Computers, Environment and Urban Systems 95 (2022) 101802

Contents lists available at ScienceDirect

Computers, Environment and Urban Systems

journal homepage: www.elsevier.com/locate/ceus

Estimating generalized measures of local neighbourhood context from multispectral satellite images using a convolutional neural network

Alex Singleton *, Dani Arribas-Bel, John Murray, Martin Fleischmann

Geographic Data Science Lab, Department of Geography and Planning, University of Liverpool, UK

ARTICLE INFO

Keywords:
Deep learning
Convolutional neural networks
Urban morphology
Multispectral satellite imagery
Cluster analysis

ABSTRACT

The increased availability of high-resolution multispectral imagery captured by remote sensing platforms provides new opportunities for the characterisation and differentiation of urban context. The discovery of generalized latent representations from such data are however under researched within the social sciences. As such, this paper exploits advances in machine learning to implement a new method of capturing measures of urban context from multispectral satellite imagery at a very small area level through the application of a convolutional autoencoder (CAE). The utility of outputs from the CAE is enhanced through the application of spatial weighting, and the smoothed outputs are then summarised using cluster analysis to generate a typology comprising seven groups describing salient patterns of differentiated urban context. The limits of the technique are discussed with reference to the resolution of the satellite data utilised within the study and the interaction between the geography of the input data and the learned structure. The method is implemented within the context of Great Britain, however, is applicable to any location where similar high resolution multispectral imagery are available.

1. Introduction

The research presented in this paper develops a new national multidimensional and generalized measure of urban spatial structure. The derivation of such contextual measures have a rich academic history, and are used for a variety of applications related to the spatial targeting of policy or other interventions. In this example we focus on the built environment through a methodologically innovative analysis of high-resolution satellite derived imagery, although as discussed later, the output measures from this work also have the potential to enhance more general-purpose geodemographic classifications through the integration of additional population and socioeconomic data. The key innovation of this work is therefore to introduce and implement an entirely new way in which the built environment context of small areas can be classified from high-resolution satellite data using a convolutional neural network. Our new measure of "context" is made available at a very high spatial resolution for a case study covering the entire extent of Great Britain. The input measures and output classification are all placed within the public domain, and available from: https://dataverse.harvard.edu/dataverse/sat_cnn and for reproducibility, the code can also be found here: https://github.com/GDSL-UL/sat_cnn.

The paper is organised as follows: we first describe the context for this work, then those data used in the study as well as the many and varied challenges encountered in its ingestion and cleaning; next, we provide an overview of the specification and implementation of our model; then, after running the model we present a case study application illustrating how the outputs might be operationalised to derive insight about urban context, exploring this representation in relation to internal and external validation metrics. We close with some concluding remarks about the potential of our framework to advance our understanding of the contexts of cities and their spatial structure.

1.1. Framing the challenges and opportunities for the measurement of local urban context

Those ways in which the outcomes of human behaviour are influenced by the context of those places in which they live, interact and work are fundamental to our understanding of urban systems. The quantitative measurement of urban spatial structure has a legacy of development within Urban Geography and Planning (Batty, 2013), both as a normative tool in practice or as an epistemological construct aiming to advance knowledge about how cities work. There are multiple ways in which researchers have attempted to represent context; summarised here through three distinctive computational approaches (Fig. 1 A-C):

* Corresponding author.
E-mail address: alex.singleton@liverpool.ac.uk (A. Singleton).

https://doi.org/10.1016/j.compenvurbsys.2022.101802
Received 18 November 2021; Received in revised form 28 March 2022; Accepted 1 April 2022
Available online 19 April 2022

A) Space Syntax OpenMapping of street betweenness centrality ("choice") showing the distribution of likely movement (https://spacesyntax-openmapping.netlify.app)

