Supplementary Material for: "Discovering Underground Maps from Fashion"

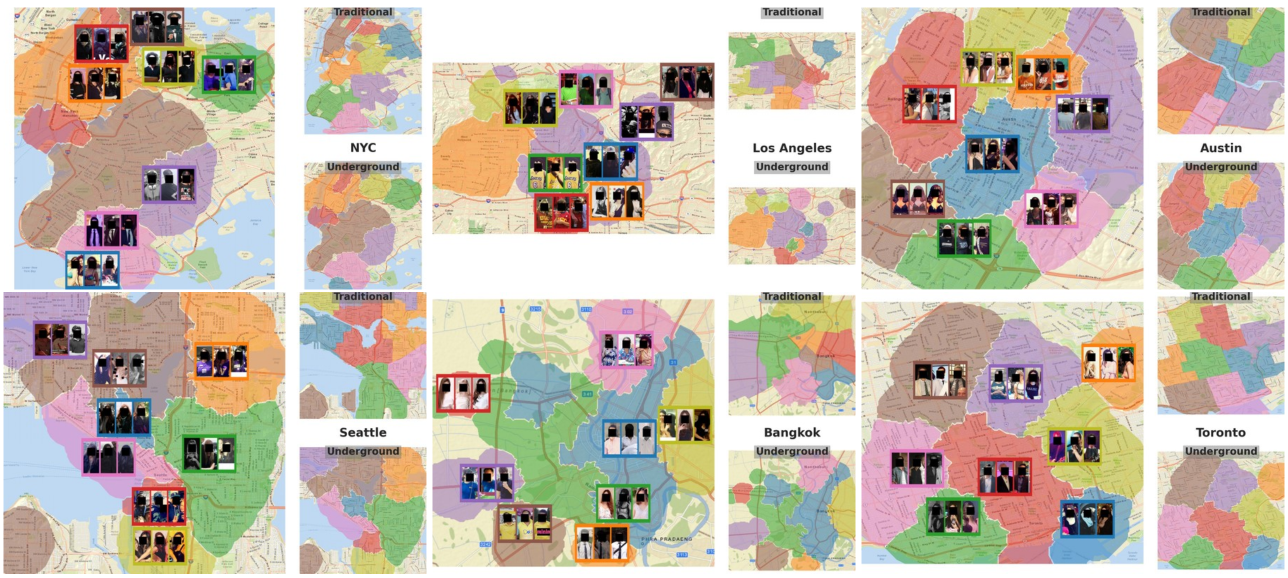

Figure 1: Underground maps in contrast with the traditional maps for 6 cities. Our method can discover neighborhoods based on activities like student neighborhoods in Los Angeles (red), Austin (orange) and Seattle (orange) by looking at the fashion sense. It also discovers tourist neighborhoods of NYC (red), Seattle (Pink).

In this supplementary document we present some of the results that we could not present in the main paper. Please also refer to the **video** along with this supplementary to explore a wide range of results.

Sec. 1 shows underground maps for many cities discovered by our method. In Sec. 2, we present additional implementation details about the benchmark creation and getting unbiased locations. Sec. 3 refers to the list of cities used for different types of analysis in our results. In Sec. 4 we break down the performance of our method further for each city and for each class of HM benchmark. In Sec. 5 we conduct various ablation studies by measuring the effect of different hyperparameters on our method. Sec. 6 shows how we collected data for experiments with local human judges. In Sec. 7 we present additional qualitative results of unique, similar, and analogical neighborhoods found by our method. Finally, in Sec. 8 we discuss other variants we tried for grouping localized features.

1. Underground Maps

Figure 1 shows the underground maps for 6 cities and contrasts them with their traditional maps. Our method discovers similarity across regions that are geographically far apart. For example, for Bangkok it finds 2 regions (colored blue) that are geographically far away, but people wear similar clothing. Also these maps have very different boundaries than a traditional map.

Our method can shed light on a city by discovering many non-obvious insights in a city. Figure 2 show underground maps of some cities where our discovered maps can reveal many things about a city. Here we discuss them in more details.

Delhi: It discovers traditional vs. Westernized neighborhoods. People in the southern neighborhood are seen in Western

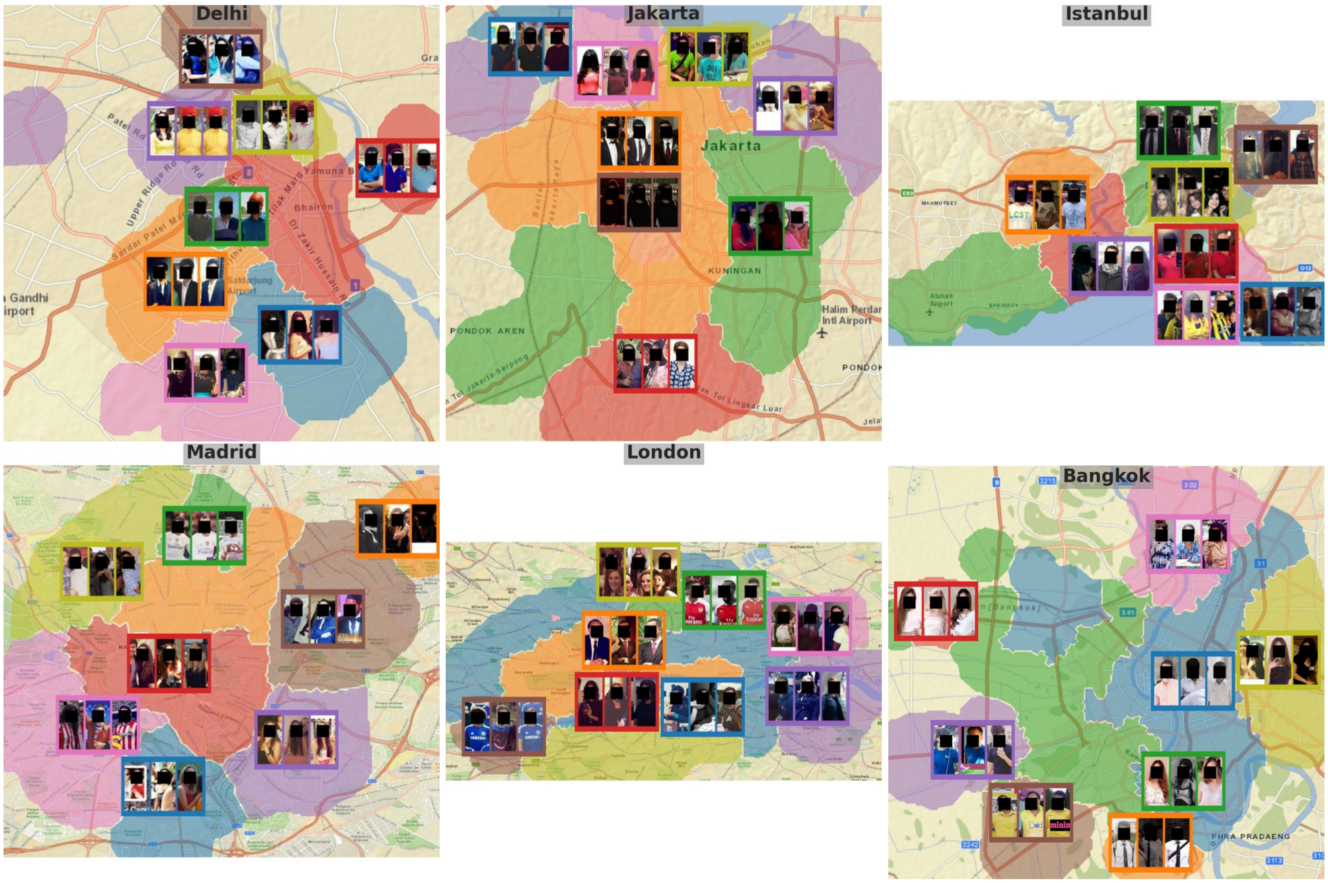

Figure 2: Our maps can reveal unexpected insights about a city. For example, in Delhi and Bangkok our method indicates how the city expanded over time. In Istanbul and London, our method finds out business neighborhoods that are geographically far apart but similar. In London and Madrid, we discover neighborhoods of popular football clubs. In Jakarta we discover tourist neigborhoods are found closer to the beaches and museums.

clothing more than the northern ones, explainable by the fact that the city expanded from north to south as it grew.[1]

Jakarta: We discover areas popular among tourists (pink and yellow) in the north closer to beaches and museums. In South Jakarta more people can be seen wearing traditional Indonesian clothing (red) as well as in southeast and southwest neighborhoods (shown in green).

