142

Beyond Axial Lines

High-Resolution Geometric Analysis of London's Urban Fabric Using LIDAR Scans

Tasos Varoudis
Space Syntax Laboratory, Bartlett School of Architecture, UCL
t.varoudis@ucl.ac.uk

Sean Hanna
Space Syntax Laboratory, Bartlett School of Architecture, UCL
s.hanna@ucl.ac.uk

Pages: xxx-xxx

ABSTRACT

Space Syntax urban analyses have typically been performed on high-level approximations of the spatial network, in the form of axial lines (Penn et al. 1998), street segments (Hillier & Iida 2005), or similar, which necessarily lack some information at the smallest scales. Yet some key theoretical principles rely explicitly on locally observable features: intelligibility (Hillier et al. 1987), for example, is the correlation between long-range structure and that of the local. This paper investigates the highest resolution model of London's geometry yet available to determine the degree to which local surface geometries correlate with long-range structural features of the urban network, particularly global measures of network centrality: integration and choice.
By employing a LIDAR scan of ground and building surface geometry encompassing 2800 km^2 around London, captured as a dense cloud of points and transformed into 200-meter squares of 200 pixels, the research first queried the degree to which long-range structural features of the network are immediately evident in the local surface geometry data through clustering and unsupervised learning methods and trace its relationship to longer-scale centrality measures.
We conclude that traditional measures of network centrality can be learned to be predicted from local features, suggesting an alternative to traditional syntactical intelligibility. Finally, we determine the scale at which the differentiation between foreground and background networks is most clearly discernible through local features is approximately 4 km, which coincides with that seen to best predict movement and is within the range of a typical journey length.

KEYWORDS

space syntax, intelligibility, LIDAR, machine learning, neural networks.

1 INTRODUCTION

Space Syntax urban analyses have typically been performed on high-level approximations of the spatial network, in the form of axial lines (Hillier et al. 1987; Hillier 1996; Penn et al. 1998), street segments (Hillier & Iida 2005; Serra and Hillier 2019), or similar, which necessarily lack some information at the smallest scales. Yet key theoretical principles rely explicitly on locally observable features: intelligibility (Hillier et al. 1987; Penn 2003), for example, is the correlation between long-range structure and that of the local, which suggests that the cognitive processes guiding our movement and behaviour in the city depend on the local.

Real human experience goes far beyond the axial line, and the most relevant local features may be far richer and more varied than such representations. We investigate here the degree to which a more detailed description of the local visible geometry at any point in the city represents relevant information about its larger network structure, using the highest resolution model of London's surface geometry yet available. LIDAR scans capture both the ground and building surfaces and vegetation canopy, representing heights at a resolution of 1.0 m throughout the city; from these we use samples of e.g. 200 m x 200 m centred around a given street segment to indicate the surface geometry that may be perceived locally from that location.

The cognitive questions, reflecting the notion of intelligibility, have to do with the degree to which important network centrality measures such as integration and choice are evident in the local surface geometry. It is not a given that they should be—they are both non-local properties and measures of a model of linear street connections rather than surfaces. Yet they are fundamental to our understanding of human action as they have been shown to describe observed movement and socioeconomic behaviour (Hillier et al. 1987; Hillier and Hanson 1984; Hillier 1996; Penn et al. 1998; Hillier & Iida 2005; Serra and Hillier 2019). The theoretical grounding within space syntax is not yet resolved as to whether these cognitive mechanisms are innate and universal, or learned through exposure to space and culture; the theory of natural movement (Hillier et al 1993) seems to lean toward the former, and that of the inverted genotype (Hillier and Hanson 1984) to imply the latter. We can formulate both questions:

a) Are any long-range structural features of the network immediately evident in the set of local surface geometry data?
b) To what degree can an association be learned of global measures of network centrality (integration and choice) from local surface geometry?

There is also the related question of morphology. Distinct morphological features fundamental to urban structure are known to emerge from global measures of the network, such as the distinct foreground and background networks that comprise the dual structure of the grid (Hillier 2012). It is not yet known whether these are evident in local surface data.

c) Are the foreground and background networks distinguished from local features? And if so, what is the scale at which best differentiated?

These three questions, of cognition and morphology, structure the work that follows and are detailed in the following sections.

2 THEORY

Representation of the urban network in space syntax is almost always an approximation of street geometry by linear elements, originally axial lines (Hillier et al 1987, Hillier 1996), and later variants including continuity lines (Figueiredo and Amorim 2005) and natural roads (Jiang et al. 2008), segments of axial lines (Hillier and Iida 2005) and road centrelines. These lend themselves to calculations of long-range and global properties of the network, in which analyses are made of the graphs formed by the connection of these basic units. Typically these are centrality measures of integration (closeness centrality, or “to-movement”) and choice (betweenness centrality, or “through-movement”), which have been observed to correspond with movement and other behaviour (Hillier et al 1987, Hillier 1996, Hillier and Iida 2005…).

A separate class of analytical methods have focused on local calculations, notably the *exosomatic visual architecture* agent model (EVA, Turner and Penn 2002) implemented in depthmapX (Varoudis 2012), which simulates a random-walk in which agents move probabilistically based solely on distances immediately visible from their current location. This uses a visibility graph based on a regular grid as its spatial representation, typically for smaller spaces within buildings, although related methods have been implemented on street segments at urban and regional scale (Hanna 2021). Like the global and long-range measures above, these are deemed useful because they correspond with observed movement.

2.1 Intelligibility as a relation between properties

The relationship between the local and long-range structure of the city is understood through a core space syntax concept of *intelligibility*. Qualitatively, this is the degree to which “the whole can be read from the parts” (Hillier et al 1987), and quantitatively, as a measure of a whole network, intelligibility was originally defined as the correlation between the local connectivity (degree) of axial lines and their global or longer-range integration. High correlations mean that the centrality of a line can be inferred from the number of its local connections, so the network is deemed intelligible.