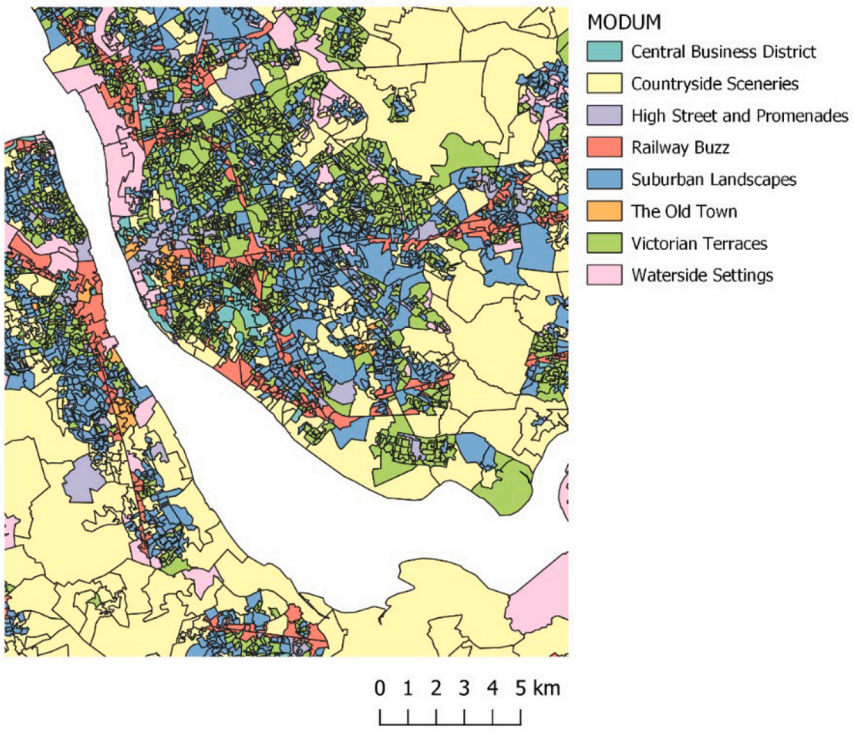

B) The MODUM typology of neighbourhoods describing multidimensional attributes of the built environment Alexiou, Singleton & Longley (2016)

Fig. 1. Different approaches to the measurement of urban context: showing the urban area of Liverpool, North West UK (Morton, Marston, O'Neil & Rowland (2020)).
Data 2021 © Space Syntax Ltd.
Contains OS data © Crown copyright and database right 2021.
Historic England 2015.
Office for National Statistics licensed under the Open Government Licence v.3.0.
© UKCEH 2017.

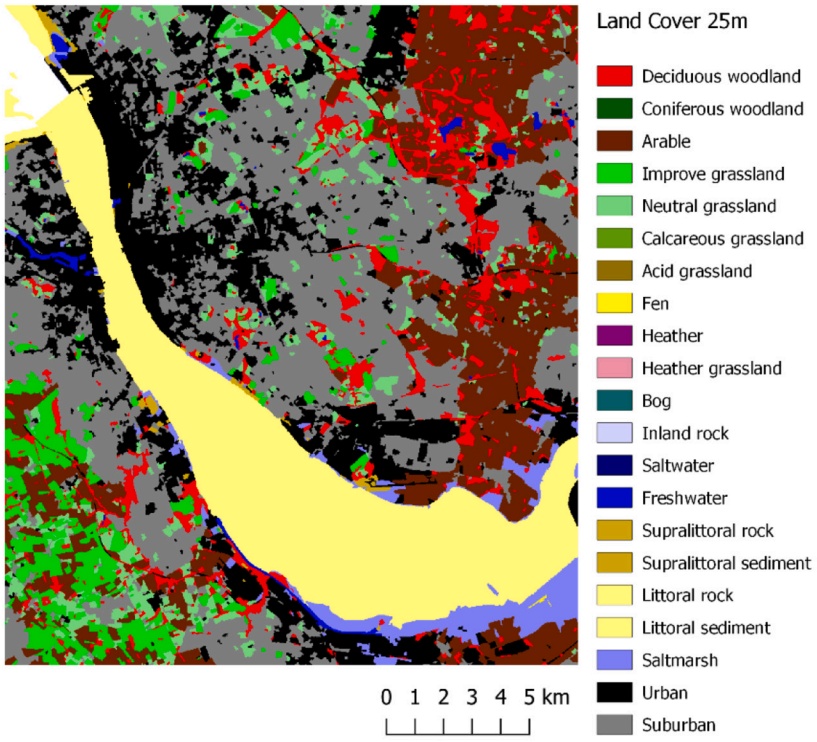

C) The 20m resolution 2019 Landcover Map from the UK Centre for Ecology & Hydrology
Morton, Marston, O'Neil & Rowland (2020)

Fig. 1. (*continued*).

those based on geometry which ascribe morphometric attributes derived from representations of the built environment, those which are based on the summary of tabular measures of built or social attributes collated for small area geography (often called geodemographics); and those derived from uniform gridded data such as remotely sensed imagery. Although the approaches vary in their theoretical rationale and implementation, they typically share common goals to better derive a representation of the complex organisational structure of and beyond urban areas. This includes the measurement of salient characteristics, the derivation of insight about specific aspects of latent structure that are not directly observable, and the relationship between such observations to function and evolutional form.

