comparisons generated by the RSS-CNN.

obtain stable ranked scores. We use this trick to generate TrueSkill scores for four attributes using pairwise comparisons predicted by a trained RSS-CNN (VGGNet) (30 per image). Figure 3 shows examples from the dataset, and Fig. 4 shows failure cases. We find that, for instance, highway images with forest cover are predicted to be highly safe, and overcast images as highly boring. Quantitatively, the correlation coefficient (R^2) of *Safe* with *Lively*, *Beautiful*, and *Wealthy* is 0.80, 0.83, and 0.65 respectively. This indicates that there is relatively large orthogonality ($(1 - R^2)$) between attributes. See supplement for details.

5.4 Predicting Urban Appearance Across Countries

Our hope is that the Place Pulse 2.0 dataset will enable algorithms to conduct automated audits of urban appearance for cities all across the world. The Streetscore [11] algorithm was able to successfully generalize to the Northeast and Midwest of the U.S., based on training data from just two cities, New York and Boston. This indicates that models trained with the PP 2.0 dataset containing images from 28 countries should be able to generalize to large regions in these countries, and beyond. For a qualitative experiment to test generalization, we download 22,282 Street View images from six cities from six continents—Vancouver, Buenos Aires, St. Petersburg, Durban, Seoul, and Brisbane—that were not a part of the PP 2.0 dataset. We map the perceived safety for these



Fig. 4. Example failure cases from the prediction results, containing images and their TrueSkill scores for attributes computed from pairwise comparisons generated by the RSS-CNN.

cities using TrueSkill scores for images computed from 30 “synthetic” pairwise comparisons generated with RSS-CNN (VGGNet). While the prediction performance of the network on these images cannot be quantified due to a lack of human-labeled ground truth, visual inspection shows that the scores assigned to streetscapes conform with visual intuition (see supplement for map visualizations and example images).

6 Discussion and Concluding Remarks

In this paper, we introduced a new crowdsourced dataset of global urban appearance containing pairwise image comparisons and proposed a neural network architecture for predicting the human-labeled comparisons. Since we focussed on predicting pairwise win/loss decisions to aid image ranking, we ignored the image pairs where the users perceive the images to be *equal* for the given perceptual attribute. However, 13.2% pairwise comparisons in our dataset are *equal*, and incorporating the prediction of equality in comparisons should be a part of future work. Future work can also explore the determinants of perceptual attributes of urban appearance (e.g., what makes an image appear safe? or lively?) Such studies would allow better visual designs that optimize attributes of urban appearance. From a computer vision perspective, understanding the geographical range over which models trained on street-level imagery from different regions of the world are able to generalize would be an interesting future direction, since the architectural similarities between cities are determined by a complex interaction of history, culture, and economics.

Our technique can be generalized for computer vision tasks of studying the style, perception, or visual attributes of images, objects, or scene categories. Our trained networks can be used to generate a global dataset of urban appearance, which will enable the study a variety of research questions: How does urban appearance affect the behavior and health of residents, and how do these effects vary across countries? How are different architectural styles perceived? How similar/different are different cities across the world in terms of perception? Can visual appearance be used as a proxy for inequality within cities? A global dataset of urban appearance will thus aid computational studies in architecture, art history, sociology, and economics. These datasets can also help policymakers and city governments make data-driven decisions on allocation of resources to different cities or neighborhoods for improving urban appearance.

