Chapter 1: Urban Complexity
Sean Hanna, Bartlett School of Architecture, UCL

Cities are arguably the most complex things we have ever built. Most of humanity now lives peacefully and productively in groups vastly larger than our natural social capacity would allow, and as this proportion only continues to grow (United Nations 2019), the problems it poses for those responsible for designing and managing the city are equally great. In part this is due to the increasing complexity of the city itself. We humans are equipped with intuitions about spaces we can see in a glance and easily traverse on foot, and about social groups of less than 200 people (Dunbar 2014), but cities are well beyond our human scale. Historically, we have often managed urban growth gradually, building by building and street by street, but we can now construct environments for millions of people almost overnight, without feedback. In part this is due to unprecedented technological change. We do not have sufficient precedents for self-driving cars, smart homes, or the setting of our social and commercial interactions away from the streets and into virtual space, yet we need to understand the nature of city so that we don't plan counter to it.

What we do have is access to data, more than ever before, although there are different approaches to how we can use it. Traditional scientific research, on which much of our understanding of cities rests, begins with theory or hypotheses. The field of Space Syntax, for example, is a scientific discipline through which accurate predictions can be made about the effect of spatial configuration on traffic density, crime, social interaction, property values and other complex phenomena (Penn 2003; Hillier 2007; Silva 2017). It consists of a set of theories about the relationship between space and society, including the influence of space on the natural movement of people. From these, along with associated representations of space and related analytical methods, specific hypotheses or predictions are made, which can then be tested against observations. Data on real human movement and behaviour enters this scientific process last, to test the hypothesis, which is sometimes corroborated, and sometimes refuted, thereby advancing the core theory. Importantly, theories themselves are valued for their clarity, and in some cases the stated theory necessarily simplifies what patterns are in the data for the sake of this clarity and understanding. Sometimes, complexities in the pattern are missed.

Machine learning approaches instead begin with the data, and attempt to discern patterns from within it, with the intent that these patterns will reliably inform decisions we make about the city. The complexity of these patterns determines how they do so: where they are clear enough for us to articulate they may guide general theory and policy, and where they are not they may be used to make highly contextual predictions. The essential problem of machine learning is that of picking out these patterns. As Jacobs (1961) articulated in the early days of complexity science, there is a profound difference between the "disorganised" complexity if millions of independent actions, which might be treated statistically, and the "organised" complexity of systems (people, spaces, data, goods) for which their mutual interactions matter. Cities are the latter.

Phenomena that matter most in a city do not lend themselves to treatment by gross statistical analysis, precisely because of these interactions between many parts. While statistics isolate one or two variables, a particular demographic of the population, for example, machine learning can be valuable just because it may find the patterns among many variables, perhaps including that population's distribution in space, movement in time, connection via technology, and so on. But each new variable potentially interacts with all the others, potentially increasing the complexity of the system to be studied exponentially. The essential problem in such an analysis is in understanding whether there exist any points, scales or levels of representation, at which a meaningful pattern can be extracted. If machine learning begins with the data, is this even possible with something so complex as a city? How is the city intelligible to the machine?

How can a machine understand the city?

For some phenomena, patterns are predictable because they converge with increasing scale or time. Agent based modelling often owes its effectiveness to the fact that a large population is used. Turner and Penn's (2002) exosomatic visual architecture agents (EVAS), for example, are extremely simplified models of pedestrians in space, which make random navigation decisions based on a probability weighted by how far they can see in any given direction. Their individual paths are entirely unrealistic, appearing often to walk in circles, but over an extended time, a population of agents will converge very closely toward the distribution of real people in a space, correlating approximately 76% (Turner and Penn 2002) with observed pedestrians. The result illustrates clearly that individual people are unpredictable, but, in some ways at least, the aggregate is predictable.