Istanbul: We discover two far-apart but similar neighborhoods (green) that are popular business neighborhoods. Both the neighborhood west (Bakirkoy)[2] and north (Sisly)[3]. We also discover tourist neighborhoods (red) near the popular tourist areas of Fatih and Beyoglu.[4]

Madrid: We discover two neighborhoods with stadiums and fans of two different football clubs: Real Madrid (green) and Athletico Madrid (pink). These neighborhoods are different from other neighborhoods in the city, like the neighborhood with many nightclubs and pubs (red).

London: We find commercial neighborhoods in Central London (Square Mile and London Bridge) and West London (West End) marked in orange, where people can be seen wearing suits. Similar to Madrid, London also has popular football clubs; hence we find neighborhoods of those clubs Chelsea (brown) in the west and Arsenal (green) in the north.

Bangkok: Expansion of a city can also explain the pattern in Bangkok from traditional neighborhood(pink) vs westernized neighborhood (yellow).

[1] https://earthobservatory.nasa.gov/images/92813/urban-growth-of-new-delhi
[2] https://wikipedia.org/wiki/Bak%C4%B1rk%C3%B6y
[3] https://wikipedia.org/wiki/%c5%9ei%c5%9fli
[4] https://wikipedia.org/wiki/Beyo%C4%9Flu

Figures 7 and 8 show the underground maps for other cities. Again, our method discovers similarity across regions that are geographically far apart, and can find some very specific neighborhoods.

2. Additional Implementation Details

2.1. Creation of BD Benchmark

In this section we present more information on how the BD Benchmark is created.

To create regions over maps using business density, we follow our method from Sec. 3 (main paper), with a few differences. First, since we need to segment on the basis of business density, the featurization of a region on the map is the normalized histogram of different businesses. Second, to keep the similarity of the two benchmarks, rather than sampling a circular radius around a point, we sample a square region of length and width 0.01°. Finally, we do not know the number of neighborhoods (unlike the HoodMaps Benchmark, where the number of labels is given). So we use affinity propagation for clustering instead of K-means to find the ideal number of neighborhoods per city. Affinity propagation produces a min/max/median of 3/6/6 types of regions per city.

2.2. Getting Unbiased Locations

The samples within the radius r are not going to be distributed uniformly around $\boldsymbol{x}$. These samples will be biased towards certain directions on the map. For example, if we sample for a location at the junction of land and sea, almost no images will be sampled from sea, hence, the location of the image samples would be biased towards land. The mean of the sampled image location would be the coordinate for which these samples are unbiased. Therefore, we use the new unbiased location in the pipeline instead of the sampling location. The histogram description $\boldsymbol{h_x}$ is an unbiased descriptor of this location $\boldsymbol{x'}$ hence these location is used in the later stages of the pipeline. $T(\boldsymbol{x})$ are the set of images that describes the sampled location $\boldsymbol{x}$:

$$T(\boldsymbol{x}) = \{I_i : ||\boldsymbol{l}_i - \boldsymbol{x}||_2 < r\} \tag{1}$$

The unbiased location is:

$$\boldsymbol{x'} = \frac{\sum_{I_i \in T(\boldsymbol{x})} \boldsymbol{l}_i}{|T(x)|} \tag{2}$$

3. List of Cities

Austin	Bangkok	Beijing	Berlin	Bogotà	Budapest	Buenos Aires	Cairo
Chicago	Delhi	Dhaka	Guangzhou	Istanbul	Jakarta	Johannesburg	Karachi
Kyiv	Kolkata	Lagos	London	Los Angeles	Madrid	Manila	Mexico City
Milan	Moscow	Mumbai	Nairobi	NYC	Osaka	Paris	Rio
Rome	São Paulo	Seattle	Seoul	Shanghai	Singapore	Sofia	Sydney
Tianjin	Tokyo	Toronto	Vancouver				

Table 1: Cities in our analysis. Cities colored blue and red are evaluated quantitatively using both the HM and BD benchmarks. Cities colored red have administrative boundary data available and so are evaluated for the right hand side of Table 1 in the main paper.

Table 1 presents the list of 44 cities from GeoStyle. 37 of these 44 cities have enough data available on HoodMaps for us to use. We do the quantitative analysis on all 37 of these cities using HM and BD benchmarks. The 8 cities marked red are where we can get publicly available administrative boundary information.

4. Quantitative Performance Breakdown

4.1. Per-Class Performance for HM Benchmark

We also compare Per-Class performance against the Admin baseline for the 8 cities. Table 2 compares MMIoU for these cities. Our method finds it difficult to discover corporate neighborhoods, while it is better at understanding hipster neighborhoods for these cities. We believe this is because fewer people are posting images of themselves from a corporate environment.

	Wea.	Hip.	Tou.	Stu.	Nor.	Cor.
Admin	0.315	0.327	0.269	0.261	0.260	0.168
Ours	0.306	0.408	0.294	0.281	0.281	0.129

Table 2: Per class accuracy on HoodMaps (MMIoU) for 8 cities where Admin data is available. On average our method performs better than admin baseline, But our method finds it difficult to find corporate neighborhoods as fewer people post pictures of themselves from a corporate environment.

4.2. Per-city Performance

Figure 3 shows the per-city purity measure comparing our method against the best-performing baseline (PID) for the HM benchmark. PID performs better than our method for the cities to the left and our method outperforms PID for the cities to the right. Figure 4 similarly compares our method against admin baseline (PID) for the 8 cities over the HM benchmark. Our method performs better than admin for all the cities.

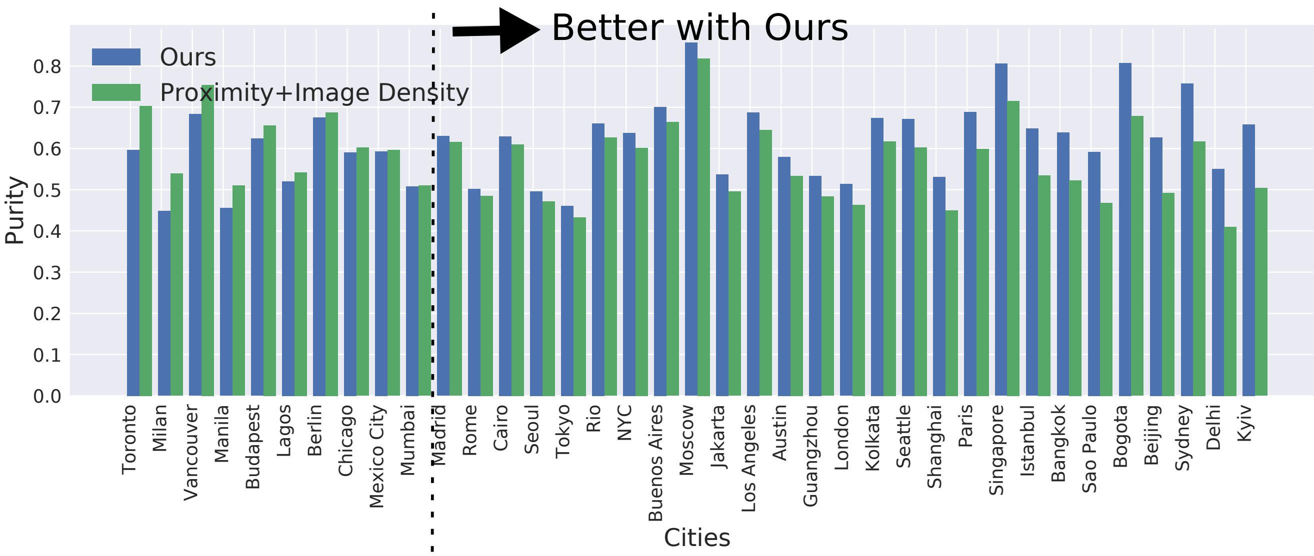

Figure 3: Performance (purity over HM benchmark) of our method against the PID baseline for individual cities. The cities are sorted from left to right by the cities where PID performs best to the cities where our method performs the best. For 27 out of 37 cities, we perform better than the PID baseline.