References

1. Wilson, J.Q., Kelling, G.L.: Broken windows. *Atlantic Monthly* **249**(3), 29–38 (1982)
2. Keizer, K., Lindenberg, S., Steg, L.: The spreading of disorder. *Science* **322**(5908), 1681–1685 (2008)
3. Milam, A., Furr-Holden, C., Leaf, P.: Perceived school and neighborhood safety, neighborhood violence and academic achievement in urban school children. *Urban Rev.* **42**(5), 458–467 (2010)
4. Cohen, D.A., Mason, K., Bedimo, A., Scribner, R., Basolo, V., Farley, T.A.: Neighborhood physical conditions and health. *Am. J. Public Health* **93**(3), 467–471 (2003)
5. Piro, F.N., Næss, Ø., Claussen, B.: Physical activity among elderly people in a city population: the influence of neighbourhood level violence and self perceived safety. *J. Epidemiol. Commun. Health* **60**(7), 626–632 (2006)
6. Sampson, R.J.: *Great American City: Chicago and the enduring neighborhood effect*. University of Chicago Press, Chicago (2012)
7. Miller, D.K.: Using google street view to audit the built environment: inter-rater reliability results. *Ann. Behav. Med.* **45**(1), 108–112 (2013)
8. Hwang, J., Sampson, R.J.: Divergent pathways of gentrification racial inequality and the social order of renewal in chicago neighborhoods. *Am. Sociol. Rev.* **79**(4), 726–751 (2014)
9. Salesses, P., Schechtner, K., Hidalgo, C.A.: The collaborative image of the city: mapping the inequality of urban perception. *PloS One* **8**(7), e68–400 (2013)
10. Quercia, D., O’Hare, N.K., Cramer, H.: Aesthetic capital: what makes London look beautiful, quiet, and happy? In: *ACM Conference on Computer Supported Cooperative Work & Social Computing*, pp. 945–955 (2014)
11. Naik, N., Philipoom, J., Raskar, R., Hidalgo, C.: Streetscore-Predicting the perceived safety of one million streetscapes. In: *IEEE CVPR Workshops*, pp. 793–799 (2014)
12. Ordonez, V., Berg, T.L.: Learning high-level judgments of urban perception. In: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (eds.) *ECCV 2014*. LNCS, vol. 8694, pp. 494–510. Springer, Heidelberg (2014). doi:[10.1007/978-3-319-10599-4_32](https://doi.org/10.1007/978-3-319-10599-4_32)
13. Porzi, L., Rota Bulò, S., Lepri, B., Ricci, E.: Predicting and understanding urban perception with convolutional neural networks. In: *ACM Conference on Multimedia*, pp. 139–148 (2015)
14. Naik, N., Raskar, R., Hidalgo, C.A.: Cities are physical too: using computer vision to measure the quality and impact of urban appearance. *Am. Econ. Rev.* **106**(5), 128–132 (2016)
15. Been, V., Ellen, I.G., Gedal, M., Glaeser, E., McCabe, B.J.: Preserving history or restricting development? the heterogeneous effects of historic districts on local housing markets in new york city. *J. Urban Econ.* **92**, 16–30 (2015)
16. Naik, N., Kominers, S.D., Raskar, R., Glaeser, E.L., Hidalgo, C.A.: Do people shape cities, or do cities shape people? the co-evolution of physical, social, and economic change in five major U.S. cities. Working Paper 21620, National Bureau of Economic Research (2015)
17. Harvey, C., Aultman-Hall, L., Hurley, S.E., Troy, A.: Effects of skeletal streetscape design on perceived safety. *Landscape Urban Plann.* **142**, 18–28 (2015)
18. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. In: *Advances in Neural Information Processing Systems*, pp. 1097–1105(2012)