This suggests one factor that makes the complexity of cities intelligible: the fact that useful patterns will appear at a sufficient scale. A rank-size relationship is one "law" that has been found in many aspects of cities; if the population size for a group of cities is plotted against their rank on a logarithmic scale, the result is nearly linear (Batty 2006; 2008). This change in population is directly relevant to social factors, in that total measured values of variables like economic output, income, patents, as well as crime, all scale reliably with population, and increase at a rate of about 15% more than linear; people even walk faster in larger cities (Bettencourt & West 2010, Bettencourt et al. 2007). The same scaling pattern appears also in properties closer to the domain of the architect and planner, such as the heights of tall buildings (Batty 2008) and the degree of connections between streets in a city (Jiang 2009). All of these are high-level patterns that apply to the aggregate only; Batty (2006) has shown that there is no discernible regularity as to where an individual city will appear in the size ranking over time, as their dominance rises and falls randomly over centuries.

Another factor is that some of these high-level patterns involve a relationship between the variables that we can design and plan, such as spatial configuration, and the social or economic factors we might desire in cities. The angular betweenness centrality of street segments within a given urban network gives a theoretical measure of how much traffic is likely to pass through any given street segment, and observations confirm that greater centrality corresponds to greater pedestrian and vehicle count, and that the network of major roads can be reliably found in any network (Hillier 2007). But, depending on the scale used, the same measurement can clearly pick out the location of local high streets, the known centres of local neighbourhoods, and where commercial activity is actually located, all as a function of the street geometry (Hillier 2007). Local shops, for example, will tend to be successful in zones of maximum centrality measured at a radius of about 1km; we can know this even for streets yet to be built, and use this information to plan. The same applies across far larger scales, even to the extent of indicating the major locations of commercial and economic activity on international street networks, in which the pattern of centrality values correlates with the economic output of nations—the higher a country's total centrality, the higher is its GDP (Hanna et al. 2013). An analysis of such centrality across the range of scales (Krenz 2017) indicates that the scale hierarchy seen among cities, and in distances between them, is also a property of the network itself. It is not matched by randomly generated networks, which suggests that these human spatial networks may be optimised for patterns of particular human activity.

Cities are optimised to make some patterns clear

Some patterns of cities appear to be discernible simply because cities are optimised to make these patterns evident. Of the many possible ways of arranging roads, spaces and buildings across a surface, real cities are a quite constrained subset. This is useful in any machine search for regularities, as it drastically reduces the space to be searched.

The agent models above, just like individual people, have a view only of their immediate surroundings, yet the movement they predict strongly resembles that of global centrality measures. Is the pattern one of the small scale or the large? Is it of the cognitive, phenomenological properties of moving pedestrians or of the structural properties of roads and space? Causally, these appear to be entirely independent of one another, and it has been noted of methods that study the structure of space that they "cannot account for the dynamics of movement" (Batty 2001). Those who analyse street structure have argued that activity in a city is driven solely by the properties of the network (Ma et al. 2018), whereas others who look at individual path choice see relevant visual cues and cognitive factors (Turner 2007, Emo 2014). The evidence supports both claims, not because they are necessarily related, but quite possibly because cities are so often shaped such that the same patterns appear at both scales.

The reasons for this are evident when considering what it means to navigate an unfamiliar part of the city without a map. Immediate visual cues frequently lead us toward longer, wider streets where our visibility is greater, and most of us can normally rely on these to lead efficiently to our destination because, if they did not, we would be lost in a labyrinth. Where natural footfall on a street does not bring many people in contact with a shop, that shop is more likely to fail as a business and disappear. The degree to which the large-scale properties are conveyed by the small can be measured, as in the space syntax measure of *intelligibility* (Penn 2003), which gives an assessment of how effective an area is at conveying this essential information.

Where we see evidence that large scale structural information correlates with small scale geometry, it need not take the complexity of the human visual system to reveal it. Agent models far simpler even than those above can be constructed to simulate random walks through a city by taking the street segment network as graph, with connections weighted by the angle at which streets meet; agents walk randomly but with a greater probability of continuing straight than making acute angled turns. The distribution converges quite rapidly, within 20 to 30 "steps", to approximate that of observed pedestrians and vehicles at distances ranging from neighbourhood to regional scales (Hanna 2020). This simple random process appears to mimic all the complexity of a city full of real human travellers.