In a general sense, morphometric analysis (Fig. 1A) aims to provide a lexicon of urban structure; derived through various approaches to metric derivation that typically examine aspects of size, shape and interconnectivity of digitised representations of urban features (Hijazi et al., 2016). Such measures are often ascribed at the scale of the features being assessed; for example, connectivity scores for street geometry (Araldi & Fusco, 2019; Boeing, 2018) or the spatial arrangement of buildings (Vanderhaegen & Canters, 2017); although may also be complied for ancillary derived features or summarised through aggregate units (Fleischmann, Feliciotti, Romice, & Porta, 2021). Although not universally the case (Araldi & Fusco, 2019), such measures tend not to be combined, and are mostly presented as univariate indicators.

By contrast, tabular approaches to the summary of only urban structure are reasonably rare (Fig. 1 B) and tend to apply either combinatorial indexes or cluster analysis to a set of small area standardized measures and may conflate both social and built characteristics into multidimensional scores or categories. Such approaches overlap with morphometric analysis when derived measures are considered within zonal geography, although typically draw on a wider range of inputs beyond those derived from geometric properties of the built environment. Data sources for such studies typically include attributes derived from Census, surveys, administrative or commercial data, and as a technique where typological representations are created, are commonly referred to as geodemographic classification (Webber & Burrows, 2018). However, divergent from the work presented here, such classifications typically describe both the social and built characteristics of small areas (Singleton, Alexiou, & Savani, 2020; Singleton & Spielman, 2014; Spielman & Singleton, 2015b), and there are much more limited examples of applying geodemographic classification techniques to purely built environment data (Alexiou, Singleton, & Longley, 2016).

Data derived from remotely sensed images such as satellites and other airborne sensors also have a legacy of use within this domain through the development of land cover / land use (LCLU) classifications (Fig. 1 C). LCLU classification typically considers the spectral 'signatures' of remotely sensed images, which are generated as a result of how different surfaces reflect distinctive wavelengths of light. Such classifications are often compiled at a pixel level and may also include measures that relate to the surrounding values within a moving window around each pixel. In many cases the classes are often predefined before the classification is carried out and are usually discrete, but can also be

Fig. 2. True Colour Image comparisons.
Modified Copernicus Sentinel data 2021.