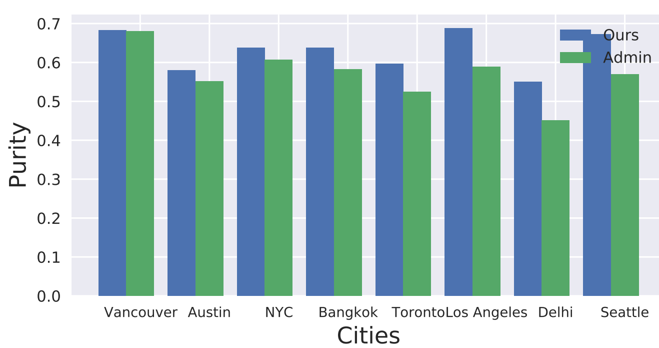

Figure 4: Performance (purity over HM benchmark) of our method against the Admin baseline for all 8 cities where Admin is available. The cities are sorted from left to right by the increasing gain of our method. Our method performs better than Admin for all the cities.

5. Ablation Studies

5.1. Effect of Radius r

We look at how changing the sampling radius affects the performance of our method on the two benchmarks. We increase and decrease the radius by a factor of $\sqrt{2}$. Table 3 shows the performance of our method with changing radius r. The best performance is at $r = 0.020$. However, as can be seen by the performance on BD benchmark, changing the radius does not have a huge impact on the performance.

	HM Benchmark			BD Benchmark		
r	NMI	Purity	MMIoU	NMI	Purity	MMIoU
0.010	0.185	0.607	0.212	0.293	0.580	0.306
0.014	0.216	0.599	0.238	0.347	0.599	**0.336**
0.020	**0.260**	**0.635**	**0.272**	**0.369**	**0.597**	**0.339**
0.028	0.246	0.573	0.257	**0.369**	**0.598**	**0.332**
0.04	0.240	0.584	0.254	0.350	0.576	0.325

Table 3: Effect of changing sampling radius r on the performance of neighborhood discovery.

5.2. Number of Discovered Clusters

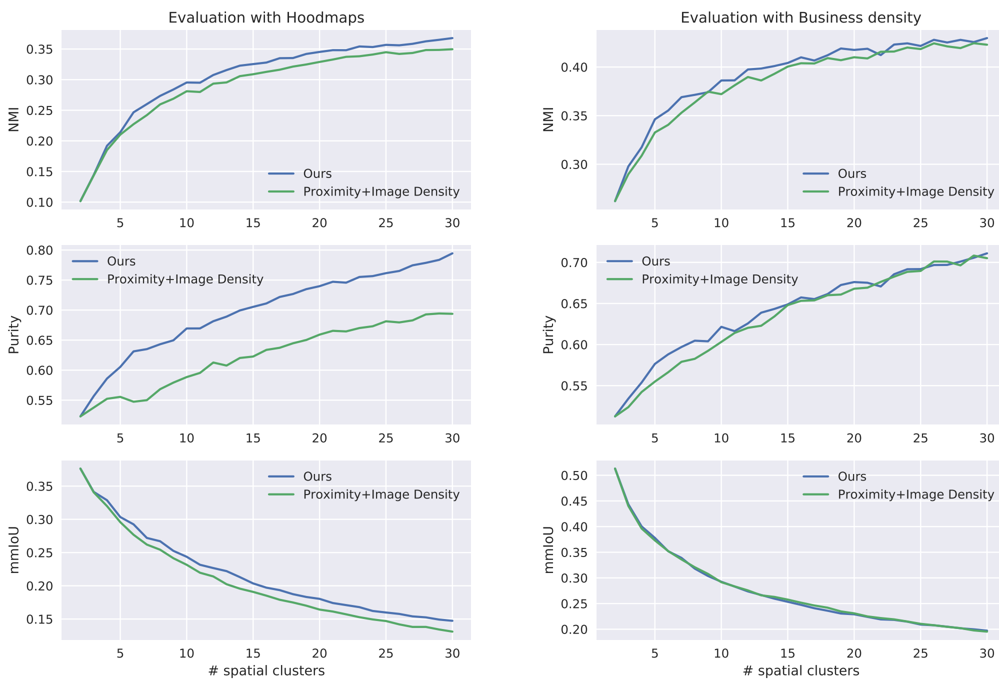

Figure 5: Performance (NMI, Purity and MMIoI) of our method against the PID baseline when varying the number of discovered clusters for HM (left) and BD (right). Our method outperforms the best baseline on all metrics and can be used to discover any number of neighborhoods in a city.

In this section we see how changing the number of spatial clusters affects performance on the evaluation metric. Figure 5 shows the performance when sweeping over the number of spatial clusters for all three unsupervised evaluation metrics over the two benchmarks. Our method consistently outperforms the baseline method for all metrics (performs on par for MMIoU for the BD benchmark). Hence our conclusions are stable with respect to this parameter setting.

6. Experiments with Human Judges

We restrict the Amazon Mechanical Turk workers to perform tasks for cities of their country (approximate way of enforcing a "local" person). Additionally, we add sentinels in our tasks, by duplicating 4 questions out of 20. If the 2 clicked points for the same question are far apart from each other, we do not consider the answers of that worker.

Figure 6 shows the interface as seen by the Mechanical Turk workers. Workers are supposed to look at the set of images as shown on the right and click on the point where they think the images come from (red marker on map). Workers are allowed to zoom in/out on the map, and redo a previously selected image. We additionally also ask for a confidence rating for a click. Note that it is not visible to the workers which method (ours or the baseline) has produced the images. They are simply asked to localize the styles they see.

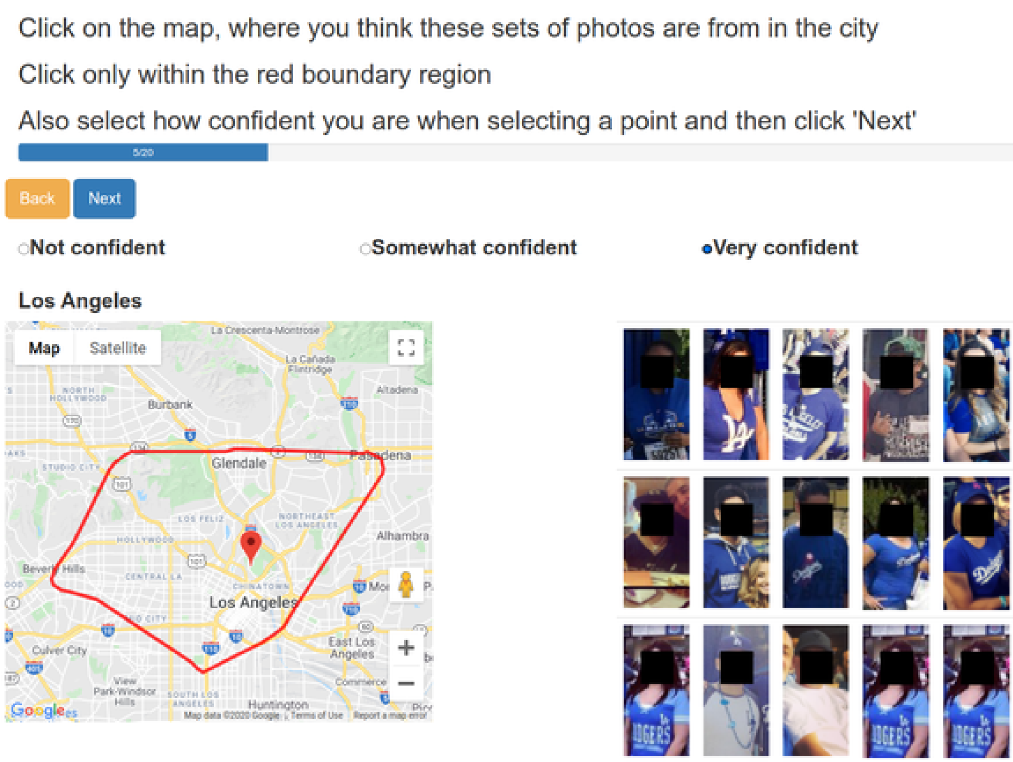

Figure 6: Interface seen by the Amazon Mechanical Turk workers. The instructions can be seen on the top.

7. More Qualitative Results

7.1. Examples of Unique Neighborhoods

Table 4 shows the list of the top 40 most unique neighborhoods found by our methods. It shows unique regions such as neighborhoods with sports stadiums (Los Angeles, Milan, Chicago, Seattle), tourist areas (Bogota, Beijing (second neighborhood)), beaches (NYC, Sydney (fourth neighborhood)) etc.