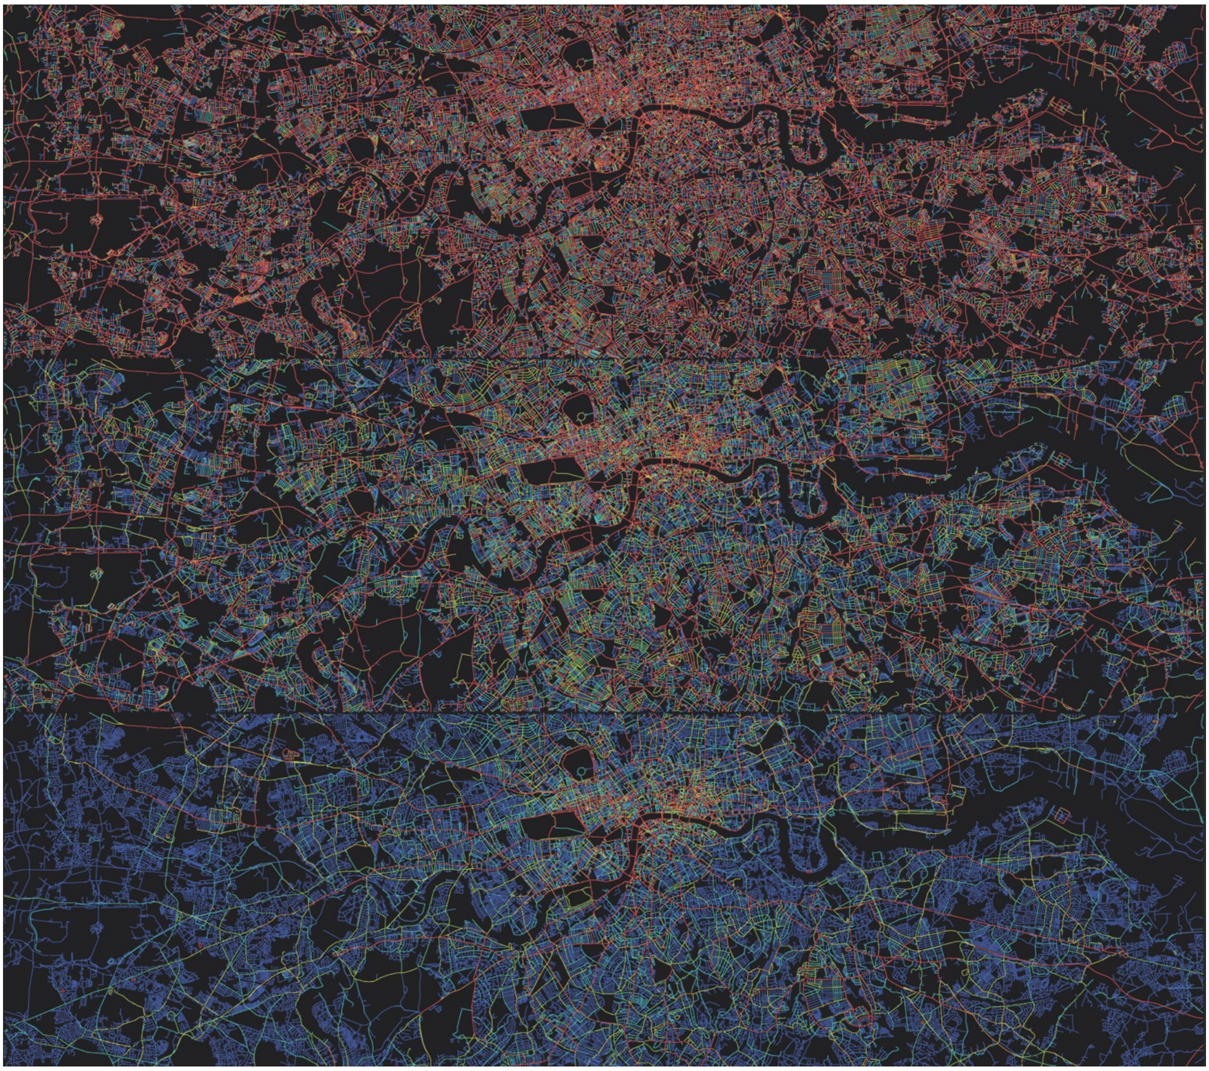

Figure 1: A simulated random walk on the street network of London, at 5, 26 and 100 iterations. Agents begin uniformly distributed (top) but aggregate on the more highly-trafficked roads after a short time (bottom), revealing the major routes and areas of traffic density.