While the theory is clear, this specific measure has some limitations both at the global and local ends of the scale. In practice, global integration is rarely used for the long-range measure, in favour of more local radius integration of two or three steps (Hillier 1996; Penn 2003). At the local end of the scale, an axial line, particularly when long, may be more usefully considered as an approximation to a series of local units, but is not itself a single local unit in the sense that is genuinely perceived from a single point (Serra and Hillier 2019; Hanna 2021). Versions of intelligibility measures have been proposed using properties of the more local street segment, using segment angle either directly (Hanna 2022), or in random walk methods (Hanna 2021, 2024).

An alternative approach to representing local spaces is as vectors defined by their graph spectra (Luo et al 2003). First used to classify local spaces (Hanna 2007) and whole cities (Hanna 2009), the spectra of specific local neighbourhoods (e.g. the nearest 100 segments from a location) have since been shown to reveal larger network properties within the city (Hanna 2012) and larger regions (Varoudis and Penn 2020), although not specifically with the aim of predicting the main centrality measures of integration and choice. A related technique has used machine learning on images of local networks extracted from segment networks (Varoudis and Penn 2019). All of these, like traditional centrality and intelligibility, are based on graph representations of lines (axial or centrelines).

If not constrained by this basic representation, and we are concerned with plausible cognitive and perceptual mechanisms, the question is relevant: at local scale, what can we actually see and experience? Often we are unable to perceive less than the complete axial line. Yet we can perceive many other features not represented at all, such as building heights, street width, building geometry, the presence of vegetation and the material surfaces of the built environment. The detailed features of building geometry vary by neighbourhood (Laskari et al. 2008) and street-view images are capable of predicting property prices (Law et al. 2019). This paper goes beyond the axial line (and segment, and centreline) to capture some of these additional perceptual features via a local sample of actual surface geometry.

2.2 Is urban structure evident or learned?

The concept of intelligibility implies two theoretical approaches to the structure and cognition of space which are quite distinct. In the first, spatial structure is imagined to be immediately perceptible from local observation, simply due to the correlation between local and global; in the second, there is a cognitive process by which it must be learned through experience in navigating the city. These two are sometimes conflated even in the same description, as in Hillier (1996, p. 94), which uses both the immediately perceived, “intelligibility […] means the degree to which what we can see […] is a good guide to what we cannot see”, and the learned, “a picture of the whole urban system can be built up from its parts […] from moving around from one part to another” on the same page. Both are plausible, both are important, and one or the other is implied in the basic measures of space above, and in technical approaches to data.

To the extent spatial patterns are immediately perceived, we expect them to be evident within the spatial data itself. The cognitive equivalent is that one perceives them with no training, as is the case with EVA agents. This is not to say that the movement of such agents should necessarily correspond to these patterns (they do have their own specific method) but only that there is not a process by which they tune the rules of their behaviour to detect a particular desired outcome, e.g. integration, optimal paths. Here we are interested in the patterns that emerge most obviously from the data, in this case local surface images. The corresponding approaches to the data are unsupervised pattern recognition, using methods like clustering and principal component analysis (PCA).

Alternatively, the relevant patterns could be learned. Both integration and choice measures depend, for their calculation, on determining the shortest path between any two given nodes, a property that requires complete knowledge of the network, which must be learned if actually used as a cognitive strategy. More plausibly, Penn (2003) suggests that human cognition might serve as a ‘correlation detector’—i.e. we learn to recognise important features in intelligible space. Even if some of our spatial cognition is innate and invariable across the human species, evidence suggests that the environment in which one is raised influences navigation ability or strategy (Coutrot et al 2022). Each of these potential strategies implies a hypothesis as to whether and to what degree patterns relating spatial features can be learned by repeated exposure to urban environments. The corresponding approaches to the data are those of supervised learning, more specifically regression to a continuous variable, in which a learning algorithm is trained to associate perceptible inputs (i.e. local surface geometry) with relevant structural outputs (i.e. long-range and global centrality measures).

We are interested here in both the patterns that emerge from the data, and those that can be found by learning through experience. The two approaches inform two distinct parts of the method (sec 3.1 and 3.2) and results (sec 4.1 and 4.2).

2.3 Morphology: foreground networks

In addition to the human aspects of the city, the cognitive, perceptual and behavioural, we are interested in the morphology of the city itself. The dual structure of the urban grid, by which it is formed of a distinct foreground and background network, is tied to movement, social and economic behaviour, and universal to cities (Hillier 2012). The dominant foreground streets, longer and more continuous, are associated with greater through-movement and economic activity, and form the larger structure around the background of shorter, more perpendicular streets of low traffic. These are distinguished by the measure of angular choice, with higher values indicating the foreground.

Higher radii are typically used, from the scale of a reasonable journey (2–5 km) to that of the whole city, but no specific scale is preferred. It is reasonable to assume that it varies with the scale of any given urban area considered or journey to be made, and Hillier (2012, p. 33) explicitly suggests that the foreground network can be defined “at all scales, from a couple of shops and a café through to whole sub-cities”. Yet some specific distances may be relevant to these properties of city form, as observed movement appears to correlate best at particular choice radii around 4km (Serra and Hillier 2019; Hanna 2021; Hanna 2024) and patterns similar to the foreground network also emerge in random walks (Hanna 2021) only after a similar distance, and remain stable thereafter. This distance is within the range of typical journey lengths, which average from 1.3 km for pedestrians to 11 km for cars (Department for Transport Statistics, 2019). Is there a particular scale at which the foreground appears?

This question may not be one of network measures, or of movement, but we might ask it from the standpoint of perception of local features. The difference between a residential street (typical of the background) and a public/commercial high street (typical of the foreground) is easy to visualise from immediately visible non-network features such as building height (etc.). If these are really perceptible features, and captured in the local surface data, can we determine the scale at which the foreground network is most clearly distinguished from the background? This is a classification problem, by which we can measure the scale at which classification is most successful. The corresponding approach to the data are those of classification based supervised learning, and is detailed in part 3.3 of the methodologies and 4.3 of the results.

3 DATASETS AND METHODS

3.1 The dataset

Our dataset comprises LIDAR surface scans of approximately 2800 km^2 around London, captured as a dense single-channel data cloud of the terrain, built environment and vegetation with 1 sqm per pixel resolution. We made use of both the Digital Terrain Model of the bare earth surface (DTM) and the Digital Surface Model (DSM) which shows the built environment and vegetation canopy.