19. Zhou, B., Lapedriza, A., Xiao, J., Torralba, A., Oliva, A.: Learning deep features for scene recognition using places database. In: *Advances in Neural Information Processing Systems*, pp. 487–495 (2014)
20. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition (2014). arXiv preprint [arXiv:1409.1556](https://arxiv.org/abs/1409.1556)
21. Joshi, D., Datta, R., Fedorovskaya, E., Luong, Q.T., Wang, J.Z., Li, J., Luo, J.: Aesthetics and emotions in images. *IEEE Sig. Process. Mag.* **28**(5), 94–115 (2011)
22. Isola, P., Xiao, J., Torralba, A., Oliva, A.: What makes an image memorable? In: *IEEE CVPR*, pp. 145–152 (2011)
23. Dhar, S., Ordonez, V., Berg, T.L.: High level describable attributes for predicting aesthetics and interestingness. In: *IEEE CVPR*, pp. 1657–1664 (2011)
24. Deza, A., Parikh, D.: Understanding image virality. In: *IEEE CVPR*, pp. 1818–1826 (2015)
25. Donahue, J., Jia, Y., Vinyals, O., Hoffman, J., Zhang, N., Tzeng, E., Darrell, T.: Decaf: a deep convolutional activation feature for generic visual recognition (2013). arXiv preprint [arXiv:1310.1531](https://arxiv.org/abs/1310.1531)
26. Doersch, C., Singh, S., Gupta, A., Sivic, J., Efros, A.A.: What makes paris look like paris? *ACM Trans. Graph.* **31**(4), 101 (2012)
27. Lee, S., Maisonneuve, N., Crandall, D., Efros, A., Sivic, J.: Linking past to present: discovering style in two centuries of architecture. In: *IEEE International Conference on Computational Photography* (2015)
28. Arietta, S.M., Efros, A.A., Ramamoorthi, R., Agrawala, M.: City forensics: using visual elements to predict non-visual city attributes. *IEEE Trans. Visual. Comput. Graph.* **20**(12), 2624–2633 (2014)
29. Glaeser, E.L., Kominers, S.D., Luca, M., Naik, N.: Big data and big cities: the promises and limitations of improved measures of urban life. Working Paper 21778, National Bureau of Economic Research (2015)
30. Zhou, B., Liu, L., Oliva, A., Torralba, A.: Recognizing city identity via attribute analysis of geo-tagged images. In: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (eds.) *ECCV 2014. LNCS*, vol. 8691, pp. 519–534. Springer, Heidelberg (2014). doi:[10.1007/978-3-319-10578-9_34](https://doi.org/10.1007/978-3-319-10578-9_34)
31. Khosla, A., An, B., Lim, J.J., Torralba, A.: Looking beyond the visible scene. In: *IEEE CVPR*, pp. 3710–3717 (2014)
32. Kuipers, M.A., van Poppel, M.N., van den Brink, W., Wingen, M., Kunst, A.E.: The association between neighborhood disorder, social cohesion and hazardous alcohol use: a national multilevel study. *Drug Alcohol Depend.* **126**(1), 27–34 (2012)
33. Dulin-Keita, A., Thind, H.K., Affuso, O., Baskin, M.L.: The associations of perceived neighborhood disorder and physical activity with obesity among african american adolescents. *BMC Pub. Health* **13**(1), 440 (2013)
34. Kelling, G.L., Coles, C.M.: *Fixing Broken Windows: Restoring Order and Reducing Crime in Our Communities*. Simon and Schuster, New York (1997)
35. Sampson, R.J., Raudenbush, S.W.: *Disorder in urban neighborhoods: Does it lead to crime*. National Institute of Justice (2001)
36. Harcourt, B.E.: Reflecting on the subject: a critique of the social influence conception of deterrence, the broken windows theory, and order-maintenance policing New York style. *Mich. Law Rev.* **97**(2), 291–389 (1998)
37. Parikh, D., Grauman, K.: Relative attributes. In: *IEEE ICCV*, pp. 503–510 (2011)
38. Parkash, A., Parikh, D.: Attributes for classifier feedback. In: Fitzgibbon, A., Lazebnik, S., Perona, P., Sato, Y., Schmid, C. (eds.) *ECCV 2012. LNCS*, vol. 7574, pp. 354–368. Springer, Heidelberg (2012). doi:[10.1007/978-3-642-33712-3_26](https://doi.org/10.1007/978-3-642-33712-3_26)