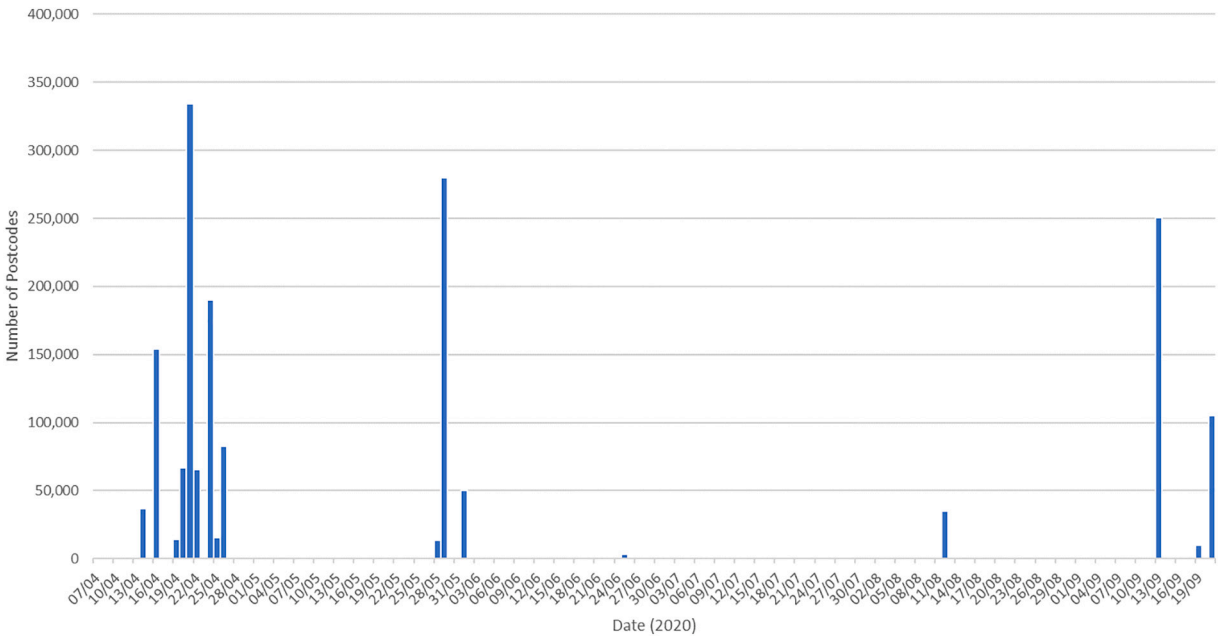

Fig. 3. The frequency of images by date that were extracted for each postcode.

fuzzy. Alternative approaches classify objects from delineated areas of imagery (Helber, Bischke, Dengel, & Borth, 2018; Shendryk, Rist, Ticehurst, & Thorburn, 2019). In some studies, vector data such as built environment features are integrated into the models in their creation or as post facto processing. LCLU methods are also differentiated from object detection methods that attempt to extract specific geographic features (e.g buildings, roads etc) from high resolution imagery (Blaschke, 2010).

With the overarching aim of this paper being to derive a measure of context of where people live, this is divergent from the majority of more traditionally orientated LCLU studies. However, there are a growing number of applications from within this literature that implement or evaluate innovative machine learning methodologies when applied to remotely sensed imagery, notably using convolutional neural networks and other related techniques (Cheng, Han, & Lu, 2017). Such methods frame our work, and have been shown to improve land cover classification, and particularly in case studies where high-resolution spatial imagery are available (Chen, Ming, & Lv, 2019a, 2019b; Gaetano, Ienco, Ose, & Cresson, 2018; Liao, Cao, Wang, & Xu, 2022; Zhang et al., 2018). Although a number of innovative techniques have been proposed to mitigate such issues (Fan, Feng, Wang, Yan, & Zhang, 2020), there remain common challenges in many studies when supervised methods are being applied, as these rely upon labelled data (Bhosle & Musande, 2019; Feng et al., 2019; Heryadi, Miranda, Heryadi, & Miranda, 2019;

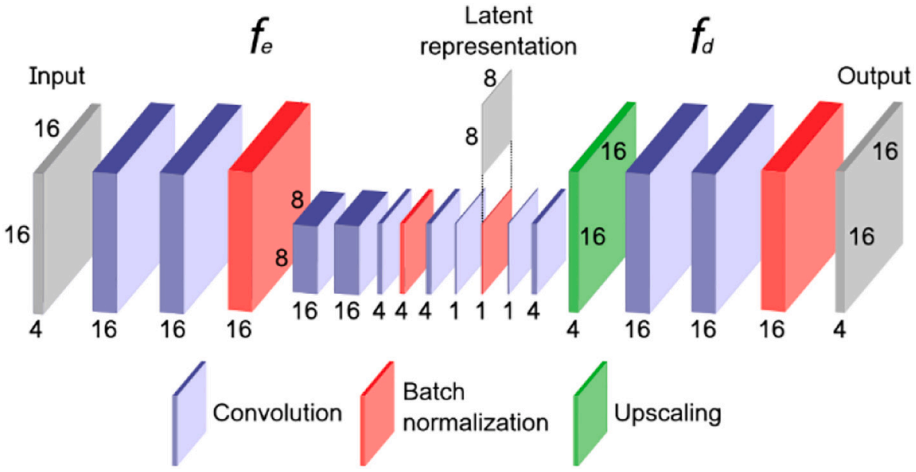

Fig. 4. Summary of the Autoencoder Architecture (Source: Authors Own).