As the target of our work is to investigate synergies between surface geometrical features and London's larger network structures, we converted the initial dataset into roughly 150,000 cropped images of 200 m square. These are represented as 200 x 200 pixels with a single channel of height information, centred around each street segment of London. Each '200x200 crop' records the details of the surrounding urban geometric patterns centred around the mid-point of each segment and rotated so the direction of the street segment is horizontal (figure 1). The crops are height-normalised based on the mid-point of the street so as to record the relative difference in meters from the centre. Throughout the different phases of our investigations, including some that didn't make it to this paper, the 150,000 topographic crops describing London's urban canyons and the surrounding areas ended up being between 80 and 380 GigaByte of raw data.

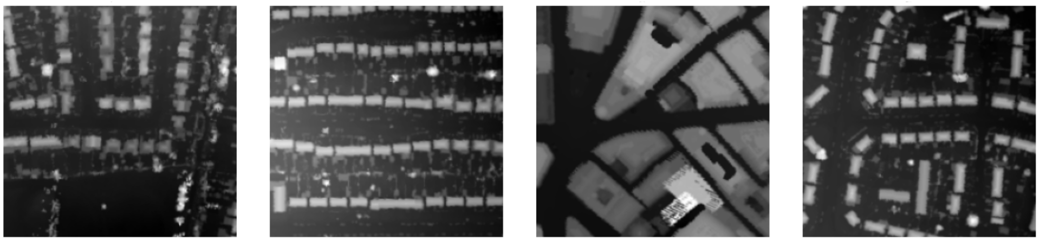

Figure 1: 200m, 200x200 pixel crops

3.2 Evident Urban Structures - Unsupervised Learning

To grasp the most apparent and easily generalizable patterns emerging from the data, we employ several unsupervised techniques. We first built a reduced-size dataset from the '200x200' pixel crops (40,000 dimensions per sample) down to '32x32' pixels (1024 dimensions per sample).

Principal components analysis (PCA) is among the most common and simplest dimensionality reduction techniques, producing the linear dimensions of greatest variance within the set. We performed a Principal Component Analysis on the data, initially reducing the datasets to 64 principal components with a retained summed explained variance of 0.776. Explained variance in PCA refers to the amount of the total variability in the data that is accounted for by each principal component, which quantifies how much of the original data's information (variance) is captured by these 64 principal components: higher values indicate a greater proportion of the data's total variance, or information, is explained by the respective component. Although some information is lost in the dimensionality reduction process, the principal component and the reconstructed images retain the most significant aspects of the original data.

Components of an image dataset can be visualised as "Eigenfaces", via a method originating in face recognition (Turk and Pentland 1991). This technique represents the eigenvectors corresponding with each component of the (PCA) analysis as an image, in which the intensity of each pixel is the variance in that dimension; the first "eigenface" of a data set of faces would be likely to visualise features of greatest variance e.g. the oval area of overall shape and size of the head. Each 32x32 crop, single-channel, sample in our set is treated as an image of the local urban topography, similar to faces, and thus we leverage the eigenface approach to extract essential features and patterns inherent in the three-dimensional urban landscape. We are decomposing the large set of LIDAR images into a smaller set of characteristic feature images, known as eigenfaces in the context of facial recognition, but in our case, could be termed as '*eigentopographies*'.

These *eigentopographies* represent the dominant structural features immediately evident from the very local surface geometry data encoded in the principal components. We can extract these integral spatial elements by utilising the topmost features. Adapting the eigenface methodology in this novel way not only establishes a robust framework for the analysis of high-resolution LIDAR data but also unveils new perspectives for comprehending urban landscapes through data-centric, eigen-based approaches. For this process, we represented each image as a 1028 (i.e. 32x32) dimensional vector, which ignores local spatial relationships within the image, applied PCA and then extracted the *eigentopographies* by reshaping the components back to 32x32 pixel images.

We used K-Means clustering to group these topographical features and explore their local and global relations. The K-Means algorithm is a popular unsupervised machine learning technique used for partitioning a dataset into distinct clusters based on feature similarity and the minimisation of within-cluster variances. For our analysis, we use values of K (number of clusters) between 2 and 6.

3.3 Supervised Learning: Convolutional Neural Networks and Residual Networks

Neural Networks (NNs) in supervised learning are powerful computational models that mimic the human brain's structure and function to process and learn from labelled data. Exploring the degree to which spatial patterns can be learned from our data, we compared several supervised learning models and found convolutional neural networks called Residual Networks (Resnet) to be effective both as regressors and classifiers. We trained the model only to a point prior to overfitting, as is good practice, and thus not memorising the training examples rather than learning the general patterns that would allow it to perform well on unseen data.

We trained these networks using our 32x32 cropped single-channel dataset as input features and the corresponding space syntax centralities as a target output, allowing the NNs to learn to map inputs to desired outputs through a process of error minimization and weight adjustment.

Convolutional neural networks (LeCun, 1989), or CNNs, are a tailor-made architecture of neural networks for processing data that have a grid-like topology similar to the data that we are working with. They are well suited to image processing because they maintain and operate on local spatial relationships within the image. Most of the successes in machine learning and artificial intelligence are heavily dependent on CNNs, like self-driving cars (Bojarski et. al. 2016). We tested the following model architectures: ResNet-8, MiniVGG and SqueezeNet in a regression task mapping our 32x32 cropped dataset to integration (section 4.2) and found ResNet-8 to have the lowest overall error; the results reported in this paper are those of ResNet-8.

The ResNet model, short for Residual Network, is a highly influential deep learning architecture that scores among the top current models in ImageNet benchmarks even though it was first introduced in 2015 (Kaiming He et. al.). ResNet addressed the problem of vanishing gradients in deep neural networks by introducing the concept of skip connections, or shortcut connections, that allow the gradient to bypass one or more layers. These connections enable the network to learn residual functions with reference to the layer inputs. While ResNet models usually have variants with depths of 50 or more layers in our research we work with a minimum ResNet implementation with 8 total layers (ResNet-8) and two residual blocks. The core idea of utilizing residual blocks has been widely adopted in various neural network architectures, making ResNet a foundational model in the development of deep learning applications, not only in image classification but also in a broad range of other computer vision tasks which is also validated in our unique dataset with the highly successful learning convergence.

