Citation for published version:

AlHalawani, S, Yang, Y, Wonka, P & Mitra, NJ 2014, 'What Makes London Work Like London?', *Computer Graphics Forum*, vol. 33, no. 5, pp. 157-165. <https://doi.org/10.1111/cgf.12441>

DOI:

[10.1111/cgf.12441](https://doi.org/10.1111/cgf.12441)

Publication date:

2014

[Link to publication](https://doi.org/10.1111/cgf.12441)

University of Bath

Alternative formats

If you require this document in an alternative format, please contact:
openaccess@bath.ac.uk

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

What Makes London Work Like London?

Sawsan AlHalawani¹

Yong-Liang Yang¹

Peter Wonka¹

Niloy J. Mitra^{2,1}

¹KAUST

²University College London

Abstract

Urban data ranging from images and laser scans to traffic flows are regularly analyzed and modeled leading to better scene understanding. Commonly used computational approaches focus on geometric descriptors, both for images and for laser scans. In contrast, in urban planning, a large body of work has qualitatively evaluated street networks to understand their effects on the functionality of cities, both for pedestrians and for cars. In this work, we analyze street networks, both their topology (i.e., connectivity) and their geometry (i.e., layout), in an attempt to understand which factors play dominant roles in determining the characteristic of cities. We propose a set of street network descriptors to capture the essence of city layouts and use them, in a supervised setting, to classify and categorize various cities across the world. We evaluate our method on a range of cities, of various styles, and demonstrate that while standard image-level descriptors perform poorly, the proposed network-level descriptors can distinguish between different cities reliably and with high accuracy.

---

1. Introduction

Think of central London. Immediately memories of slow moving traffic and bustling streets come to mind. Many such cities have unique characteristics, see Figure 1: some are best enjoyed on foot, some are easier to tour in a car, while others offer a chaotic web of streets posing a traffic nightmare to both tourists and residents. Among all the elements in an urban layout (e.g., streets, buildings, parks, etc.), transportation networks (e.g., street network, train network) form its lifeline, and largely define how the space works and functions. For example, grid networks encourage well-balanced traffic flow; while radial layouts provide easy access to the city centers but lead to irregularly shaped building blocks.

In this paper, we ask if cities can be characterized and identified *directly* based on their street network patterns. To answer this question, we would like to extract features that can capture the most important aspects of street networks. We focus on capturing the functionality of cities, rather than their visual appearance, i.e., how they *work* rather than how they *look*. In the past, a lot of effort went into analyzing images and geometry by proposing appropriate descriptors. However, a direct application of existing feature extraction and analysis frameworks is not suitable for two reasons. First, street networks are a specialization of geometric graphs, and it is important to adapt existing features. Second, in contrast to images and geometric models, very important functional

features are related to questions about paths and transport network efficiency.

In urban planning, particularly in space syntax research, a large body of work qualitatively discusses various factors af-

Figure 1: Cities are known for their different characteristics. In this work, we propose and evaluate various descriptors for street networks, and identify the ones that are particularly discriminative and characterize uniqueness of cities.

Figure 2: Starting from OpenStreetMap (OSM) data for various cities, we propose a set of street-level descriptors to characterize ‘city-ness’ of various cities, and learn the relative importance of the proposed street-level descriptors. Beyond retrieval and grouping, the ranked features can also be used as a measure of uniqueness of cities.

fecting how a city functions. Beyond a few case studies, the relative importance and effectiveness of the various features have not been qualitatively studied. In this work, based on a large-scale analysis, we investigate effectiveness of various street network features. First, we propose a set of descriptors to characterize both the topology and layout of street networks. Then, in a supervised setting, we learn city-specific classifiers and identify the relative importance of the proposed descriptors. The extracted descriptors and their value ranges provide interesting insights into city characteristics, and inform us which features are predominantly important and discriminative. While in practice various secondary factors influence network flow patterns (e.g., street width, peak traffic volume, speed limits), as a first attempt, we focus *only* on the topology and geometric layout of the street networks.