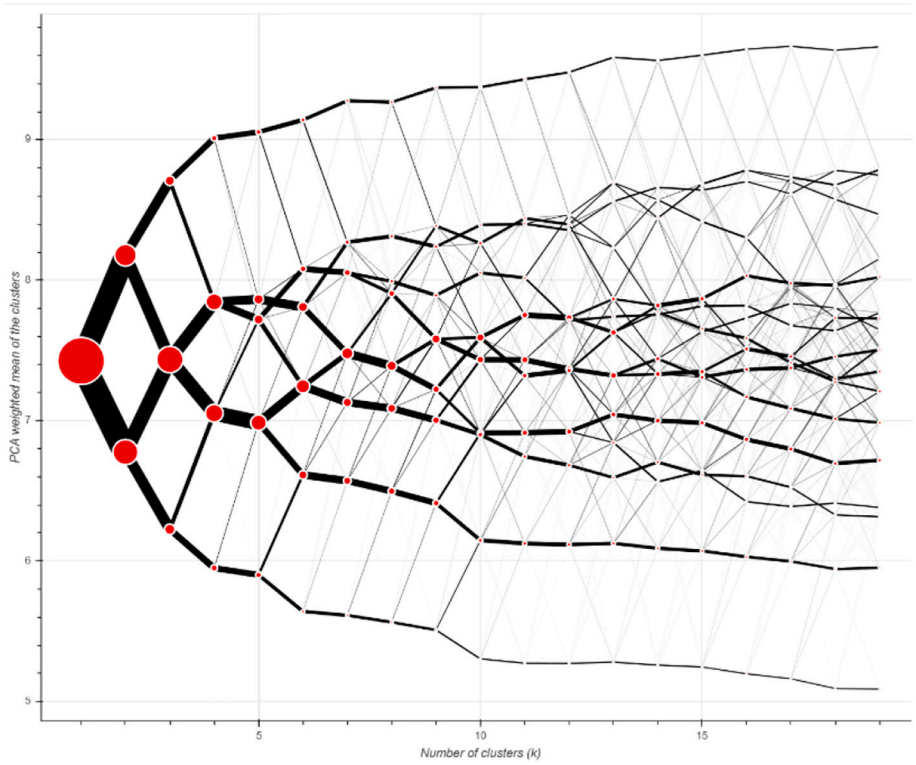

Fig. 5. A Clustergram showing the potential suitability of k = 2 to 19 solutions.

Kareem et al., 2021). Similar issues care found within this paper, both as a result of the data used, but also by design as our theoretical framework assumes no prior knowledge about the geography of context. A more related literature aligned to the work presented here is semantic segmentation (Balarabe & Jordanov, 2021; Nogueira, Penatti, & dos Santos, 2017; Weng, Mao, Lin, & Guo, 2017) which involves the partitioning of an image into regions / objects sharing similarity where models can be applied to extract semantic meaning (Kotaridis & Lazaridou, 2021). As will be described in the next section, a key difference here is that our objects are kept to a uniform size; and there are no prior labels semantically differentiating content, with this structure being learnt and then inferred through the application of our specified method.

The approaches briefly presented above all emerged from very different histories and disciplinary backgrounds, however share commonality in their ambitions to translate rich spatial data into insight about the structure of the built or natural environment. Within Great Britain, morphometric and tabular approaches have typically offered greater detail in urban over rural areas, as highlighted in Fig. 1. Morphometric approaches are more prevalently derived for specific policy / planning applications located within urban areas (Dibble et al., 2017; Hillier, Leaman, Stansall, & Bedford, 2016; Salvadori, Badas, di Bernardino, Querzoli, & Ferrari, 2021; Venerandi, Zanella, Romice, Dibble, & Porta, 2016), whereas tabular geodemographic approaches are often derived for a national extent, and in addition to capturing some built environment characteristics, will often aim to describe a brevity of population characteristics or behaviours (Gale, Singleton, Bates, & Longley, 2016). Perhaps unsurprisingly given the focus on the natural environment, the presented LCLU (Fig. 1 C) only differentiates between urban and suburban areas, with much greater detail provided for non-urban areas. Elsewhere, there are more urban focused applications of

A) UPRN Coloured by their associated postcode cluster (generated from the raw CNN output)

B) UPRN Coloured by their associated postcode cluster (generated from the smoothed CNN output)

Fig. 6. Pilot cluster analysis with the raw and smoothed data inputs.
High Resolution (25 cm) Vertical Aerial Imagery [JPG geospatial data], Scale 1:500, Tiles:sj4289, sj4288, sj4189, sj4188, Updated: 29 October 2018, Getmapping, Using: EDINA Aerial Digimap Service, <https://digimap.edina.ac.uk>, Downloaded: 2021-05-28 09:35:02.49.
Contains OS data © Crown copyright and database right 2021.
Contains GeoPlace data © Local Government Information House Limited copyright and database right 2021.
Source: Office for National Statistics licensed under the Open Government Licence v.3.0.