7.2. Examples of Similar Neighborhoods

Table 5 shows the list of the top 20 most similar neighborhoods found by our methods. It shows similar neighborhoods with tourists (Kyiv-Moscow, Chicago-NYC (first pair)), nightlife (Chicago-NYC (second pair), NYC-Toronto). Note that this measure of similarity only find pairs that are geographically close (due to similar weather and culture).

7.3. Examples of Analogical Neighborhoods

Table 6 shows the list of the top 20 tuples of contextually similar neighborhoods, where non-contextual similarity fails to give good similar neighborhoods. The first and second column are our contextually similar neighborhoods. The third column shows the neighborhood produced by simpler similarity search. Note that most of the pairs are from cities that are geographically far (where non-contextual similarity faces more challenges). The first example finds similar tourist regions in culturally different cities Istanbul (neighborhood around Hagia Sophia) and Rome (neighborhoods around Colosseum and Vatican City).

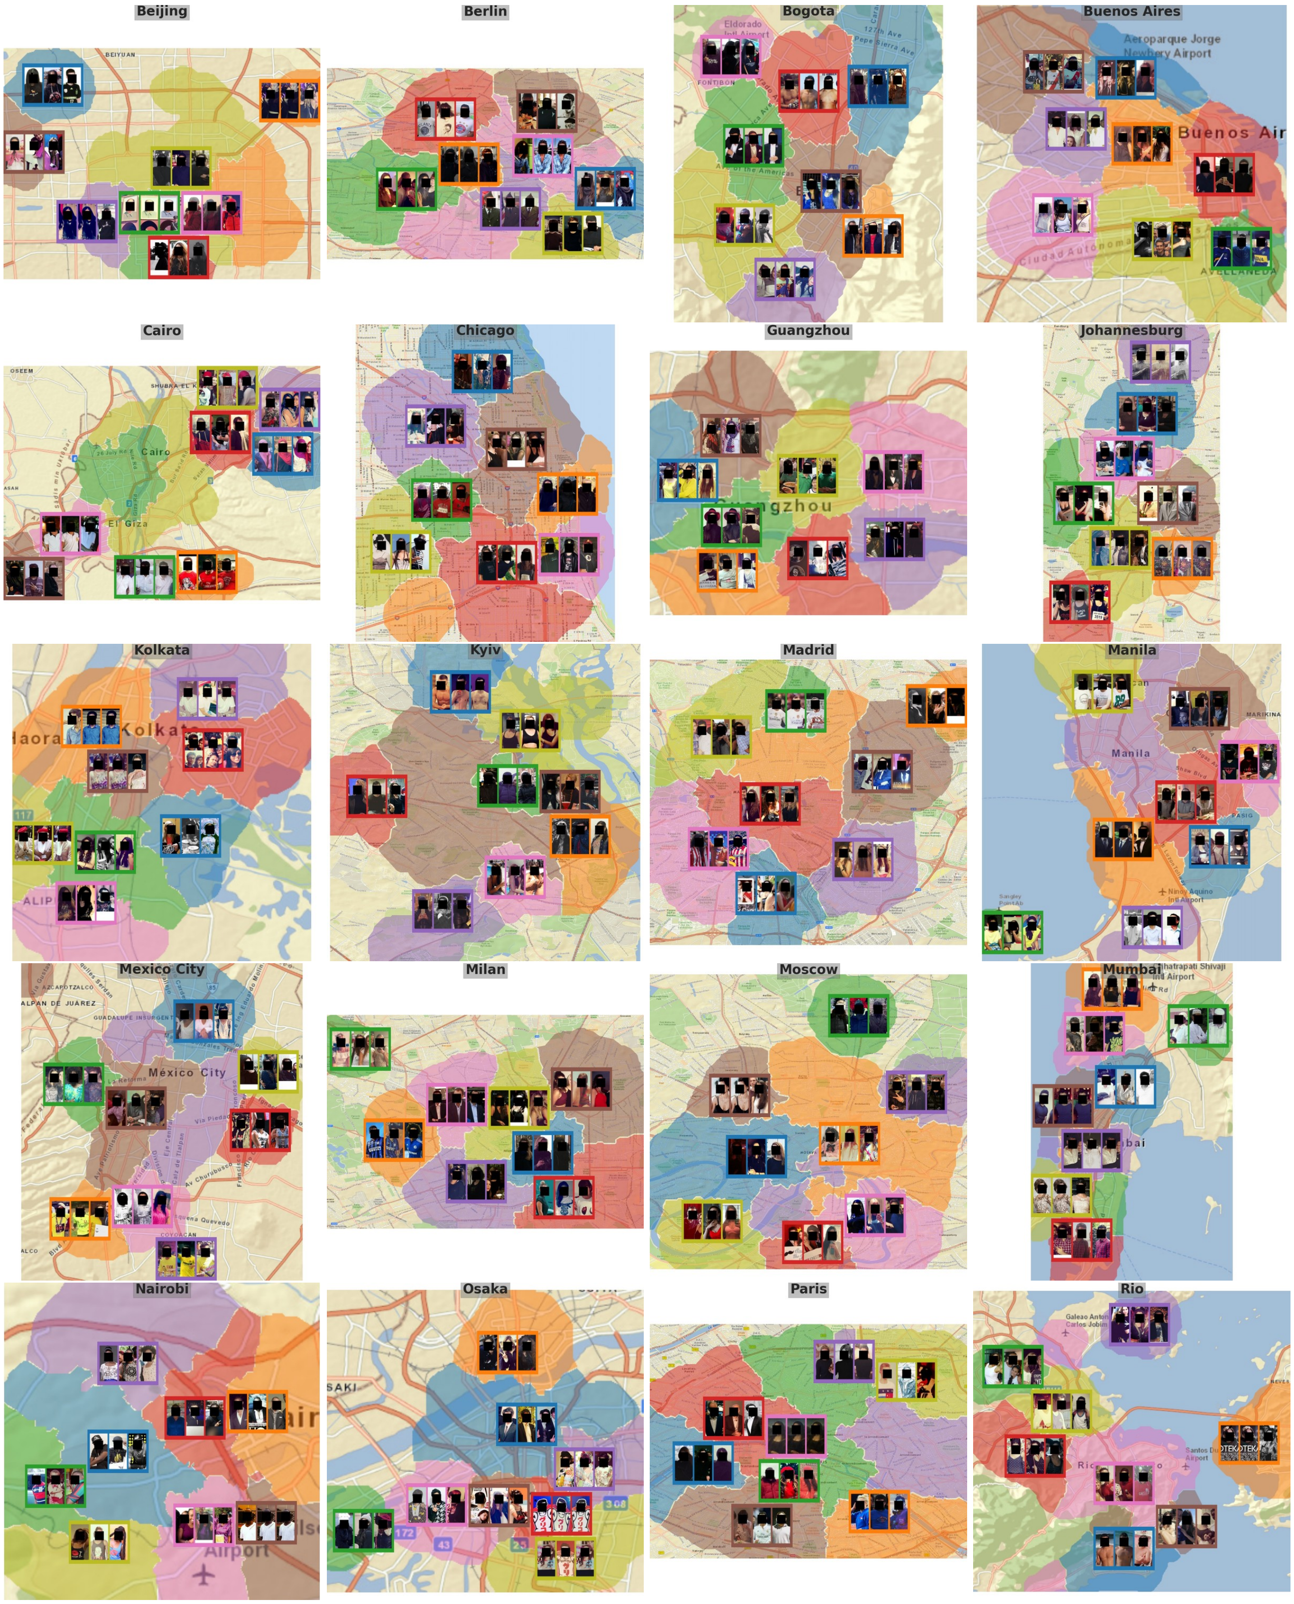

Figure 7: Underground maps for 20 of the cities. Our model uses fashion to discover neighborhoods of cities. For example, in Madrid (green and pink) and London (brown and green) we discover neighborhoods with sports arenas. We also discover tourists areas in cities such as Beijing (red).