We used the proposed descriptors to analyze 10 cities (London, Paris, New York, Cardiff, etc.) across the world. For each city, we extracted 100 street network patches as GIS measurements (from OpenStreetMap) at three scales (0.25 km², 1 km², and 4 km²) to analyze the primary and secondary streets of the networks. We trained classifiers (SVMs and discriminative analysis), both inter-city and intra-city, based on the proposed network descriptors.

For inter-city classification, density of 4-way crossings and connectivity index play dominant roles; while, for London intra-city classification, intersection and street densities are most discriminative. Based on the corresponding precision-recall curves, we report that the descriptors perform significantly better than state-of-the-art image descriptors working directly on the city (street) maps. We also determine the relative importance of the different proposed features and the importance of scale.

Beyond classification, the proposed descriptors provide the first measures of what characterizes cities at a functional level (although only using the transportation network). This could be used for the generation of similar cities by ensuring the behavior of the proposed areas (e.g., neighborhoods, cities, game levels, etc.). We provide initial results in this direction.

2. Related Work

We review the related work in urban street network modeling and street network analysis in urban planning.

Street network modeling. Urban networks can be generated from scratch by iteratively adding streets to an existing street network [PM01, WMWG09]. The algorithm can be controlled by setting parameters such as a distribution of angles at an intersection, the street length between two intersections, and snapping distances. Procedural modeling can be combined with interactive techniques to have more control over the result [CEW*08, LSWW11]. Yang et al. [YVWV13] propose a framework for street network design that jointly considers the quality of streets and parcels. Another approach to network modeling is data driven. The work of Aliaga et al. [AVB08] investigates the problem of synthesizing urban layouts by example by combining aerial-view images of urban areas with vector-based data describing the street and parcel network. Another work of Aliaga et al. [ABVA08] supports interactively editing an urban layout by translating, rotating, copying, cutting, and pasting (a group of) tiles in order to edit or create a new city arrangement. Most related to our work is the idea to tune the parameters of a procedural model to optimize for certain functional or behavioral characteristics [VABW09, VGDA*12, IGD14]. The modeling of large-scale road networks has been tackled by Galin et al. [GPGB11]. At a lower level, Maréchal et al. [GPGB11] studied how a single road adapts to the terrain.

Street network analysis and urban planning. There is a large amount of urban planning literature that addresses different aspects of street network planning and analysis. Most literature is concerned with discussing goals and relevant factors on a high level and the descriptions are abstract [AIS77, Lyn60, SBJ03, SO93, Mar05]. The urban planning literature itself rarely includes specific details on computation for generating new street networks or analyzing existing street networks. The most specific results are typically design standards [Ass06] and some simple statistics used to describe road networks (e.g. the number of city blocks per area). We use these statistics in our work. A typical strategy for analysis in these books is to argue by example. For example, multiple authors categorize street patterns into multiple

categories. Each category is explained by one or multiple examples and a textural description. Unfortunately, a lot more information is necessary to perform an actual classification by automatic computation. This problem is also recognized in the literature, e.g. Marshall [Mar05] chapter 4. Another attempt uses street-view images to identify distinctive image-space facade elements, and use them to distinguish images of Paris from other cities [DSG*12].

Several researchers in urban planning attempted a more formal and computational analysis of street networks. A lot of these papers stem from the seminal work in space syntax [Hil96]. Two prominent ideas in this work are to consider visibility and turning angles. For example, researchers conjecture that a person is more likely to walk straight rather than taking turns and therefore they augment distance measurements by considering the number of turns a person has to take. We also integrate an adapted version of reach and directional reach from a recent paper [JP08] in our work. The study of smaller scale aspects or street networks typically falls in the realm of engineering. Therefore, the study of geometric road details (e.g. highway on-ramps, intersections, turning path of trucks) and the study of local traffic flow (e.g. one vehicle following another in a single lane) have been well investigated and there are a larger number of mathematical models available [Boa10].