4 RESULTS

This study employed a high-resolution LIDAR dataset to unravel the intricate relationship between local surface geometry and the broader network structure of the city. This section presents the findings focusing on the easily emerging structures in the data through unsupervised learning methods and the extent to which larger network structures can be learned using supervised learning techniques. Finally, the impact of physical topography on urban intelligibility and the identification of the foreground network at varying scales are also discussed.

4.1 Network structure immediately evident

Our analysis begins with the application of PCA to our dataset, to reveal any aspects of London's network structure that are immediately evident from the local surface geometry. The first principal component (PC1) shows decent correlations with angular segment Integration indicating that the most significant variance in our dataset aligns well with the concept of space syntax integration. This correlation substantiates our hypothesis that the structural features of the network are, to a degree, immediately evident within the local surface geometry. However, the relationship with angular segment Choice is less pronounced, suggesting that larger-scale through-movement patterns are not as readily captured by PC1.

We determine the scale of the long-range measure that is most evident in the local data by comparing the correlation across a range of radii. Figure 2 depicts the R-squared (r2) correlation between PC1 and both Integration and Choice across various radii and it highlights the differential sensitivity of our data to these two key space syntax measures. Integration peaks at a radius between 2km and 4km (with 10km still high) and, while significantly lower, choice peaks at a very local 1km to 2km.

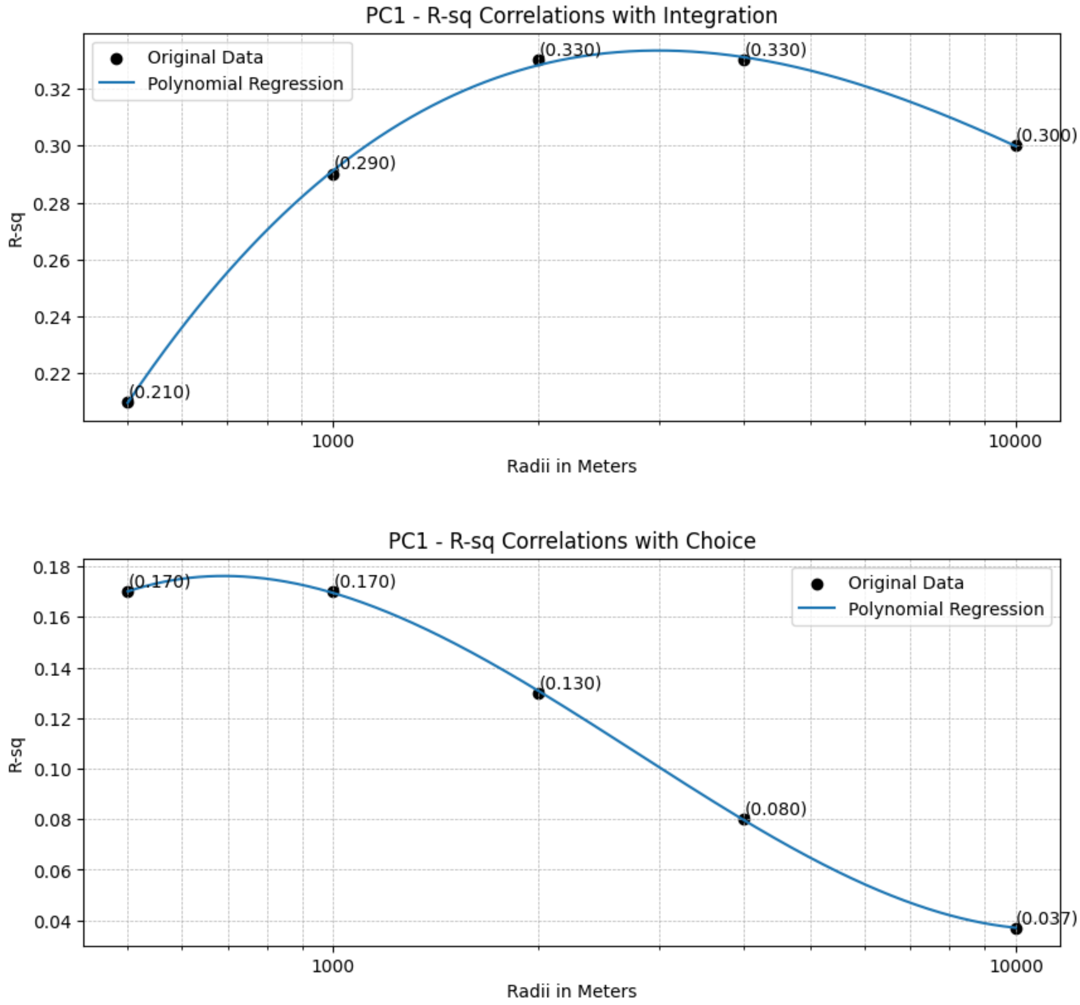

Figure 2: R-squared correlations for Integration and Choice with PC1 at multiple radii.

Eigentopographies (figure 3) provide a visual description of the features that vary most between samples, and thereby clearly generalise the very local dominant spatial characteristics encoded in the data. By visualising the top 4 eigentopographies ranked by magnitude of their variance [eigenvalues: 0.229, 0.119, 0.089, 0.021], it is evident that these are the linear elements, sloping topographies and some distinct street configurations. The first component (eigentopography 1) consists primarily of a strong linear band aligned with the street segment, indicating the strongest local features consist of the difference in height between the street level and its surrounding buildings, with potentially some variance in the street width. The similar band of eigentopography 4 also highlights the region immediately bordering the street segment, giving information on the depth of the buildings fronting the street and the land behind. Components 2 and 3 display the relative slope of the land overall, both perpendicular and parallel to the street.