This is significant because the route decisions made are entirely local, with agents unable to "see" beyond their current intersection, yet the model predicts real movement similarly to others that explicitly optimise longer range routes. Methods such as betweenness centrality (Hillier Iida 2005), which calculates optimal paths through graph nodes, and network analyses of continuity lines (Figueiredo & Amorim 2005) or natural streets (Jiang et al. 2008), which group together sequences of segments with minimal angles of turn, exploit information at some distance across the city to model movement. The implication is that some longer-range knowledge of the network is necessary for navigation, and that people optimise their routes accordingly. But both these longer-range methods and the locally informed random walks correlate well (with Pearson coefficients > 0.7) with movement, and with one another, which suggests, at least for the urban networks studied, that in real cities there is a rough equivalence between navigating based on knowledge of the street map and navigating based on immediate visual cues, and that the geometry of the street network is optimised such that it conveys the relevant information about distant routes to a naive traveller at any intersection (Hanna 2020). The same patterns are clear at both large and small scales.

Explicit optimisation is rarely likely to have been the cause of such intelligibility. In the case of New York's central park, for example, centrality analyses of a hypothetical street grid indicate that if the park did not exist the streets in its place would be less central and rarely used, simply due to asymmetries in the shape of the island (Al Sayed et al. 2009). Manhattan's planners placed the park exactly where modern computational methods would recommend, but without any such methods being available at the time. The causes of such decisions in real cities are often too complex to be known for certain: market forces and competition might determine the location of commercial property or of parks; cultural precedents might suggest resemblances to other known cities; the political pressures on design and planning are numerous. But to the extent that cities are intelligible, they are so because their patterns are obvious even when we are not certain of their underlying cause.

Non-discursive features also appear in data

Many of the qualities relevant to us in our own experience of the city are more complex even than we can precisely articulate. The style of buildings in a particular neighbourhood is intuitively recognisable to us as different from another, yet the precise features of those buildings that determine the difference are not easily described. It might seem that these non-discursive properties are elusive, or impossible to quantify, but the patterns are no less real for their complexity, and also there to be found by the machine.

To investigate properties of building form in Athens and London, Laskari (Laskari et al. 2008) used the shape of the combined footprint of buildings within an urban block as the unit for comparison, which captured essential properties of building width, density and uniformity in the shape and size of the internal courtyards hidden away from the street facade. Fourteen different measurements were taken for each unit, including straightforward values of perimeter and area, and more complex ones such as fractal dimension and quantities derived from the lines of sight within the courtyard spaces. All such measurements are legible automatically, using no more than rudimentary machine vision, to yield a fourteen-dimensional data point for each building block. When these are compared, 25 points for each of four neighbourhoods in Athens, in addition to Bloomsbury in London, clear clusters of points were seen to differentiate one neighbourhood from the next. While these are not perfectly separable, with some points overlapping, the differences coincide well with our own intuitive assessments of building style or type. The clusters within Athens are closer to one another than any of them is to London, and those neighbourhoods of a similar age are closer than those built a century apart. The result quantifies, for the machine, just those complex aesthetic properties that we would find so hard to describe.