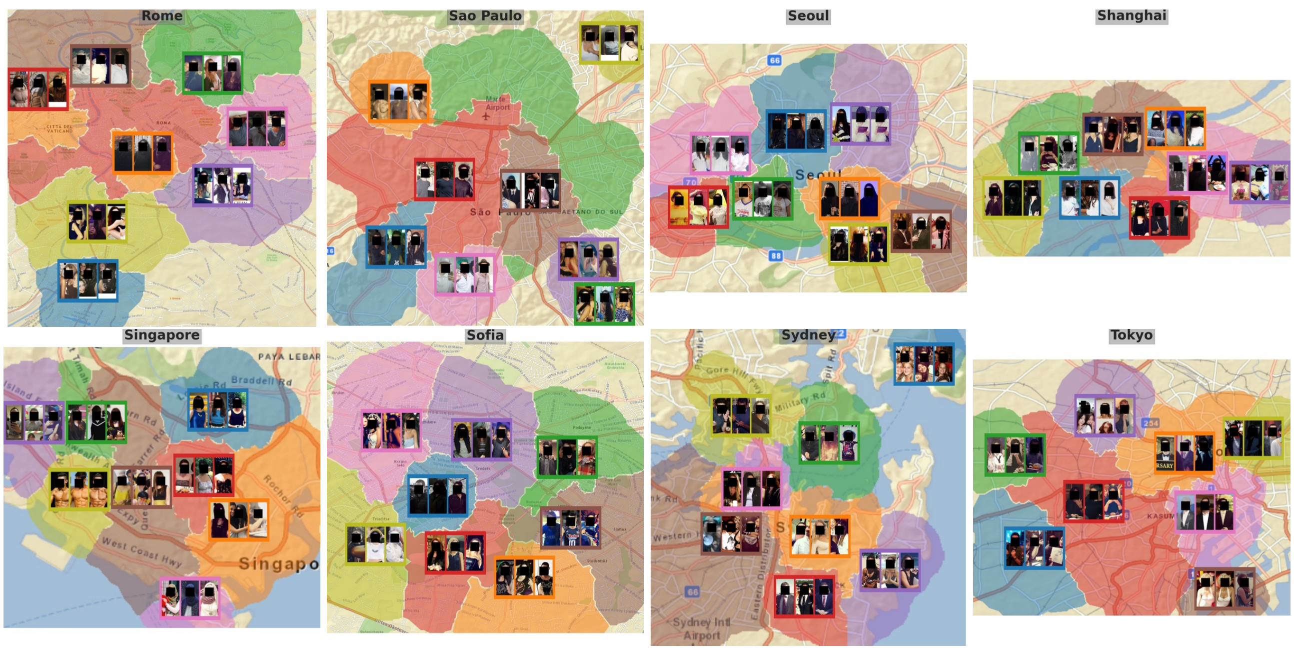

Figure 8: Underground maps for 8 more cities. We discover popular beaches in Rio (blue) and Sydney (purple). We also discover two geographically apart tourist regions in Rome (orange).

Unique Neighborhood/City	Images/Top Attributes	Unique Neighborhood/City	Images/Top Attributes
 Los Angeles	 major color::Blue wearing hat::Yes sleeve length::Short sleeve clothing category::T-shirt clothing pattern::Graphics	 Milan	 major color::Red clothing pattern::Graphics wearing hat::Yes clothing category::T-shirt major color::Blue
 Shanghai	 major color::Green clothing pattern::Spotted major color::Red major color::Gray clothing category::Sweater	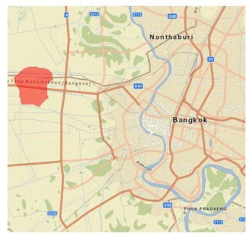 Bangkok	 clothing pattern::Striped major color::More than 1 color clothing pattern::Plaid clothing category::Dress sleeve length::No sleeve
 Bogota	 wearing scarf::Yes wearing glasses::Yes clothing category::Outerwear wearing jacket::Yes major color::Gray	 Kyiv	 wearing hat::Yes major color::Green major color::Orange wearing necktie::Yes neckline shape::Folded
 Chicago	 major color::Red clothing pattern::Graphics clothing category::T-shirt collar presence::No clothing category::Sweater	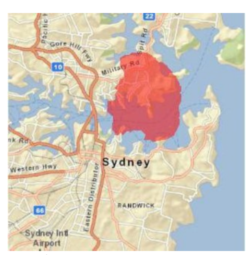 Sydney	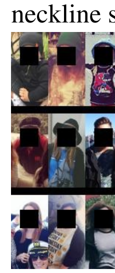 wearing hat::Yes major color::Pink major color::Blue wearing scarf::Yes major color::More than 1 color

Unique Neighborhood/City	**Images/Top Attributes**	**Unique Neighborhood/City**	**Images/Top Attributes**
 Rome	 clothing category::Tank top sleeve length::No sleeve wearing hat::No major color::Black neckline shape::V-shape	 Nairobi	 major color::Brown major color::White wearing scarf::Yes clothing pattern::Striped major color::Pink
 Nairobi	 sleeve length::Short sleeve major color::Green clothing category::T-shirt clothing category::Outerwear neckline shape::V-shape	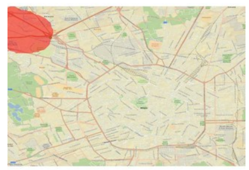 Milan	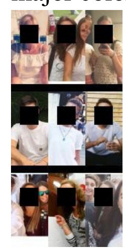 wearing glasses::Yes major color::More than 1 color major color::Cyan major color::Purple wearing necktie::No
 Nairobi	 major color::Blue wearing glasses::Yes neckline shape::Round major color::Red clothing category::Tank top	 Singapore	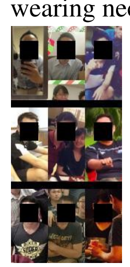 major color::Gray sleeve length::Short sleeve major color::Green multiple layers::One layer clothing category::T-shirt
 Moscow	 wearing hat::Yes major color::Orange collar presence::No major color::Blue wearing necktie::No	 Beijing	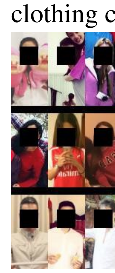 clothing category::Shirt clothing pattern::Spotted major color::White clothing pattern::Graphics wearing hat::No

Unique Neighborhood/City	Images/Top Attributes	Unique Neighborhood/City	Images/Top Attributes
 Seattle	 wearing hat::Yes major color::Green clothing pattern::Graphics major color::Yellow clothing category::T-shirt	 Dhaka	 major color::Cyan clothing category::Outerwear major color::Black major color::Brown multiple layers::Multiple layers
 Sofia	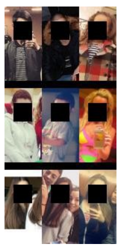 major color::Orange clothing category::Sweater clothing pattern::Floral clothing pattern::Spotted clothing pattern::Plaid	 Cairo	 wearing glasses::Yes major color::Cyan wearing hat::Yes major color::Black major color::More than 1 color
 Madrid	 major color::Purple major color::Cyan neckline shape::Round wearing jacket::No multiple layers::One layer	 Johannesburg	 major color::Cyan neckline shape::Round clothing category::Tank top major color::White collar presence::No
 Kolkata	 multiple layers::Multiple layers clothing category::Outerwear wearing jacket::Yes major color::Purple wearing scarf::Yes	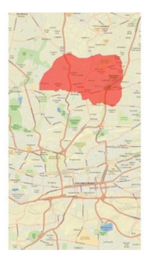 Johannesburg	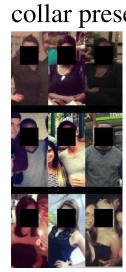 clothing category::Dress clothing pattern::Plaid major color::Orange clothing category::Shirt wearing glasses::No

Unique Neighborhood/City	Images/Top Attributes	Unique Neighborhood/City	Images/Top Attributes
 Kyiv	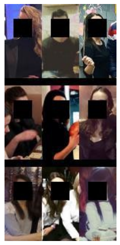 major color::Blue clothing pattern::Plaid major color::Cyan wearing glasses::No major color::Black	 Dhaka	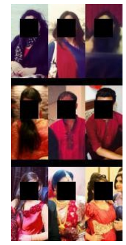 wearing scarf::Yes major color::Orange neckline shape::V-shape major color::Red major color::Pink
 Sydney	 wearing necktie::Yes clothing category::Suit major color::Yellow major color::Red collar presence::Yes	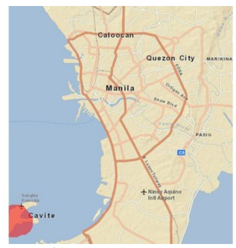 Manila	 major color::Yellow major color::Blue major color::Pink wearing hat::Yes clothing pattern::Floral
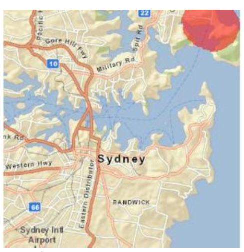 Sydney	 neckline shape::Round clothing pattern::Floral wearing jacket::No collar presence::No multiple layers::One layer	 Sydney	 major color::Brown clothing category::Tank top sleeve length::No sleeve neckline shape::V-shape multiple layers::One layer
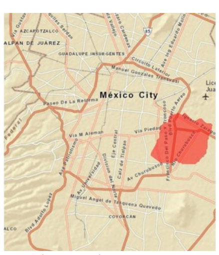 Mexico City	 clothing pattern::Graphics wearing necktie::No wearing glasses::Yes clothing category::T-shirt wearing hat::Yes	 Sofia	 major color::Green major color::Gray major color::Brown major color::Yellow clothing category::Sweater