3. Overview

Given multiple street networks $\{L_1, L_2, \dots, L_K\}$ encoded in standard GIS data format, our goal is to characterize the cities at the street network level based on a set of proposed features (see Figure 2). In this work, we propose both topological and geometric features for such network layouts. In the process, we are interested in two key questions: (i) what are good descriptors for classifying street networks; and (ii) what applications are possible using the proposed feature descriptors.

In a preprocessing step (see Section 4), we parse input GIS data from 1000 map tiles (with sizes of 0.25 km^2 , 1 km^2 , and 4 km^2) and extract the underlying street information (e.g., start and end points, sequential road segments, etc.) along with block information (i.e., areas enclosed by streets, street length, etc.). Each map tile gets mapped to a high-dimensional feature vector (see Section 5). Then, in a supervised learning setting (using SVMs and Discriminative analysis), we identify the distinctive features and the uniqueness of the various cities (see Section 6), both at inter-city and intra-city (for London) level. The analysis provides insights about characteristic network-level features and enables novel retrieval possibilities along with rating the validity of synthesized networks (see Section 7).

4. Collecting City Street Network Data

We collected urban layout patches for 10 different cities (Beijing, Camp Durant, Cardiff, London, Los Angeles, Moscow, New York, Paris, Toronto, and Vienna) from OpenStreetMap to represent different urban patterns, e.g., grid-like, radial, hierarchical, and curved (see Figure 1). For each city, we collected 100 layouts at three different scales (0.25 km^2 , 1 km^2 , and 4 km^2). These layouts come in the '.osm' files under standard XML format. We use *readosm*, an open source library to parse the input to extract street information. We process a total of these 1000×3 map tiles in the subsequent analysis. In this paper, we focus on the primary and secondary streets, while ignoring other connections namely 'stairs', 'footways', 'bike paths.' The whole parsed dataset can be downloaded from the project page: <http://geometry.cs.ucl.ac.uk/projects/2014/whatMakesLondon/>.

For each layout, we obtain a set of nodes (i.e., intersections) $\{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_N\}$ described by their 2D location, together with the individual street connections $\{r_1, r_2, \dots, r_M\}$.

Based on the node set and the connectivity of individual streets, we construct a graph $G = \{V, E\}$ for the entire street layout. The vertex set is defined as $V := \{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_N\}$ and the edge set as $E := \{e_{ij}\}$ where $e_{ij} = \overline{\mathbf{v}_i \mathbf{v}_j}$ denotes a street segment.

5. Features for Street Networks

In this section, we propose a set of features, both topological and geometric, for characterizing the underlying street networks. Many of the proposed features are motivated by qualitative measures that have been proposed in space syntax and urban planning literature (see Section 2). We utilize topological features to capture the street connectivity, and geometric features to encode the physical layout within the region. Note that only later analysis reveals the relative importance for different tasks along with their right scales.

5.1. Topological (Connectivity) Analysis

Urban street networks exhibit a graph structure, making connectivity a unique signature to represent the topology characteristics of a street network. Clearly, one can improve access to a place by adding roads to it. This scenario is regularly observed in the real world as cities continue to grow; or, densely versus sparsely inhabited cities. Similarly, in a radial layout, the center often has many connecting streets so that it can be easily accessed. On the other hand, residential areas usually use dead-end roads to protect neighborhood privacy, and T-junctions to reduce through-traffic. All the measures are normalized by area of the map tile.

Figure 3: Analysis of the local connectivity of a layout based on valence statistics over its nodes, i.e., intersections. Later, our analysis reveals, that valence distribution is one of the most important street-network features.

Valence (t_{val}). The simplest connectivity measure is based on the valence of a vertex. We characterize the local connectivity of graph G based on valence statistics over dead ends ($val = 1$), T-junctions ($val = 3$), 4-way crossings ($val = 4$), and others ($val > 4$), denoted by $t_{v=1}(G)$, $t_{v=3}(G)$, $t_{v=4}(G)$ and $t_{v>4}(G)$ respectively. For simplicity, we use t_{val} to represent the vector of four numbers. Figure 3 shows vertices with different valences in a patch layout.