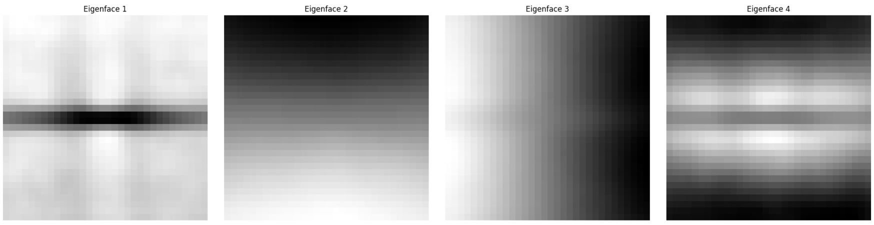

Figure 3: Top 4 Eigentopographies

Regarding the long-range and global properties of the network, we observed a significant similarity between the K-Means clusters and global integration (figure 4, k=4 clusters). This suggests that the very local topography encapsulated in the 200m x 200m crops has a profound impact on the spatial configuration and intelligibility of urban areas. This was a strong indication for us that features of built geometry more directly visible than axial connectivity can be shown to play an important role in our understanding of the city. As a result, we continued with more rigorous supervised methods that also work better with 2D arraignments of data and preserve spatial structures.

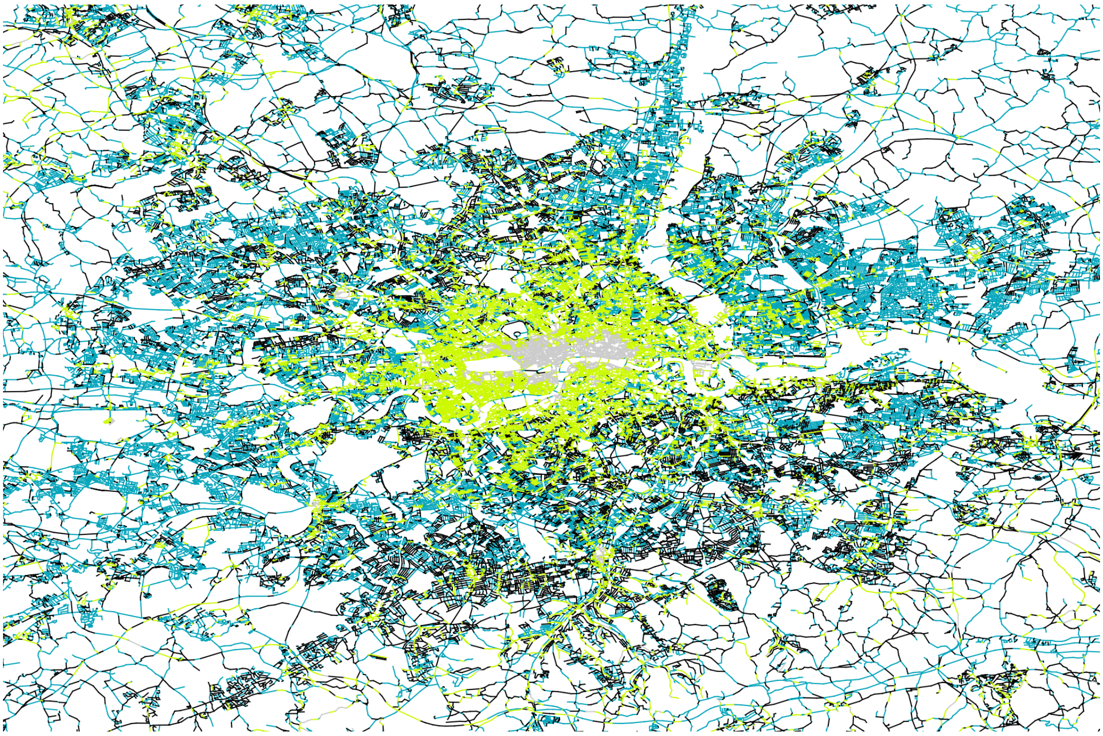

Figure 4: K-Means – K=4 clusters.

4.2 Network structure learned

Supervised learning models were used to determine the degree to which complex spatial relationships, especially long-range and global centrality measures, can be learned and predicted from local surface geometries. We trained our regressor ResNet-8 on Integration and Choice choosing five radii: 500m, 1km, 2km, 4km and 10km. Figures 5 and 6 present the R-squared values for Integration and Choice and an indicative quadratic interpolation.

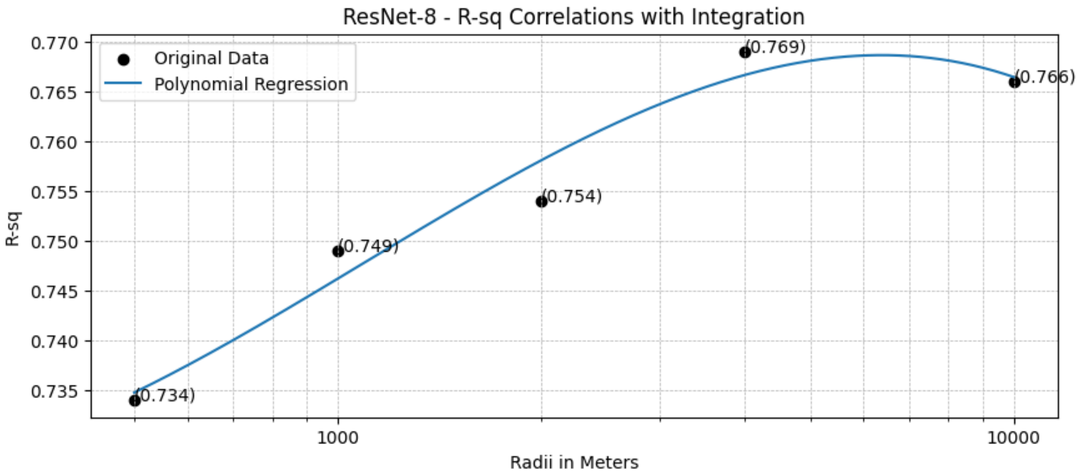

Figure 5: Resnet-8 prediction R-sq at multiple radii – Integration

Integration has (Fig. 5) a very strong correlation with learned features from the local 200x200 meter geometry. This was a strong indication for us that features of built geometry more directly visible than axial connectivity can be shown to play an important role in our understanding of the city and its global integration structures. The peak correlation is observed to be around the 4km radius, indicating that it is at this scale that the long-range structure of the city can best be learned as an association with local topographical features. This scale is comparable to the unlearned patterns of section 4.1, and that which is known to predict movement (section 2.3).