39. Kovashka, A., Parikh, D., Grauman, K.: Whittlesearch: Image search with relative attribute feedback. In: IEEE CVPR, pp. 2973–2980 (2012)
40. Kiapour, M.H., Yamaguchi, K., Berg, A.C., Berg, T.L.: Hipster wars: discovering elements of fashion styles. In: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (eds.) ECCV 2014. LNCS, vol. 8689, pp. 472–488. Springer, Heidelberg (2014). doi:[10.1007/978-3-319-10590-1_31](https://doi.org/10.1007/978-3-319-10590-1_31)
41. Zhu, J.Y., Agarwala, A., Efros, A.A., Shechtman, E., Wang, J.: Mirror mirror: crowdsourcing better portraits. *ACM Trans. Graph.* **33**(6), 234 (2014)
42. Wang, J., Song, Y., Leung, T., Rosenberg, C., Wang, J., Philbin, J., Chen, B., Wu, Y.: Learning fine-grained image similarity with deep ranking. In: IEEE CVPR, pp. 1386–1393 (2014)
43. Zagoruyko, S., Komodakis, N.: Learning to compare image patches via convolutional neural networks. In: IEEE CVPR, pp. 4353–4361 (2015)
44. Persson, P.O., Strang, G.: A simple mesh generator in MATLAB. *SIAM Rev.* **46**(2), 329–345 (2004)
45. Stewart, N., Brown, G.D., Chater, N.: Absolute identification by relative judgment. *Psychol. Rev.* **112**(4), 881 (2005)
46. Bijmolt, T.H., Wedel, M.: The effects of alternative methods of collecting similarity data for multidimensional scaling. *Int. J. Res. Mark.* **12**(4), 363–371 (1995)
47. Jou, B., Bhattacharya, S., Chang, S.F.: Predicting viewer perceived emotions in animated GIFs. In: ACM International Conference on Multimedia, pp. 213–216 (2014)
48. Sartori, A., Yanulevskaia, V., Salah, A.A., Uijlings, J., Bruni, E., Sebe, N.: Affective analysis of professional and amateur abstract paintings using statistical analysis and art theory. *ACM Trans. Interact. Intell. Syst.* **5**(2), 8 (2015)
49. Herbrich, R., Minka, T., Graepel, T.: TrueSkill: a bayesian skill rating system. In: Advances in Neural Information Processing Systems, pp. 569–576 (2006)
50. Joachims, T.: Optimizing search engines using clickthrough data. In: ACM International Conference on Knowledge Discovery and Data Mining, pp. 133–142 (2002)
51. Chopra, S., Hadsell, R., LeCun, Y.: Learning a similarity metric discriminatively, with application to face verification. In: IEEE CVPR, vol. 1, pp. 539–546 (2005)
52. Karpathy, A., Toderici, G., Shetty, S., Leung, T., Sukthankar, R., Fei-Fei, L.: Large-scale video classification with convolutional neural networks. In: IEEE CVPR, pp. 1725–1732 (2014)
53. Chapelle, O., Keerthi, S.S.: Efficient algorithms for ranking with SVMs. *Inf. Retrieval* **13**(3), 201–215 (2010)
54. Jia, Y., Shelhamer, E., Donahue, J., Karayev, S., Long, J., Girshick, R., Guadarrama, S., Darrell, T.: Caffe: convolutional architecture for fast feature embedding. In: ACM International Conference on Multimedia, pp. 675–678 (2014)
55. Fan, R.E., Chang, K.W., Hsieh, C.J., Wang, X.R., Lin, C.J.: Liblinear: a library for large linear classification. *J. Mach. Learn. Res.* **9**, 1871–1874 (2008)
56. Oliva, A., Torralba, A.: Modeling the shape of the scene: a holistic representation of the spatial envelope. *Int. J. Comput. Vis.* **42**(3), 145–175 (2001)
57. Malik, J., Belongie, S., Leung, T., Shi, J.: Contour and texture analysis for image segmentation. *Int. J. Comput. Vis.* **43**(1), 7–27 (2001)
58. Xiao, J., Hays, J., Ehinger, K.A., Oliva, A., Torralba, A.: Sun database: large-scale scene recognition from abbey to zoo. In: IEEE CVPR, pp. 3485–3492 (2010)