Unique Neighborhood/City	Images/Top Attributes	Unique Neighborhood/City	Images/Top Attributes
 Tokyo	 major color::More than 1 color clothing category::Outerwear major color::Blue clothing pattern::Plaid wearing glasses::Yes	 Seattle	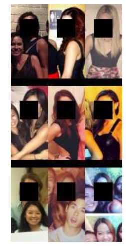 neckline shape::V-shape clothing category::Suit wearing necktie::Yes sleeve length::No sleeve clothing category::Dress
 Madrid	 sleeve length::Short sleeve clothing pattern::Graphics collar presence::Yes clothing category::T-shirt major color::White	 Chicago	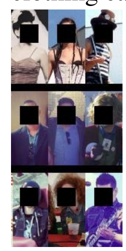 major color::Yellow wearing glasses::Yes major color::White clothing pattern::Plaid sleeve length::Short sleeve
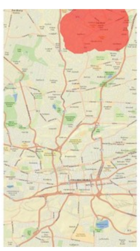 Johannesburg	 neckline shape::V-shape wearing hat::No clothing pattern::Spotted wearing scarf::No collar presence::Yes	 Beijing	 major color::More than 1 color wearing glasses::Yes clothing pattern::Striped wearing scarf::Yes clothing category::Outerwear
 NYC	 wearing glasses::Yes collar presence::No wearing necktie::No clothing category::Tank top neckline shape::V-shape	 Johannesburg	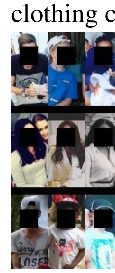 major color::Orange wearing glasses::Yes multiple layers::One layer clothing category::Tank top wearing jacket::No

Table 4: (Left to right, then top to bottom) List of top 40 unique neighborhoods. Top attributes are the relatively most frequent attributes in a neighborhood. For instance, the second example shows the neighborhood of Milan with two famous football clubs with their stadiums and people wearing colors of the clubs. The fourth example shows a neighborhood in Bangkok with white and striped dresses.

Neighborhood/City	Images	Similar Neighborhood/City	Images	Top Attributes
 Kyiv		 Moscow	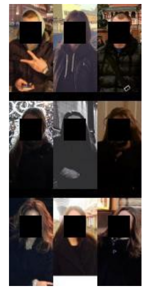	major color::Black major color::Brown neckline shape::Folded multiple layers::Multiple layers wearing jacket::Yes
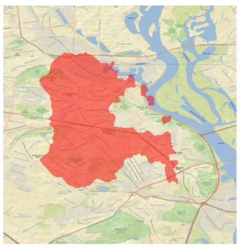 Kyiv	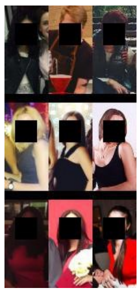	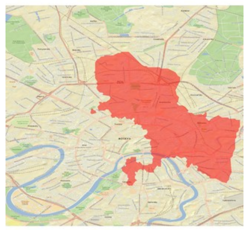 Moscow		clothing category::Shirt clothing pattern::Plaid collar presence::Yes wearing scarf::No neckline shape::V-shape
 NYC		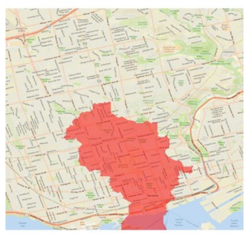 Toronto	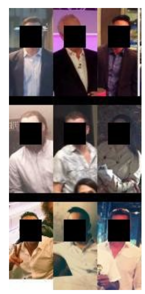	wearing hat::No clothing pattern::Spotted clothing pattern::Solid major color::Black clothing category::Suit
 Chicago		 NYC	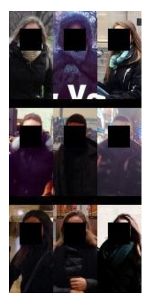	wearing scarf::Yes neckline shape::Folded wearing jacket::Yes multiple layers::Multiple layers clothing category::Outerwear
 Chicago	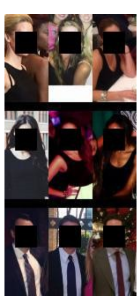	 NYC		neckline shape::V-shape wearing hat::No clothing pattern::Solid clothing pattern::Spotted clothing category::Dress

Neighborhood/City	Images	Similar Neighborhood/City	Images	Top Attributes
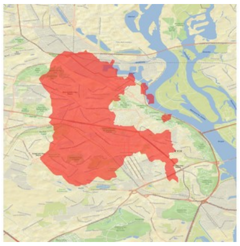 Kyiv	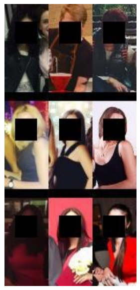	 Moscow	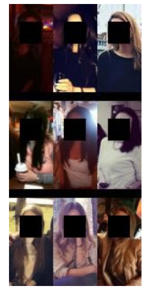	major color::Black clothing pattern::Solid clothing category::Shirt neckline shape::V-shape clothing category::Dress
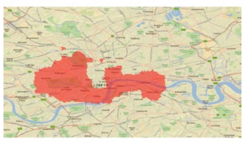 London		 Paris		clothing category::Suit wearing necktie::Yes collar presence::Yes clothing pattern::Solid multiple layers::Multiple layers
 Chicago	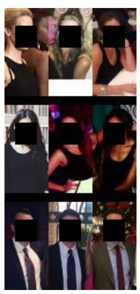	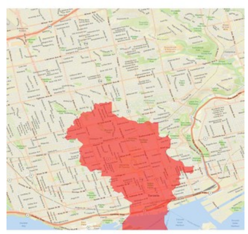 Toronto	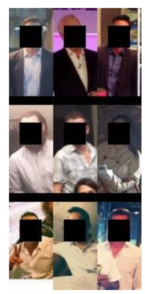	clothing category::Suit wearing hat::No collar presence::Yes wearing necktie::Yes neckline shape::V-shape
 Chicago		 Seattle		wearing scarf::Yes neckline shape::Folded sleeve length::Long sleeve wearing jacket::Yes multiple layers::Multiple layers
 Osaka		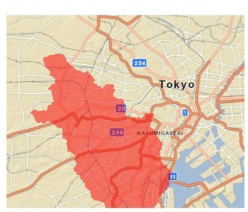 Tokyo		clothing category::Dress clothing category::Sweater major color::Black clothing pattern::Solid neckline shape::V-shape

Neighborhood/City	Images	Similar Neighborhood/City	Images	Top Attributes
 Milan	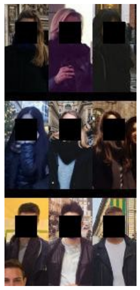	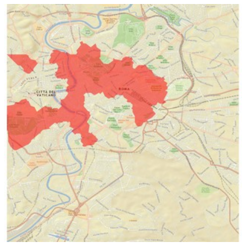 Rome	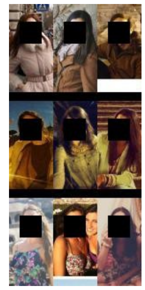	major color::Brown sleeve length::Long sleeve wearing jacket::Yes multiple layers::Multiple layers wearing scarf::Yes
 Madrid		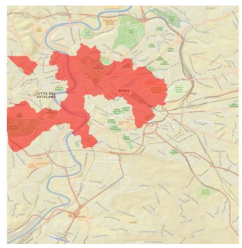 Rome	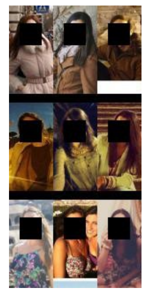	wearing glasses::Yes wearing scarf::Yes major color::Brown clothing category::Outerwear multiple layers::Multiple layers
 Buenos Aires		 Madrid		wearing scarf::Yes sleeve length::Long sleeve multiple layers::Multiple layers wearing jacket::Yes wearing glasses::Yes
 London		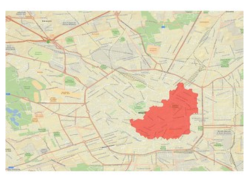 Milan	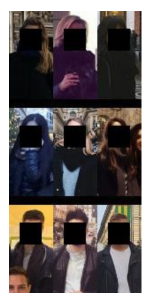	sleeve length::Long sleeve multiple layers::Multiple layers major color::Brown wearing jacket::Yes wearing scarf::Yes
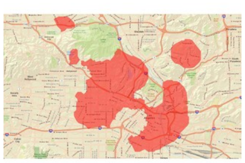 Los Angeles		 Sydney	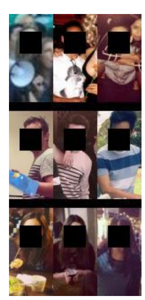	clothing pattern::Plaid sleeve length::Long sleeve clothing category::Outerwear clothing category::Shirt major color::Black