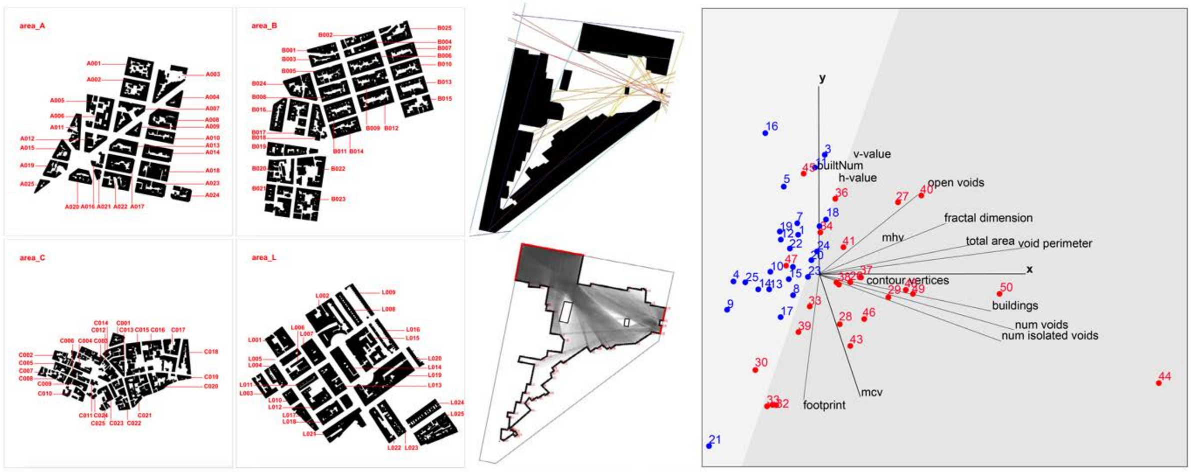

Figure 2: Building footprints, which differ from neighbourhood to neighbourhood, naturally form distinct clusters associated with these neighbourhoods in the space given by various automated measurements of their shape. (Image: redrawn from Laskari et al. 2008).

Although its judgements coincide with our own perception of stylistic categories, most of the measurements do not resemble something a human observer would notice. The machine "sees" the buildings in a plan view of the entire block, a view which is not given to its occupants or passers-by. The measure of fractal dimension might be thought of as a degree of complexity of this plan, but only approximately. Some measures of the lines of sight of the courtyards could be considered as a degree of convexity, but only approximately. The fact that machine and human judgements of the categories agree despite this difference between human vision and machine measurement suggests that such patterns are not dependent upon the selection of particular features to be used as inputs, and that they are likely to be found with relative ease regardless of which method is used.

Such a result is exactly what we would hope if we are concerned with a machine's capacity to pick out these relevant clusters, because we needn't be too concerned about choosing the correct input features. The best strategy, in this case at least, seems simply to have as many different features as possible. When analyses are compared using different groups of features (Hanna 2010), the correct clusters become more clearly differentiated as more features are used. This is not a case of having more dimensions in which to divide the classes, as is done in supervised learning; a fixed number of principal components is used to ensure each clustering is made in a space of equal dimensions. The results show that different machines classifying the sets of buildings converge both with one another, and with the correct identification of neighbourhoods, with greater numbers of inputs. The relevant patterns appear readily in the data drawn from the buildings, regardless of the particular representation used.

More complex patterns can be learned

When the relevant patterns do not so readily fall out of the data but are a more complex function of the input dimensions available, machine learning can be used to find them. Like the form of building plans above, the local configuration of roads differs depending on the land use, but the precise features relevant to these differences are not obvious. In recent analyses of the UK road network, Tasos Varoudis has used a type of neural network known as an auto-encoder, which is trained to extract the principle non-linear variances in the data from local samples of streets, and thereby map them to a subspace in which the most relevant differences are clear. The data are not pre-processed by taking any pre-determined measurements, but are instead presented directly to the network as street graphs, represented as adjacency matrices. The natural clusters of similarity that appear correspond almost exactly with the land use regions designated as urban areas, farmland, or natural landscapes in surveys such as the Corine Land Cover Inventory (2018).