Street density (t_{sd}). A street or link is defined as a sequence of graph edges that connects two graph nodes. Note that both the graph nodes should have valence different from 2. We measure density t_{sd} as the total number of streets (links) within a map tile, which has fixed area.

Connectivity index (t_{ci}). Connectivity index measures how well a roadway network connects destinations. We define the connectivity index of a graph as the ratio of number of (connecting) links to the number of nodes, i.e.,

$$t_{ci}(G) = \sum_i val(\mathbf{v}_i) / \text{number of nodes}. \quad (1)$$

Please note that here we only consider intersection nodes and dead end nodes, i.e., the node valence is not 2. Note that urban design guidelines often prescribe a minimum connectivity index of 1.4.

Intersection density (t_{id}). measures the number of intersections within a patch layout, i.e., the total number of graph nodes with valence higher than 2.

4-way crossing proportion (t_{cp}). is the proportion of 4-way crossings with respect to all the intersections. For a Manhattan network, this ratio tends to 1.

Number of blocks (t_{nb}). In any urban layout, a block is a polygonal region bounded by streets. This measures how many blocks reside in a patch layout (of fixed area). Essentially, it is the dual of the street network graph.

5.2. Geometric (Layout) Analysis

Street connectivity can effectively measure whether a place can be easily accessed or not. However, it fails to capture the layout behavior within a region. For example, how easily one can travel from one place to another. We therefore also consider geometric properties of a street network that are complementary to topology features.

Total street length (g_{sl}). Transportation efficiency in a patch layout is directly influenced by the total length of the streets $g_{sl}(G)$. For example, a crowded region usually has dense roads, while sparse roads are common across rural areas.

Average street length (g_{al}). is the average street/link length between any two graph nodes whose valences are not 2.

Figure 4: A comparison between the Euclidean shortest path (in blue) against the shortest path along the road network connecting the source and destination nodes (in red).

Transportation convenience (g_{tc}). measures how easy it is to travel from source $\mathbf{s}$ to destination $\mathbf{d}$. It is obvious that the most convenient way to travel is along the Euclidean shortest path (we assume the travel time is proportional to the path length). However, this is not an option if there is no straight path directly connecting $\mathbf{s} \rightarrow \mathbf{d}$. In a grid-like street network, the length of the shortest path follows the well-known Manhattan distance. For a general layout, we compute the shortest path using Dijkstra's algorithm. Figure 4 shows the two different paths on a sample map. Suppose the Euclidean distance between $\mathbf{s} \rightarrow \mathbf{d}$ is $d_E(\mathbf{s}, \mathbf{d})$ and the shortest graph distance is $d_D(\mathbf{s}, \mathbf{d})$, the transportation convenience between $\mathbf{s}$

Figure 5: Topological and geometric features for the same layout: (a) Valence t_{val} , red ($val = 1$), purple ($val = 3$), green ($val = 4$), and yellow ($val > 4$); (b) Redundancy g_r for two pairs of nodes; (c) Metric reach g_{mr} with $x = 0.5$ km; and (d) Travel distance histogram g_{th} for the 25 sampled nodes.

and $\mathbf{d}$ is defined as:

$$g_{tc}(\mathbf{s}, \mathbf{d}) := d_E(\mathbf{s}, \mathbf{d}) / d_D(\mathbf{s}, \mathbf{d}). \quad (2)$$

We estimate the global transportation convenience $g_{tc}(G)$ for the entire layout by sampling several random source and destination locations (1000 in our experiments) and compute the average. To avoid local estimation, the source and destination are required to have certain minimal distance ($d_E(\mathbf{s}, \mathbf{d}) > 0.25$ km).