Neighborhood/City	Images	Similar Neighborhood/City	Images	Top Attributes
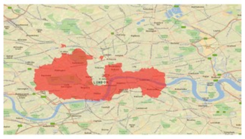 London		 Paris	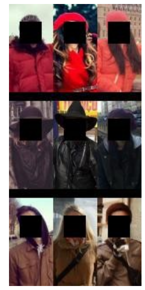	wearing scarf::Yes major color::Brown multiple layers::Multiple layers clothing category::Outerwear wearing glasses::Yes
 London		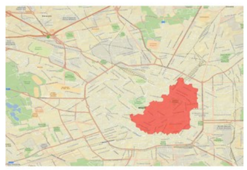 Milan	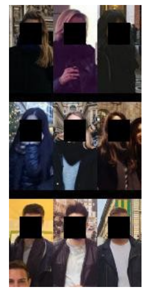	sleeve length::Long sleeve major color::Brown wearing jacket::Yes multiple layers::Multiple layers wearing scarf::Yes
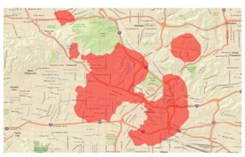 Los Angeles		 NYC	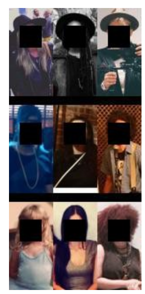	major color::Gray major color::Brown major color::More than 1 color clothing pattern::Plaid clothing category::Sweater
 Berlin		 Paris		major color::More than 1 color major color::Cyan clothing pattern::Striped major color::Gray clothing pattern::Plaid
 Los Angeles		 NYC		neckline shape::V-shape clothing pattern::Solid clothing pattern::Spotted major color::Black clothing category::Dress

Table 5: List of top 20 most similar neighborhoods. Top attributes are the relatively most frequent attributes common in both the neighborhoods. For instance, the fourth row, shows that the tourist regions of Chicago and NYC are similar.

Neighborhood/City	Images		Contextually Similar Neighborhood/City	Images	Similar Neighborhood	Images
 Istanbul **Common Attributes:**		↔	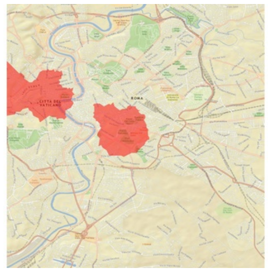 Rome wearing scarf::Yes wearing hat::Yes	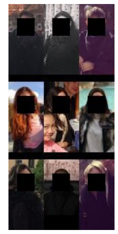	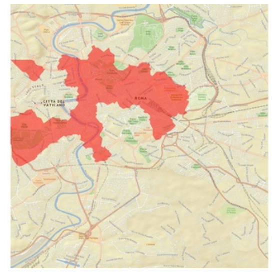 Rome wearing glasses::Yes clothing category::Outerwear	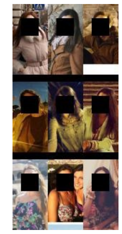
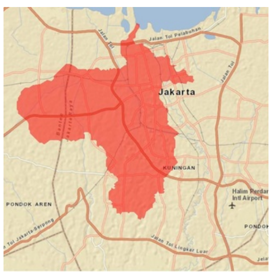 Jakarta **Common Attributes:**		↔	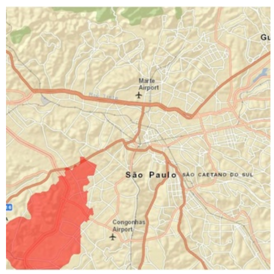 Sao Paulo collar presence::Yes clothing category::Shirt		 Sao Paulo clothing pattern::Solid major color::Black	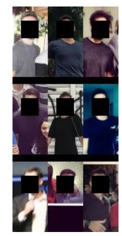
 Los Angeles **Common Attributes:**		↔	 Singapore neckline shape::V-shape sleeve length::No sleeve		 Singapore clothing category::Dress clothing pattern::Solid	
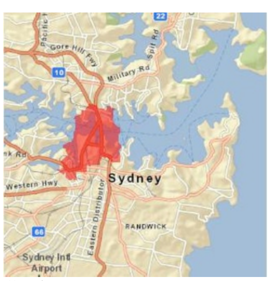 Sydney **Common Attributes:**		↔	 Sao Paulo clothing pattern::Solid sleeve length::Long sleeve	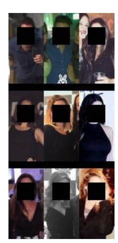	 Sao Paulo major color::Black wearing hat::No	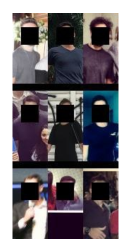

Neighborhood/City	Images		Contextually Similar Neighborhood/City	Images	Similar Neighborhood	Images
 NYC **Common Attributes:**		↔	 Mexico City neckline shape::V-shape wearing hat::No		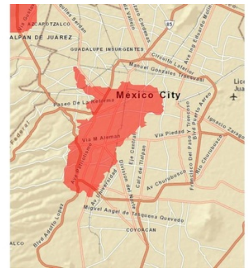 Mexico City clothing pattern::Solid major color::Purple	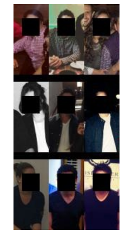
 Berlin **Common Attributes:**		↔	 Rome wearing glasses::Yes multiple layers::Multiple layers	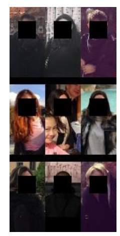	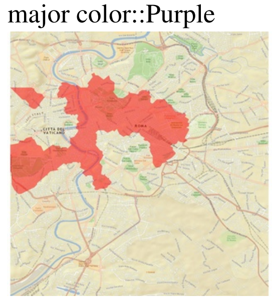 Rome wearing scarf::Yes clothing category::Outerwear	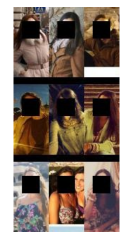
 Osaka **Common Attributes:**		↔	 Bogota clothing pattern::Solid wearing hat::No	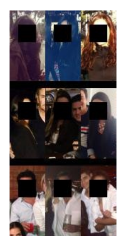	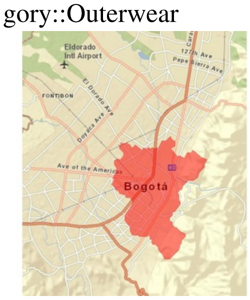 Bogota clothing category::Shirt clothing category::Suit	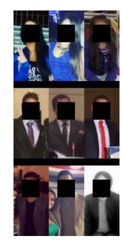
 Buenos Aires **Common Attributes:**		↔	 Berlin sleeve length::Long sleeve neckline shape::Folded	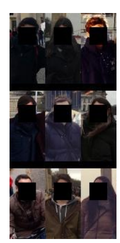	 Berlin multiple layers::Multiple layers wearing jacket::Yes	