LCLU; for example, the European Environment Agency Urban Atlas[1], which delineates built up urban areas into multiple categories of density, industrial and other aspects of supporting infrastructure.

For many normative applications of contextual measures, and particularly in tabular approaches, these are concerned with the characterisation of different types of "neighbourhood", which although might be argued as varying in their definition / ontology (Forrest & Kearns, 2001; Galster, 2001), they share an ambition to capture important effects on behaviours or outcomes that extend beyond those which are observable at the individual scale (Galster, 2013; Sampson, 2019; van Ham, Manley, Bailey, Simpson, & Maclennan, 2012).

[1] https://land.copernicus.eu/local/urban-atlas/urban-atlas-2018

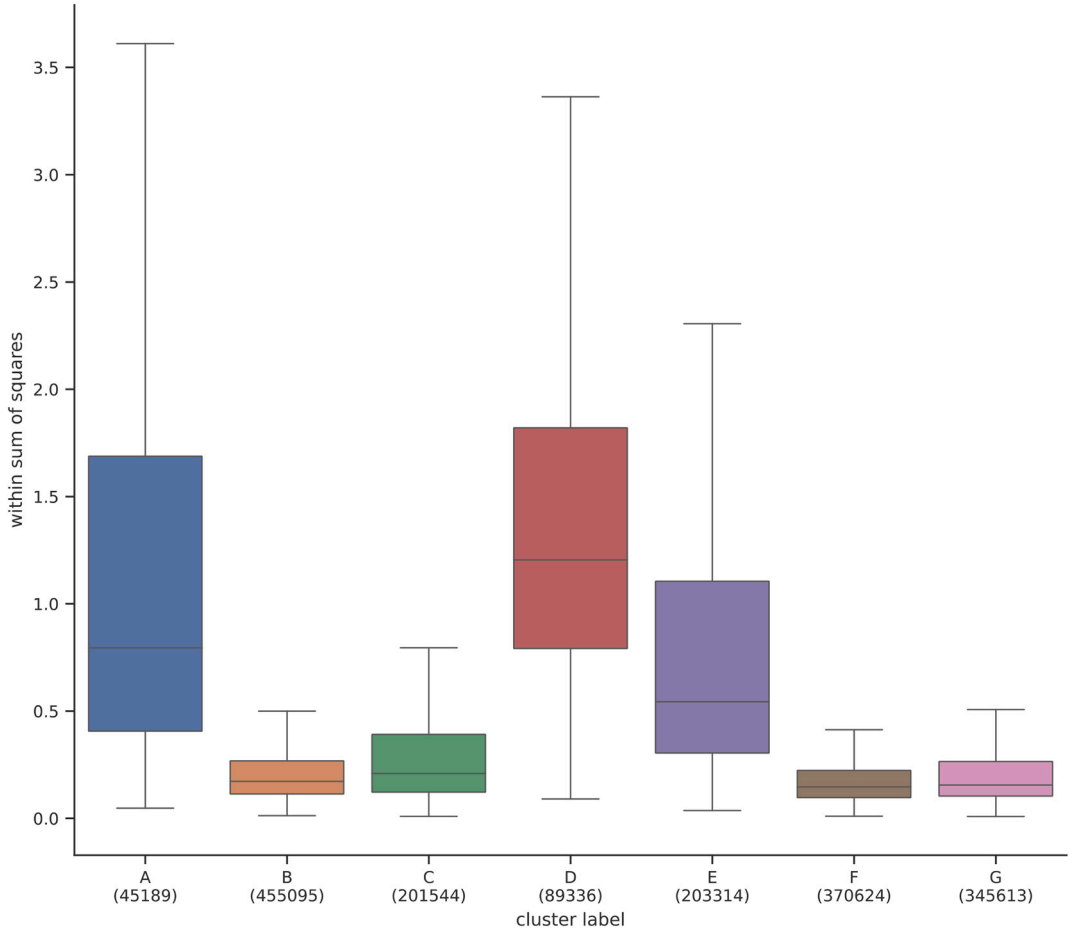

Fig. 7. Within Sum of Squares aggregated by Cluster; with the frequency of postcodes in brackets.