Redundancy (g_r). Given a pair of nodes, source $\mathbf{s}$ and destination $\mathbf{d}$, we propose to measure *redundancy* of the network as the number of different routes connecting $\mathbf{s} \leftrightarrow \mathbf{d}$. We compute this metric based on a breath-first-search from $\mathbf{s}$ to $\mathbf{d}$, and then back-track all the paths from $\mathbf{d}$ to $\mathbf{s}$. We filter out the inefficient routes identified if their length is more than twice of the shortest graph distance. We estimate redundancy of the whole layout as the average of the top 100 pairs of nodes with the largest Euclidean distance. This measure captures the ease of rerouting traffic without incurring severe (time) penalty. As we see later, this descriptor can help distinguish between busy and less busy cities (e.g., New York versus Camp Durant). See Figure 5b.

Metric reach (g_{mr}), which is commonly discussed in the urban planning literature, measures the total length of streets if one goes a total of x km (0.5 km in our implementation) along all possible directions. Essentially, it is the total length of a street that one can access from any point as long as the total length is x km (without loops). See Figure 5c.

Travel distance histogram (g_{th}) captures the distribution of travel distances over the entire layout. (We designed this descriptor inspired by the shape distributions [OFCD02] in the case of meshes.) We uniformly generate 5×5 sample points and project them to the nearest graph nodes. Then, we compute the shortest graph distance for each pair of the nodes. The travel distance histogram is built based on the shortest distances over all the pairs and quantized into 10 bins and

each of which has the same distance range across all the layouts. See Figure 5d.

Figure 6: Grouping cities based on a 2D-embedding of the corresponding features. All features are used in this experiment without confidence weighting. A random sampling of cities is shown although the embedding is generated using all the images from the respective cities.

6. Layout Classification

Based on the extracted topological and geometric features, we construct a feature vector $\mathbf{f}$ for each layout L as:

$$\mathbf{f}(L) := [t_{val}, t_{sd}, t_{ci}, t_{id}, t_{cp}, t_{nb}, g_{sl}, g_{al}, g_{tc}, g_r, g_{mr}, g_{th}]. \quad (3)$$

The dimension of the feature vector is 24, given that t_{val} has 4 components and g_{th} has 10 bins. In all the subsequent analysis, we normalize the range for the different descriptors based on the extent found in the corresponding training sets (see later). Figure 6 shows a 2D embedding of all the features for the 10 cities using multi-dimensional scaling (MDS). Note that at the smaller scale (top figure), New York clearly stands out, while London and Paris are easy to confuse. The embedding with more cities (bottom figure) shows a hierarchical nature, where some cities are more different from others (see discussion on the confusion matrix in Section 7), but New York remains quite distinctive.

The input map data have associated labels: (i) city name, and (ii) city region. For city region, we manually labeled the input map tiles in the training set into 3 zones: *downtown*, *mid-zone*, and *outskirts*.

We employ supervised learning using Support Vector Machine (SVM) [CST00] and Discriminative Analysis (DA) [Fis36] linking the feature vectors and labels. For SVM, we used two types of kernels: linear and radial basis functions (RBF). In SVM with RBF we extracted parameters by 10-fold cross validation using different partitions of the training set. For DA, we used linear and quadratic versions.

For training, we experimented with 5%, 10%, 20%, and 50% of the input data. We found the classifiers to have similar performance beyond 5% (see Section 7). Hence, as default, we used 10% training data, unless mentioned. Multi-class classifiers performed better compared to class-specific classifiers (see Figure 7), but the general accuracy was higher for binary classifiers (see Table 1). Note that the analysis was performed in the $\mathbf{f}(L)$ space, rather than the 2D-embedded space, which is only used for visualization.

7. Evaluation and Discussion

In this section, we first rate the proposed features, discuss the effect of scale of the selected map tiles, and then use the proposed features for analysis and synthesis applications.