Figure 6 depicts R-squared values of Choice for different radii; correlations are lower than with integration, but still significant. The peak, 2km, is at a significantly lower radius compared to Integration.

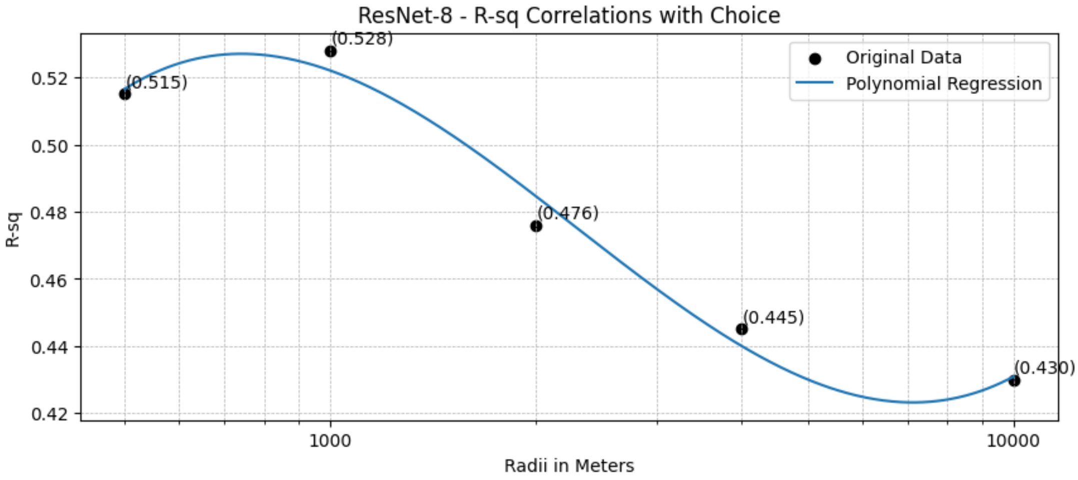

Figure 6: Resnet-8 prediction R-sq at multiple radii - Choice

Choice is associated with the foreground network, and through-movement with longer, straighter streets that are often morphologically distinct from even their immediate neighbours (e.g. a residential street of terraced houses often directly joins a high street). This suggests a hypothesis about the detailed local features that may be relevant to choice: that these have less to do with local neighbourhood surroundings perpendicular to the street and more to do with those along its length. To test this we generated a complimentary dataset deriving for 400x100 meter crops along the axis of the segment line and resized down to 64x16 pixels. The two datasets, 32x32 used till now and 64x16, have an equal number of features (1024). The purpose was to explore how this spatial information extended parallel with the axiality of the street segment affects local and long-range correlations. Figure 7 below depicts the results, R-squared correlation, from training a ResNet-8 with an input size of 64x16 and Choice as a target value. Predicting Choice from learned features is significantly better with elongated samples with correlations rising to r2=0.524. The radii at which this peaks is also greater, estimated between 2km and 4km. This suggests, for choice, that information along the street is more important than information perpendicular to it. This is the information we can see in the distance where long straight streets or minimal angular changes are more easily evident in the data.

The spatial information aligned with the axiality of street segments, emphasizing longitudinal rather than perpendicular visibility, suggests this is more pertinent to human perceptual experiences and navigation patterns within urban canyons.

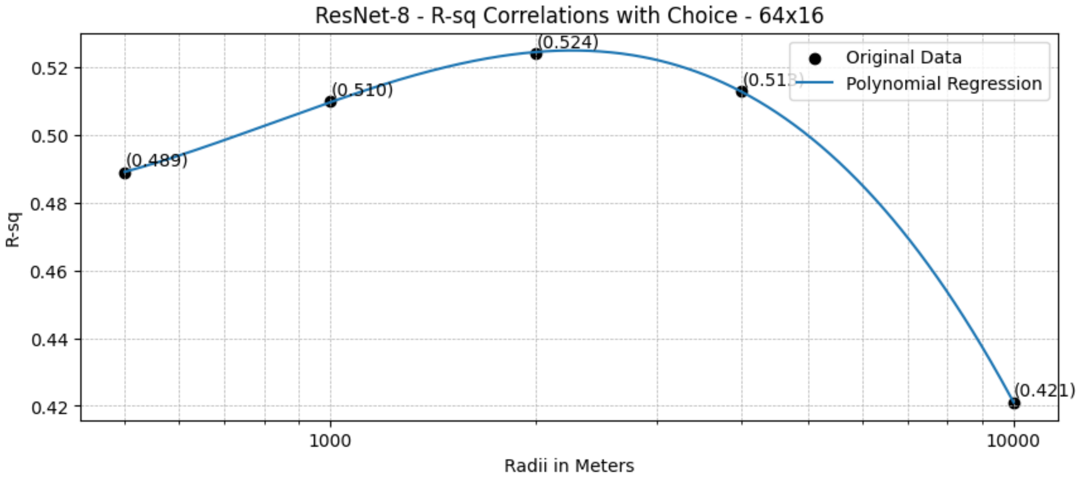

Figure 7: Resnet-8 prediction R-sq at multiple radii – 64x16 input - Choice

4.3 The foreground network and its scale

To investigate the association of local data with fundamental network morphology, a modified ResNet-8 model was trained to predict two classes: the foreground and the background network.

We first worked to identify the percentage of the cut-off split between foreground and background networks that is better distinguished in our dataset. We tested 10%, 15% and 20% splits at a 2km radius and '10%' performed better overall at classification accuracy. In defining the top 10% of segments as the foreground network, our classification models aimed to distinguish between the more significant, through-movement-oriented streets and the rest of the urban fabric. In our approach, we ensured that the model was trained and validated on datasets with evenly distributed representations of both foreground and background networks. This balanced distribution was crucial to prevent any bias towards one category over the other, therefore avoiding any form of artificial favouritism in the model's learning process. The classification accuracy on multiple radii is depicted in Figure 8 below.