Although conceptually this might be argued as suboptimal (Wei, Rey, & Knaap, 2021), for analytic simplicity or necessity many urban applications utilise zonal geographies created for statistical, political or administrative purposes. Such shorthands for delineated "neighbourhood" zones lend their characterisation through tabular forms of data where it may be possible to compile measures directly for the chosen geographic units. Comparatively, although LCLU can also be derived for discrete zonal geography or land parcels; this is less common as their inputs are primarily sourced from continuous fields (raster images), and indeed, such representations also remain a common output format for these measures. For morphometric measures, these tend to be ascribed to either discrete built environment objects or their atomistic components, and again do not easily translate into zonal geographies.

Morphometric, tabular and LCLU approaches bring complementary insight to our understanding of urban contexts that hitherto remain predominantly decoupled given their different disciplinary origins and divergent uses. We position our new framework as a supplement rather replacement of existing approaches, and as a new theoretical lens through which local contextual differentiation can be measured more flexibly alongside a nuanced definition of "neighbourhood" geography. As discussed earlier, the use of remotely sensed data for LCLU is well established, and their use in urban applications more generally is not new (Imhoff, Lawrence, Stutzer, & Elvidge, 1997; Mesev, Longley, Batty, & Xie, 1995). However, we make the case that applications of remotely sensed data are likely to gain increased importance and utility in urban settings: firstly given the proliferation of newer and higher resolution data that are being generated as a result of advancing space technology (Esch et al., 2020; Zhu et al., 2019), and secondly as a result of advances in the fields of computer vision and machine learning. However, realising this potential requires notable methodological advancement (Ibrahim, Haworth, & Cheng, 2020; Singleton & Arribas-Bel, 2019) given that the vast amount of information that is potentially uncoverable within such data are stored in highly unstructured forms, provided as images, colour patterns, and additional bands recorded by other on-device sensors. We would argue that potential value contained within such data are also increasing given that they are framed by an era when the funding of many traditional large-scale data collection exercises such as the census are under fiscal constraint; and when there are growing methodological challenges in large surveys such as declining response rates (Spielman & Singleton, 2015a).

Much of the recent work in urban geography / planning exploiting imagery data has tended to focus on single issues related to measures of urbanization, economic activity or population, with more recent work beginning to model their evolution over time (Bennett & Smith, 2017). Such studies are typically implemented to either derive global comparative measures (Arribas-Bel, Patino, & Duque, 2017; Lloyd et al., 2019; Pokhriyal & Jacques, 2017; Proville, Zavala-Araiza, & Wagner, 2017); or are applied within low- and middle-income countries (LMICs) where other sources of granular data may be less prevalent (Jean et al., 2016; Mahabir, Croitoru, Crooks, Agouris, & Stefanidis, 2018; Steele et al., 2017; Wurm & Taubenböck, 2018). There are emerging but more limited examples of the application of remotely sensed data to create more detailed measures of the socio-economic landscape (Weng, Quattrochi, & Gamba, 2018).