How about image features? We first rate the proposed network features against image-level features. We take the parsed street networks and rasterize them to form image patches. We use bag-of-word features with GIST descriptor [TMFR03]. Clearly, GIST fails to detect network level changes, while our topological descriptors are sensitive to such changes. We also tested GIST versus the proposed descriptors $\mathbf{f}(L)$ for classification in supervised settings and found the proposed descriptors to be consistently better. For

Table 1: Cross validation precision (%) using different classification methods with 10% data as training size. SVML and SVMR refer to SVM linear and SVM RBF, respectively. LDA and QDA refer to linear discriminant analysis and quadratic discriminant analysis, respectively. Note that for binary classifiers, higher precision comes at the cost of lower recall (refer to Figure 7).

classification type	SVML	SVMR	LDA	QDA
multi-class	35.8	41.7	47.3	53.4
Beijing and others	90.2	90.6	88.7	76.0
Camp Durant and others	90.2	91.1	88.8	89.9
Cardiff and others	91.1	91.3	83.8	84.1
LA and others	88.8	89.0	82.0	83.0
London and others	90.8	91.2	86.6	88.4
Moscow and others	89.3	89.1	80.3	83.7
NY and others	93.6	94.7	90.7	90.5
Paris and others	92.9	93.2	87.2	89.4
Toronto and others	89.0	89.1	67.9	66.2
Vienna and others	90.0	90.0	74.6	74.0

example, using SVM-linear, binary classification accuracy suffered by 20% or more for each of the 10 cities.

Figure 7: Precision-recall curves for various classifiers with 10% training data. Note that the performance is better with selected features (see text for details).

Figure 8: Finding ‘sister’ cities. Based on a query layout from central London, we show a sampling of very similar and dissimilar examples from other cities. Please refer to text for more details.

What is a good scale? We tested the proposed features at three scales, namely 0.25 km^2 , 1 km^2 , and 4 km^2 blocks, while normalizing by the tile areas. For most of the cities, we found the middle-scale, i.e., 1 km^2 to be the most effective. For the lowest scale, performance degraded rapidly (by 20-30%), while for the higher scale performance degradation was more gradual (approximately 10%).

What are the distinguishing features? Not all the proposed features are equally effective. Further, some of the features are mutually dependent, e.g., street density, connectivity index, and block size are mutually dependent. Hence, it is useful to identify the most discriminative features. We again use SVM-linear to learn the weights for the various descriptors. The top five features for city classification are: *4-way crossing proportion*, *connectivity index*, *# valence 4*, *#valence 3*, and *number blocks* — this partially answers the question raised in the paper title. It is interesting that metric reach does not feature among the top few descriptors.

Precision-recall and confusion matrix. Leaving aside the less distinctive features (we ignored metric reach, redundancy, node density, street density, total street length, histograms due to low weights) has a distinct impact on the precision-recall curves. As shown in Figure 7, the performance for all the classifiers improves with Multiclass SVM RBF performing best (only London example shown here).

Similarly, using our most important features outperforms classification using the image descriptors (GIST). This is due to the fact that our descriptors capture the essence of the topology and geometry of the street network while image descriptors do not. See Figure 9.

The observed confusion matrix, Figure 10, also reveals interesting patterns. At the selected feature level, the top confusion happens between New York↔Los Angeles; London↔Toronto; and London↔Moscow. The least confusing set was London↔Vienna, which is not surprising.

However, we found the selected descriptors could reliably distinguish between London↔New York, which was interesting; although London↔Cardiff were found to be confusing for the descriptors.

What features distinguish London from London? We now ask how to distinguish different areas of London, working at three levels: downtown (zone 1-4), mid-zone (zone 5-8), and outskirts (zone 8+). Our analysis reveals that now the distinctive features are: *nodes density*, *street density*, *valence 3*, *valence 4*, and *total street length*. These are exactly the features that get *downvoted* for intra-city classification. Qualitatively, these findings are easy to justify, although were not obvious before our analysis. For ranking other map

Figure 9: Precision-recall curves for linear SVM classifier with 10% training data. Note that the performance is better with selected features than using GIST features (see text for details).

Figure 10: Confusion matrix of 10 cities using 10% training data and best parameter RBF SVM.

tiles and for judging synthesized networks, these measures prove to be useful, as described next.