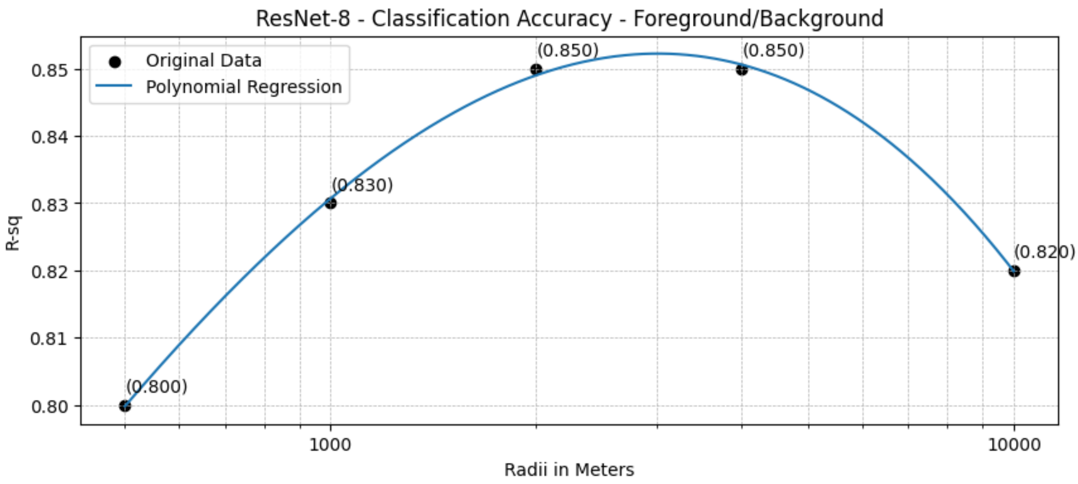

Figure 8: Resnet-8 Classification Accuracy – Foreground/Background Network.

From the plot is evident that distinct morphological features fundamental to large-scale urban structures are present in our local geometric data. Separating the distinct foreground and background networks can be done with high accuracy from local features and it peaks close to radius 4km, a scale consistent from prior studies of movement prediction (section 2.3) and with the results of regression in the results above. Even though the combined model accuracy in both 2km and 4km is the same, if we look at the per-cluster precision, 4km gives as equal values for both the foreground and the background network. At a radius of 2km the model performs slightly better in distinguishing the background network than the foreground. The optimal scale of classification around 4km, and the optimal threshold of 10% of streets as the foreground network, provide numerically precise values which define the empirical propositions of Hillier (2012) of the foreground and background networks; they demonstrate such structures may have a cognitive basis through the lens of high-resolution spatial data.

5 CONCLUSIONS

Our research is framed around three critical inquiries, using novel applications of high-resolution LIDAR scans, and seek to establish the extent to which local surface geometries correlate with long-range structural features of the urban network, particularly global measures of network centrality: integration and choice. We have examined both the *cognitive* question, of intelligibility, and also the *morphological*, specifically the scale at which the differentiation between foreground and background networks, is most clearly discernible through local features.

We first queried the degree to which long-range structural features of the network are immediately evident in the local surface geometry data. A moderate correlation (r^2=0.33) was found between high-radius integration and the first principal component (PC1) derived from our data. This relationship underscores the potential of local surface geometry to reflect broader structural characteristics of the urban fabric, i.e. analogous to syntactic intelligibility (Hillier et al. 1987; Hillier 1996; Penn 2003), although with a limited ability to capture the dynamics of choice. The first four components (eigentopography 1 to 4) suggest that integration is associated not only with properties of the street itself but also with features perpendicular to it, including the height difference between street and buildings, and perhaps surrounding tree cover.

To the question of whether such global measures of network centrality can be learned to be predicted from local features, a marked improvement was observed in supervised learning models. These models showed very strong performance in predicting integration ($r^2 \geq 0.769$), and choice ($r^2 \geq 0.524$), indicating the capacity of urban geometries to inform us about the broader spatial structures, which was seen to peak around 2-4 km. Choice prediction was improved when employing elongated, street-aligned geometrical inputs, as opposed to square inputs with more information perpendicular to the street. These findings suggest both that local topographical features are associated with the perception of intelligibility within the urban fabric, and also highlight the directional and relational aspects of urban form as crucial components in people's spatial cognition and behaviour in space.

Our morphological question was the relationship between local features and the known "dual structure" of the grid (Hillier 2012). We identified a specific scale, around 4km, at which the differentiation between foreground and background networks becomes most pronounced, and can be learned from local surface features. This scale is consistent with distances at which movement patterns are optimally predicted in the literature (Serra and Hillier 2019; Hanna 2021, 2024), which indicates that, while the foreground network may be considered "at all scales" (Hillier 2012) it is particularly evident at one particular scale. We suggest that this affirms the existence of the foreground network not only as an artefact of network centrality values, but also as a feature that is accessible to individual cognition.

This research underscores the profound potential of high-resolution LIDAR data in advancing our understanding of urban spaces beyond traditional approximations of spatial networks. By demonstrating the capacity of local surface geometry to inform on global network structures and urban morphology, our work contributes to a deeper comprehension of urban intelligibility and spatial behaviour. From these three results we conclude that there are indeed patterns perceptible in the local surface geometry that can aid larger understanding of the network centrality, which may be used in navigation, perception of space and other spatial behaviour. These suggest that the city is intelligible with respect to its local geometry, and that the large-scale morphological features of foreground and background networks are intrinsic to this intelligibility.

ACKNOWLEDGEMENTS

Petros Koutsolampros assisted with the initial data pre-processing scripts. The research was supported by the UKRI Research Capital Investment Fund (RCIF) through 'Large-scale computational modelling' and 'Machine Intelligence Lab' bid (2022).

REFERENCES

Coutrot A, Manley E, Goodroe S, Gahnstrom C, Filomena G, Yesiltepe D, Dalton RC, Wiener JM, Hölscher C, Hornberger M and Spiers HJ (2022) Entropy of city street networks linked to future spatial navigation ability. *Nature* **604**, 104–110. https://doi.org/10.1038/s41586-022-04486-7

Bojarski M., Del Testa D., Dworakowski D., Firner B., Flepp B., Goyal P., Zhang X. (2016). End to end learning for self-driving cars. arXiv preprint arXiv:1604.07316.