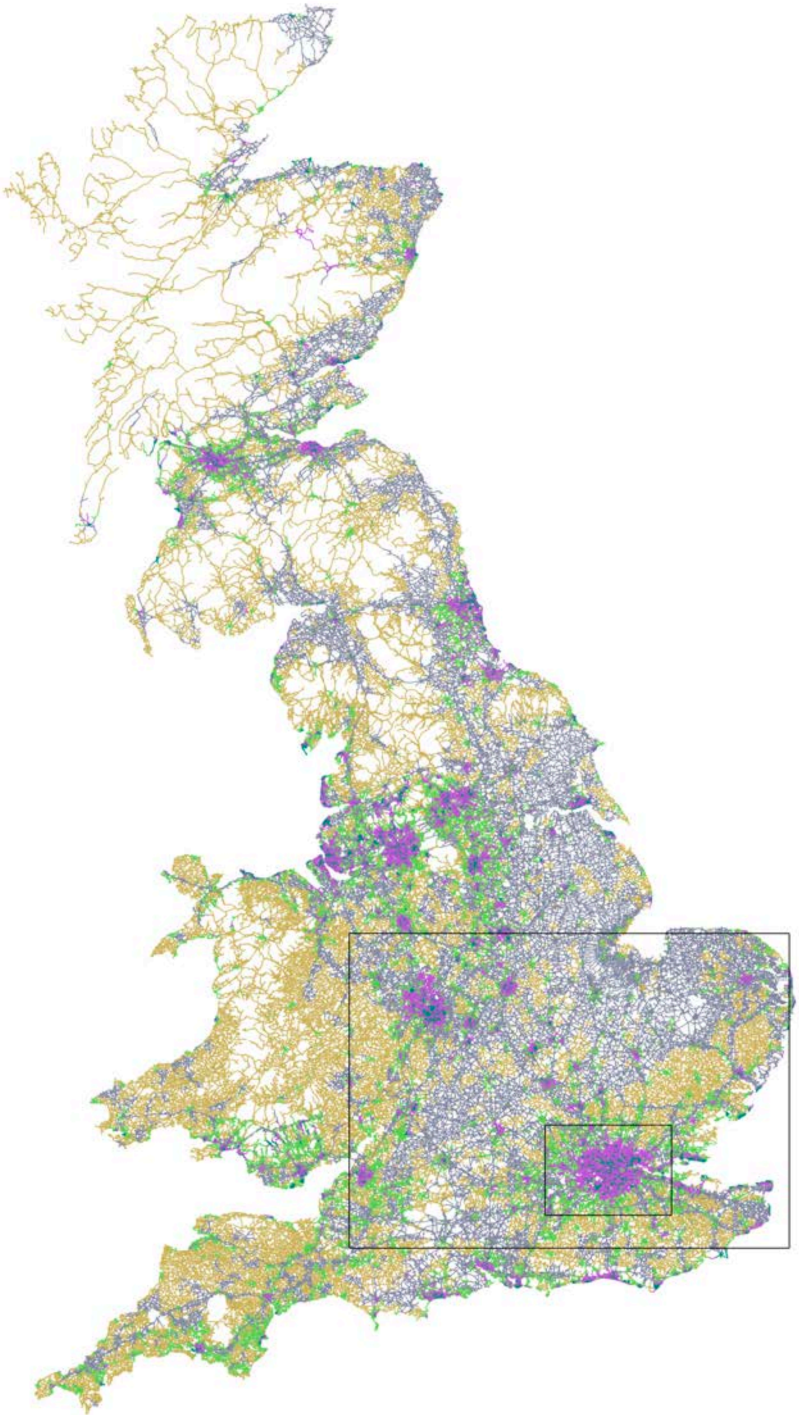

Figure 3: Differences in road morphology are clearly distinguished by a neural network, here revealing natural clusters which correspond with actual land use designations. (Image: Varoudis and Penn, 2020).

If the relevant classes or features we are looking for are known beforehand, but there is complexity in the input, supervised methods can be used to train the machine learning algorithm. Much larger urban graphs have been classified using their graph spectra, which represents the entire city as a vector in many more dimensions than the local samples above. Clustering cities in this high-dimensional space results in very little discernible pattern, but the geographical location of the cities can be used to tell the supervised learning algorithm what to look for. In Hanna (2009), a training set of cities is presented as input to a support vector machine, identifying each one as, for example, a European, or Asian, or North American, city. Once trained, the algorithm correctly classifies new cities with an accuracy between 75% to 85%, based entirely on their form.