Neighborhood/City	**Images**		**Contextually Similar Neighborhood/City**	**Images**	**Similar Neighborhood**	**Images**
 Singapore **Common Attributes:**		↔	 Mexico City neckline shape::V-shape clothing category::Suit	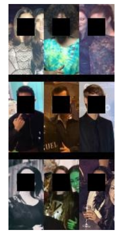	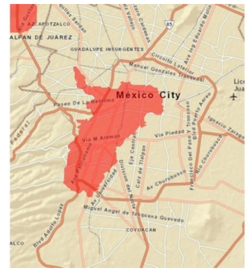 Mexico City clothing category::Dress clothing pattern::Floral	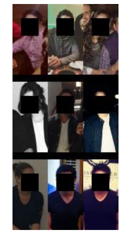
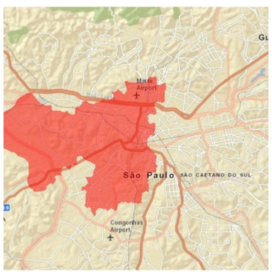 Sao Paulo **Common Attributes:**		↔	 Buenos Aires major color::Gray multiple layers::Multiple layers	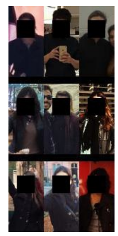	 Buenos Aires sleeve length::Long sleeve wearing jacket::Yes	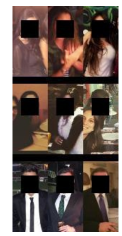
 Seoul **Common Attributes:**		↔	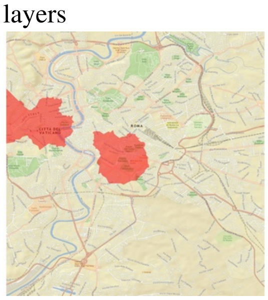 Rome wearing scarf::Yes clothing category::Outerwear	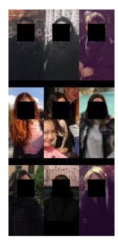	 Rome wearing glasses::Yes neckline shape::Folded	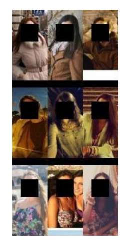
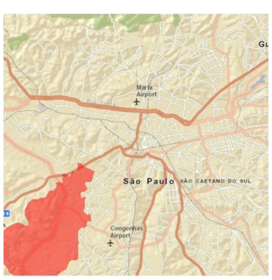 Sao Paulo **Common Attributes:**	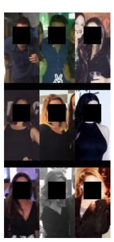	↔	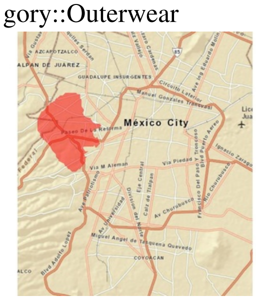 Mexico City neckline shape::V-shape wearing hat::No	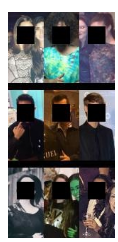	 Mexico City clothing pattern::Solid clothing category::Shirt	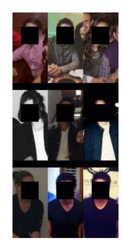

Neighborhood/City	**Images**		**Contextually Similar Neighborhood/City**	**Images**	**Similar Neighborhood**	**Images**
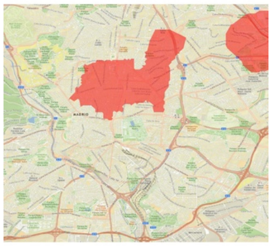 Madrid **Common Attributes:**		↔	 Mexico City neckline shape::V-shape wearing hat::No	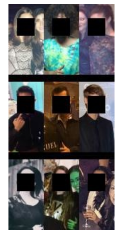	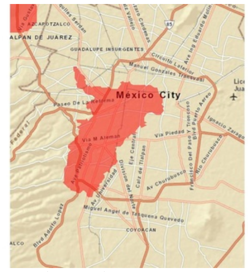 Mexico City clothing category::Suit wearing necktie::Yes	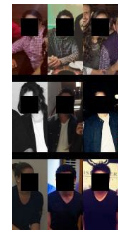
 Bogota **Common Attributes:**		↔	 Singapore clothing pattern::Solid clothing category::Dress		 Singapore wearing hat::No clothing pattern::Floral	
 Moscow **Common Attributes:**	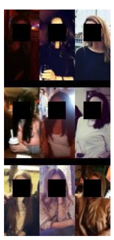	↔	 Sao Paulo clothing pattern::Solid wearing hat::No		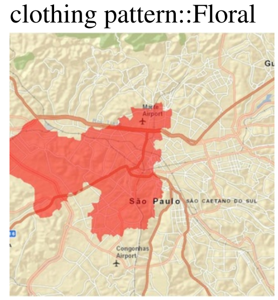 Sao Paulo major color::Black major color::Brown	
 Rio **Common Attributes:**		↔	 Buenos Aires collar presence::Yes sleeve length::Long sleeve	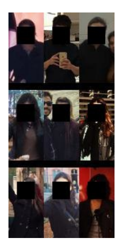	 Buenos Aires neckline shape::Folded multiple layers::Multiple layers	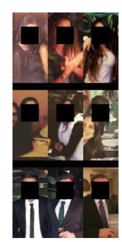

Neighborhood/City	Images		Contextually Similar Neighborhood/City	Images	Similar Neighborhood	Images
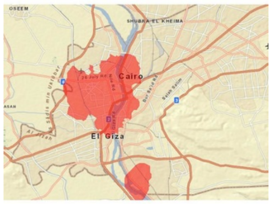					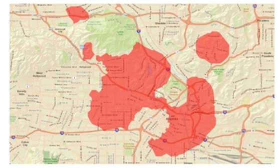	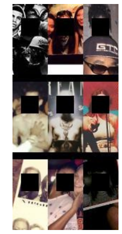
Cairo **Common Attributes:**		↔	Los Angeles neckline shape::V-shape wearing scarf::No		Los Angeles clothing pattern::Solid sleeve length::No sleeve	
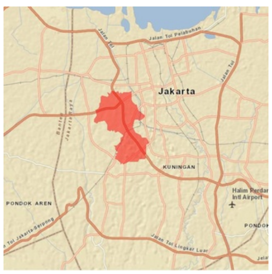					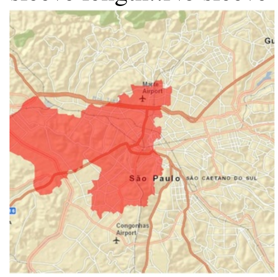	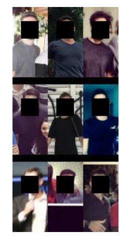
Jakarta **Common Attributes:**		↔	Sao Paulo major color::Black clothing pattern::Plaid		Sao Paulo clothing category::Shirt collar presence::Yes	
					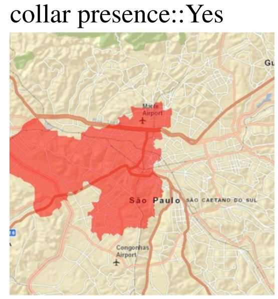	
Manila **Common Attributes:**		↔	Sao Paulo major color::Brown major color::Black		Sao Paulo clothing pattern::Solid clothing category::Shirt	
				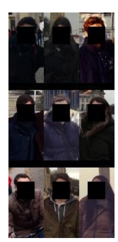		
London **Common Attributes:**		↔	Berlin neckline shape::Folded wearing scarf::Yes		Berlin multiple layers::Multiple layers wearing jacket::Yes	

Table 6: Top 20 most contextually similar neighborhoods (left pairs), where simpler similarity results in a different neighborhood (right). As can be seen from the examples, this is typically the case when 2 cities are geographically and culturally far apart. For instance, in the first example two tourist neighborhoods of Istanbul and Rome are correctly identified as similar, but they are not found without contextual encoding. Top attributes are the relatively most frequent attributes common in the two neighborhoods.

8. Grouping step

The grouping step in Sec. 3.3 mapping locations to neighborhoods via their style encodings is simple but effective—more effective than other more elaborate variants we explored. For example, we tried affinity propagation instead of K-means for clustering, but its neighborhoods tended to perform worse than K-means when adjusted to produce the same number of clusters on our benchmarks. We also tried an exponential weighting scheme for feature $\boldsymbol{h_x}$, where the contribution of an image with location $\boldsymbol{l_i}$ to a histogram for location $\boldsymbol{x}$ is weighted inversely by an exponentiation of $||\boldsymbol{l_i} - \boldsymbol{x}||_2$. This can be thought of as a softer version of the features in Sec. 3.2 in the main paper. It performed similarly to hard features, and hence we kept the simple version for evaluation.