7.1. Applications

Downtown classification. We can use the learned features to query similar networks. Essentially, this is a nearest neighbor query in the corresponding features space (we use Fast Approximate Nearest Neighbour search [HAYSZ11] - FANN - for the queries). For example, using a sample patch (1 km²) from London downtown, we retrieve the *closest* and the *farthest* layouts across all *other* city patches (see Figure 8). Note that the match with Cardiff is an error as the network inside the park also is marked as a street in the input data.

Synthetic city generation. We can now use the extracted feature vectors to create London-like street layouts. As a first try, starting from random street networks (Cardiff, Toronto, Paris), we use a simulated annealing based approach to move the networks more to target London styles (outskirts, mid-zone, downtown, respectively). As ‘London-ness’ score, we use distance from the respective target London cluster in the feature space, while random movements allow addition and deletion of edges (between original nodes). If edges intersect, we introduce an intersection points (i.e., node in the graph). We then use the generated street network in CityEngine to create renderings. Even though the proposed features measure street network performance, it in turn introduces interesting visual appearance. Figure 11 shows three examples. We believe that combined with procedural parceling and facade modeling, this can lead to procedural generation of cities with prescribed behaviors both in performance and appearance.

8. Conclusion

We presented various topological and geometric features, both local and statistical, to characterize and classify street networks across various cities. We investigated the relative importance of the proposed features for different classification and clustering tasks based on data obtained from 10 cities, each with 100 different map tiles collected at three different scale levels. Beyond classification, we used the proposed features for various retrieval tasks, e.g., find cities with

Figure 11: Starting from random city tiles, we procedurally make the respective tiles more London-like by lowering feature distance (using features proposed in this work) to those of the targets. Inset shows the target London tiles (top-to-bottom: outskirts, mid-zone, downtown, respectively).

similar characteristics as that of ‘central London’. This, to the best of our knowledge, is the first attempt to capture functional behavior of street networks, in parallel to popular image or geometric descriptors.

The proposed features allow measurement of ‘functional validity’ of street networks. This has immediate relevance to city planning and procedural modeling (of street networks). By focusing both on connectivity and geometry, the features capture the *essence* of networks and also their data-driven validity. In other words, the work presents a way to rank different synthesized street networks based on their feature descriptions. In the future, beyond synthesis applications, we would like to investigate the effect of other factors like population density, street width, etc. on the behavior of cities. The relation to traffic simulation also poses interesting research questions from an inverse modeling perspective.

9. Acknowledgements

We thank the reviewers for their comments and suggestions for improving the paper; Jun Wang for creating renderings in CityEngine; Han Liu for proofreading the paper. This work was supported in part by the Marie Curie Career Integration Grant 303541, the ERC Starting Grant SmartGeometry

(StG-2013-335373), an Anita Borg Goggle PhD scholarship award, and gifts from Adobe Research.