Complexity often comes not from the scale of the sample, but from a considerable overlap in the input dimensions. In Thirapongphaiboon and Hanna (2019), centrality measures at varying scales were seen to correspond to different types of urban land use. Commercial buildings tend to be located on street segments with high values of closeness centrality at low radii, under 1.8 km, whereas business and industrial buildings correspond to higher radius measures from 1.8 to 7.2 km, but node count, or density of streets, is a more relevant measure. Residential use, by contrast, is marked by low centrality across the full range of radii. With much overlap, no single measure, scale, or selected group of such, makes the distinction between these uses clear in itself, but supervised learning uses the known classes (in this case commercial, business or residential) to derive a particular spatial signature that best describes each class. With this, the proportion of land use can be predicted using a multi-layer perceptron for street segments with an accuracy of more than 80% (Thirapongphaiboon and Hanna, 2019).

Much more specific land uses can also be identified by such spatial signatures, such as particular business types, or even locations of an individual chain. Silva (2017) used a random forest algorithm to predict, for a range of centrality measures of a given street segment, whether it is likely to contain types of business such as pubs, cafes or travel agencies, each of these being correctly identified more than 70% of the time. Some particular business chains, such as Waterstones' bookshops, could not be placed any better than chance, but Starbucks locations (and solicitors) were positively predicted at a rate of more than 80%.

Putting the patterns to use

If the preceding examples have focused entirely on the search and understanding of regularities in the data rather than the task of managing, intervening in, or designing the city, it is in part because this pattern recognition is the strength of machine learning. But it is also because pattern recognition is such a natural and intuitive part of our own cognition that its importance is overlooked, and because the necessity of coping with novel, larger and more complex patterns in cities has never been more acute. The rapidly changing requirements of cities, and the speed of their construction mean that decisions are more costly than ever—not only in the present but also long term socially and economically. The ability to project these patterns into new scenarios allows them to be tested in silico before committing, to use this knowledge to place a business where natural footfall will mean it will thrive, to target changes to streets so that the city can be navigated effectively, or build new sections of the city that remain naturally connected and vibrant.

The apparent limitlessness of complexity may seem to be a problem in that we will always find more of it, if we look deeper, if we have more data. The examples above predict only the long-term behaviour of many individuals, but individual behaviour (thankfully) and many lower level patterns may be forever beyond our ability to determine. It is fortunate, then, that the scale of the regularities we are able to discover happens to coincide with the scale of our intervention. We design for aggregates of many people, not for single individuals. We design for the climate over the span of years, not for the weather of a single day. Even to the extent that we could in principle forecast individual behaviour, as we might with increased access to large sets of personal data, this is not the level at which our design and planning decisions are made. The aim, for example, to design cities that bring diverse individuals into contact with one another, is a higher-level goal, which requires descriptions of many people in many spaces over extended times, and this is the level of description of complexity with which we need to contend. These complex patterns of the aggregate, which are most important, are also those which are most easily found in the data.

This trait suggests the reason why machine learning is useful in the context of the complex city, a reason too easily overlooked in the day to day training of learning models, which are judged on their success in prediction. The problem with prediction, in the sense of foretelling the outcome of specific events, is that it is not possible in the context of the wicked problems of the complex city, nor should we aspire to it. What is more useful to us is understanding, which, in the best case, is what the patterns extracted by machine learning will provide. Like the theory-led approach, the data-led approach can give us models that generalise sufficiently to tell us how the phenomena we care about, including social interaction, economic activity, movement and more, will occur in new and different urban environments. When faced with complexity, the recognition of patterns otherwise invisible due to the scale or form of data in which they appear is a way of seeing the city more clearly, of extending the limits of our own natural understanding, and, ideally, a means to inform better decisions.