Department for Transport Statistics (2019) *National Travel Survey*, Table NTS0601. https://www.gov.uk/government/collections/national-travel-survey-statistics

Figueiredo, L. and Amorim, L. (2005), Continuity lines in the axial system, in A Van Nes (ed), *5th International Space Syntax Symposium*, TU Delft, Faculty of Architecture, Section of Urban Renewal and Management, Delft, pp. 161-174.pp. 163.

Hanna S (2007) Representation and generation of plans using graph spectra. In: Kubat, A.S. and Ertekin, O. and Guney, Y.I. and Eyuboglu, E., (eds.) 6th International Space Syntax Symposium, 12-15 Jun 2007. (pp. 099.01-091.14). Istanbul Technical University: Istanbul, Turkey.

Hanna S (2009) Spectral comparison of large urban graphs. In: Koch, D. and Marcus, L. and Steen, J., (eds.) *Proceedings of the 7th International Space Syntax Symposium.* (pp. p. 39). Royal Institute of Technology (KTH): Stockholm, Sweden.

Hanna S (2012) Comparative analysis of neighbourhoods using local graph spectra, in Greene M, Reyes J, Castro A (eds.) *Proceedings of the Eighth International Space Syntax Symposium* (pp. 8034:1–8034:18) Santiago de Chile. ISBN: 978-956-345-862-6

Hanna S (2021) Random walks in urban graphs: A minimal model of movement, *Environment and Planning B: Urban Analytics and City Science*, 48(6), pp. 1697–1711. doi: 10.1177/2399808320946766.

Hanna S (2022) Route Choice from Local Information: Comparing Theories of Movement and Intelligibility. In: Van Nes A and de Koning RE (eds.) *Proceedings of the 13th Space Syntax Symposium*, Bergen, Norway, pp. 377:1–22.

Hanna S (2024) The natural scale of urban networks: Intelligibility is not scale-free. In *Proceedings of the 14th Space Syntax Symposium*, Nicosia, Cyprus.

He K, Zhang X, Ren S, Sun J (2016) Deep Residual Learning for Image Recognition, Computer Vision and Pattern Recognition (CVPR).

Hillier B (1996) *Space is the Machine: A Configurational Theory of Architecture*. Cambridge University Press.

Hillier B (2012) The city as a socio-technical system: a spatial reformulation in the light of the levels problem and the parallel problem. In: Arisona, SM and Wonka, P and Aschwanden, G and Halatsch, J, (eds.) *Digital Urban Modeling and Simulation*. Springer

Hillier B, Burdett R, Peponis J and Penn A (1987) Creating life, or, does architecture determine anything? *Architecture and Behavior/Architecture et Comportement*, 3, 233-250.

Hillier B and Hanson J (1984) *The Social Logic of Space*. Cambridge University Press.

Hillier B and Iida S (2005) Network and Psychological Effects in Urban Movement. In A. Cohn & D. Mark (Eds.), *Spatial Information Theory* (Vol. 3693, pp. 475–490). Springer Berlin / Heidelberg

Hillier B, Penn A, Hanson J, Grajewski T, and Wu J (1993) Natural Movement: or, Configuration and Attraction in Urban Pedestian Movement. *Environment and Planning B*, vol (20), 29-66.

Jiang B, Zhao S and Yin J (2008) Self-organized Natural Roads for Predicting Traffic Flow: A Sensitivity Study, *Journal of Statistical Mechanics: Theory and Experiment*, July 2008, P07008.

Laskari A, Hanna S and Derix C (2008) Urban identity through quantifiable spatial attributes: Coherence and dispersion of local identity through the automated comparative analysis of building block plans. In: *Design Computing and Cognition '08* (eds. J.S. Gero and A. Goel), Springer. https://doi.org/10.1007/978-1-4020-8728-8_32.

Law S, Paige B and Russell C (2019) Take a look around: using street view and satellite images to estimate house prices, *ACM Transactions on Intelligent Systems and Technology*, 10(5) pp. 1–19.

LeCun, Y. (1989). Generalization and network design strategies. Technical Report CRG-TR-89-4, University of Toronto.

Luo, B., Wilson, R.C., Hancock, E.R., 2003, “Spectral Embedding of Graphs”, Pattern Recognition, vol. 36, pp. 2213-2233.

Penn A (2003) Space syntax and spatial cognition or why the axial line? *Environment and Behaviour*, 35(1), 30–65.

Penn, A., Hillier, B., Banister, D., & Xu, J. (1998). Configurational modelling of urban movement networks. *Environment and Planning B: Planning and Design*, *25*(1), 59–84.

Serra M and Hillier B (2019) Angular and metric distance in road network analysis: A nationwide correlation study. Computers, Environment and Urban Systems 74: 194–207.

Sirovich, L., & Kirby M., (1987). Low-dimensional procedure for the characterization of human faces. *Journal of the Optical Society of America A.* 4 (3): 519–524. doi:10.1364/JOSAA.4.000519

Turk M and Pentland A (1991) Eigenfaces for Recognition, *Journal of Cognitive Neuroscience*, 3(1) pp. 71-86.

Turner A and Penn A (2002) Encoding natural movement as an agent-based system: An investigation into human pedestrian behaviour in the built environment. Environment and Planning B: Planning and Design 29(4), pp. 473–490.

Varoudis, T. (2012). depthmapX multi-platform spatial network analysis software. OpenSource, http://varoudis.github.io/depthmapX.

Varoudis T and Penn A (2019) Variational Beauty of Space: Machine Intuition and Non-Linear Neural Aggregations. In: Duan, Jin, (ed.) *Proceedings of the 12th International Space Syntax Symposium (12SSS)*. Beijing,China.

Varoudis T and Penn A (2020) Spectral Clustering and Integration: The inner dynamics of computational geometry and spatial morphology, in *Formal Methods in Architecture: Advances in Science, Technology & Innovation,* Sara Eloy, David Leite Viana, Franklim Morais, Jorge Vieira Vaz (Editors), Springer.