References

- [ABVA08] ALIAGA D. G., BENEŠ B., VANEGAS C. A., ANDRYSCO N.: Interactive Reconfiguration of Urban Layouts. *IEEE Computer Graphics and Applications* 28, 3 (2008), 38–47. [2](#)
- [AIS77] ALEXANDER C., ISHIKAWA S., SILVERSTEIN M.: *A Pattern Language: Towns, Buildings, Construction*. Oxford University Press, New York, 1977. [2](#)
- [Ass06] ASSOCIATION A. P.: *Planning and Urban Design Standards*. Wiley, 2006. [2](#)
- [AVB08] ALIAGA D., VANEGAS C., BENES B.: Interactive Example-Based Urban Layout Synthesis. *ACM Trans. on Graph.* 27, 5 (2008). [2](#)
- [Boa10] BOARD T. R.: *Highway Capacity Manual*. Transportation Research Board, 2010. [3](#)
- [CEW*08] CHEN G., ESCH G., WONKA P., MÜLLER P., ZHANG E.: Interactive procedural street modeling. *ACM Trans. on Graph.* 27, 3 (2008), 103:1–9. [2](#)
- [CST00] CRISTIANINI N., SHAW-TAYLOR J.: *An Introduction to Support Vector Machines and Other Kernel-based Learning Methods*. Cambridge: Cambridge University Press, 2000. [6](#)
- [DSG*12] DOERSCH C., SINGH S., GUPTA A., SIVIC J., EFROS A. A.: What makes paris look like paris? *ACM Transactions on Graphics (SIGGRAPH)* 31, 4 (2012). [3](#)
- [Fis36] FISHER R. A.: The use of multiple measurements in taxonomic problems. *Annals of Eugenics* 7 (1936), 179–188. [6](#)
- [GPGB11] GALIN E., PEYTAIE A., GUÉRIN E., BENES B.: Authoring hierarchical road networks. *Computer Graphics Forum* 29, 7 (2011), 2021–2030. [2](#)
- [HAYSZ11] HAJEBI K., ABBASI-YADKORI Y., SHAHBAZI H., ZHANG H.: Fast approximate nearest-neighbor search with k-nearest neighbor graph. In *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence - Volume Two* (2011), IJCAI’11, AAAI Press, pp. 1312–1317. [8](#)
- [Hil96] HILLIER B.: *Space Is The Machine: A Configurational Theory Of Architecture*. Cambridge University Press, 1996. [3](#)
- [IGD14] IGNACIO GARCIA-DORADO DANIEL G. ALIAGA S. V. U.: Designing large-scale interactive traffic animations for urban modeling. *Computer Graphics Forum* (2014). [2](#)
- [JP08] J. PEPONIS S. BAFNA Z. Y.: The connectivity of streets: reach and directional distance. *Environment and Planning B: Planning and Design* (2008). [3](#)
- [LSWW11] LIPP M., SCHERZER D., WONKA P., WIMMER M.: Interactive modeling of city layouts using layers of procedural content. *Computer Graphics Forum (Proceedings EG 2011)* 30, 2 (Apr. 2011), 345–354. [2](#)
- [Lyn60] LYNCH K.: *The Image of the City*. MIT Press, 1960. [2](#)
- [Mar05] MARSHALL S.: *Streets & Pattern*. Spon press, New York, 2005. [2](#), [3](#)
- [OFCD02] OSADA R., FUNKHOUSER T., CHAZELLE B., DOBKIN D.: Shape distributions. *ACM Trans. Graph.* 21, 4 (Oct. 2002), 807–832. [5](#)
- [PM01] PARISH Y. I. H., MÜLLER P.: Procedural modeling of cities. In *Proceedings of SIGGRAPH 2001* (2001), pp. 301–308. [2](#)
- [SBJ03] SOUTHWORTH M., BEN-JOSEPH E.: *Streets and the Shaping of Towns and Cities*. Island Press, Wasington DC, 2003. [2](#)
- [SO93] SOUTHWORTH M., OWENS P. M.: The evolving metropolis: Studies of community, neighborhood, and street form at the urban edge. *Journal of the American Planning Association* 59, 3 (1993), 271–287. [2](#)
- [TMFR03] TORRALBA A., MURPHY K. P., FREEMAN W. T., RUBIN M.: Context-based vision system for place and object recognition. pp. 273–280. [6](#)
- [VABW09] VANEGAS C. A., ALIAGA D. G., BENEŠ B., WADDELL P. A.: Interactive design of urban spaces using geometrical and behavioral modeling. *ACM Trans. on Graph.* 28, 5 (2009). [2](#)
- [VGDA*12] VANEGAS C. A., GARCIA-DORADO I., ALIAGA D. G., BENES B., WADDELL P.: Inverse design of urban procedural models. *ACM Trans. on Graph.* (2012). [2](#)
- [WMWG09] WEBER B., MÜLLER P., WONKA P., GROSS M. H.: Interactive geometric simulation of 4d cities. *Computer Graphics Forum* 28, 2 (2009), 481–492. [2](#)
- [YWVW13] YANG Y.-L., WANG J., VOUGA E., WONKA P.: Urban pattern: Layout design by hierarchical domain splitting. *ACM Trans. on Graph.* 32, 6 (Nov. 2013), 181:1–181:12. [2](#)