References:

Al Sayed K, Turner A & Hanna S (2009) Cities as Emergent Models: The Morphological Logic of Manhattan and Barcelona, *Proc. 7th International Space Syntax Symposium*. pp. 001:1–001:12.

Batty M (2001) Agent-based pedestrian modeling. *Environment and Planning B: Planning and Design* **28**(3): 321–326.

Batty M (2006) Rank Clocks, *Nature* **444**, pp. 592–596.

Batty M (2008) Cities as Complex Systems: Scaling, Interaction, Networks, Dynamics and Urban Morphologies, in Robert A. Meyers (Ed.) *Encyclopedia of Complexity and Systems Science*, Springer.

Bettencourt LMA, Lobo J, Helbing D, Kühnert C & West GB. (2007) Growth, innovation, scaling, and the pace of life in cities, *Proc. Natl Acad. Sci. USA* **104**(17), pp. 7301–7306.

Corine Land Cover Inventory (2018) European Union, Copernicus Land Monitoring Service 2018, European Environment Agency (EEA)

Dunbar, R (2014) *Human Evolution*. Pelican Books

Emo B (2014) Seeing the axial line: Evidence from wayfinding experiments. *Behavioral Sciences* **4**(3): 167–180.

Figueiredo L & Amorim L (2005) Continuity lines in the axial system, *Proc. 5th International Space Syntax Symposium*. pp. 161–174.

Hanna S (2009) Spectral Comparison of Large Urban Graphs, *Proc. 7th International Space Syntax Symposium*. pp. 039:1–039:11.

Hanna S (2010) Design agents and the need for high-dimensional perception, in Gero JS (ed) *Design Computing and Cognition DCC'10,* Springer.

Hanna S, Serras J & Varoudis T (2013) Measuring the Structure of Global Transportation Networks, *Proc. 9th International Space Syntax Symposium*. pp. 060.1–060.13.

Hanna (2020) Random walks in urban graphs: A minimal model of movement. *Environment and Planning B: Urban Analytics and City Science*, 48 (6): 1697–1711.

Hillier B (2007) Spatial Sustainability in Cities: Organic Patterns and Sustainable Forms, in Koch D, Marcus L and Steen J, *Proc. 7th International Space Syntax Symposium*. pp. K01:1–K01:20.

Hillier B and Iida S (2005) Network and psychological effects in urban movement. *Proceedings of Spatial Information Theory*, pp.475–490. Berlin: Springer.

Jacobs J (1961) *The Death and Life of Great American Cities.* New York: Random House.

Jiang B (2009) Ranking spaces for predicting human movement in an urban environment. *International Journal of Geographical Information Science* **23**(7): 823–837.

Jiang B, Zhao S & Yin J (2008) Self-organized natural roads for predicting traffic flow: a sensitivity study, *Journal of Statistical Mechanics*. **2008**: P07008.

Krenz K (2017) Regional Morphology: The Emergence of Spatial Scales in Urban Regions, *Proc. 11th International Space Syntax Symposium*. pp. 74.1–74.23.

Laskari S, 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 Gero JS and Goel A (eds) *Design Computing and Cognition '08,* Springer.

Ma D, Omer I, Osaragi T, et al. (2018) Why topology matters in predicting human activities. *Environment and Planning B: Urban Analytics and City Science* **46**(7): 1297–1313.

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

Silva JP (2017) The Syntactic Signature of Starbucks' Locations: Toward a machine-learning approach to location decision-making, *Proc. 11th International Space Syntax Symposium*. pp. 171:1–171:8.

Thirapongphaiboon T and Hanna S (2019) Spatial Distribution of Building Use: Recognition and Prediction of Use with Machine Learning, *Proc. 12th International Space Syntax Symposium*. pp. 194:1–194:30.

Turner A (2007) To move through space: Lines of vision and movement. *Proc. 6th International Space Syntax Symposium*, pp. 037.1–037.12.

Turner A & 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): 473–490.

Varoudis T & 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.

United Nations, Department of Economic and Social Affairs, Population Division (2019). *World Urbanization Prospects: The 2018 Revision (ST/ESA/SER.A/420)*. New York: United Nations.