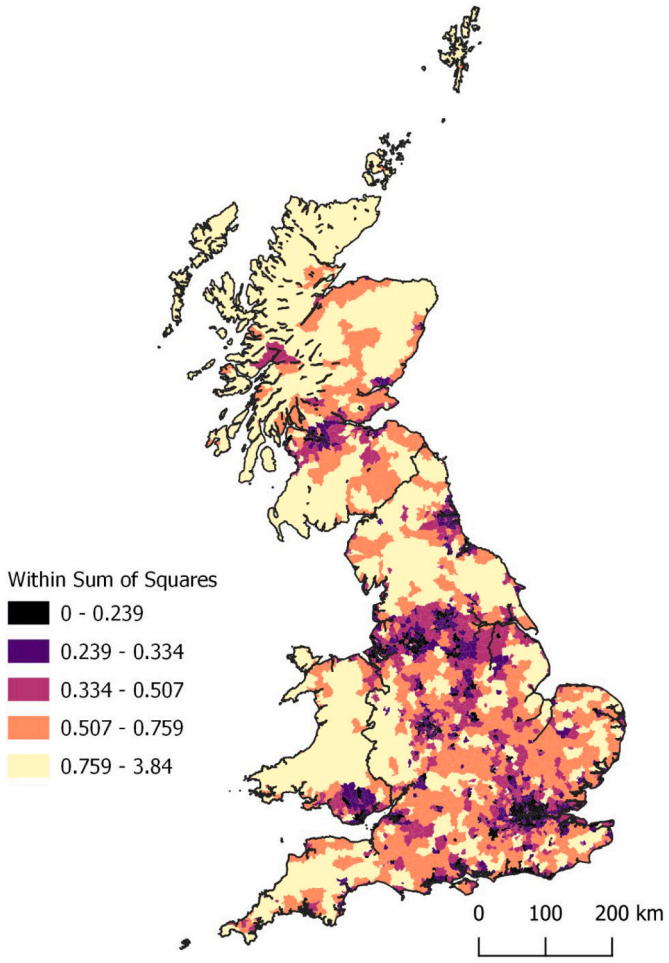

Fig. 8. Within Sum of Squares aggregated by Postcode District.
Postal Boundaries © GeoLytix copyright and database right 2012.
Contains Ordnance Survey data © Crown copyright and database right 2012.
Contains Royal Mail data © Royal Mail copyright and database right 2012.
Contains National Statistics data © Crown copyright and database right 2012.

2. Methodology

2.1. Data - measuring local spatial structure from space

The overarching objective of this paper was to present a new method of measuring local context that could be flexibly applied within different settings to meet variegated definitions of neighbourhood. As such, the aim here was not to ascribe a measure of context with universal coverage for all land area within the UK, as might be the case in LULC, but rather to focus on those localities where people live, describing these areas in terms of their varying context. Within Great Britain, unit postcodes provide a useful base geography upon which to build such a measure which on average identify around 13 residential addresses; although postcodes do capture both residential and business addresses, and also are not always mutually exclusive. In Great Britain there are 1,710,715 current postcodes which broadly follow a spatial distribution aligned with residential density; so more are found within closer proximity in urban versus rural areas. A database of georeferenced postcodes was supplied by the Office for National Statistics under an open licence.[2]

The second main source of data are high resolution multi spectral satellite imagery. The highest resolution open satellite data available with coverage for Great Britain are supplied by the European Space Agency (ESA) and derived from Copernicus Sentinel 2. ESA publish the data in 12 spectral bands, at 60 m, 20 m and 10 m resolution, as well as a true colour image (TCI) generated from the visible red, green and blue (RGB) bands. For the purpose of this exercise, we only used the 10 m resolution bands for RGB and Near Infrared; utilising the RGB bands to create a bespoke TCI. We opted to build our own TCI as the ESA composite observed a clip range of values at the 0.1% and 99.9% percentiles which resulted in loss of detail and clarity at the lower and higher ends of the spectral range. Although this issue affects only 1 in 1000 pixels, these appeared disproportionately represented on both dark and light roofs of buildings resulting in the edges of these features becoming fuzzy (Fig. 2). To avoid this potential information loss a bespoke TCI was generated where clipping at various percentiles were evaluated, with the final choice being to not clip at the lower end since this comprised very few observations. We then apply white balancing at four standard deviations from median, scaled to the range of 0 to 1 where values greater than one were considered outliers. There were additional benefits to this re-processing. For example, because ESA clipping operations are applied on a tile-by-tile basis, this resulted in different levels of brightness across adjacent TCI tiles; and as such, our re-processing which covered the whole of the GB extent returned a more consistent level of brightness across all images; with such normalisation being acutely important for nationally comparative purposes. An example of the ESA TCI and our re-processed image is shown in Fig. 2 A and B; with the latter offering greater clarity.

As discussed, the target geography for the output of this work was a measure derived for each unit postcode, and as such, a lookup was required between these locations and the ESA Copernicus Zone corresponding to the image tile extent. ESA supply a set of vector multi-polygons that were used in a large point in polygon operation relative to all postcode locations. These are provided in WGS84 so were converted to UTM to avoid spherical distortion (Karney, 2011) and loaded into a spatial database.

The analytical framework discussed in the next section was implemented with TensorFlow, which is an Open Source library for the development and training of machine learning models. We utilised this within a GPU (NVIDIA RTX8000) environment to enable more rapid (and feasible) computation. Input to TensorFlow is called a Tensorset, which is stored in GPU memory; and in this case comprised a separate tensor (multidimensional array) for each postcode composed of four dimensions (Near Infra Red–NIR, Red, Green, Blue). Sentinel 2 data were matched to each postcode using our database of ESA manifest tables against a date range and a cloud free filter. The optimal solution for temporal consistency would be full coverage of Great Britain for single cloud free day, however this is not realistic for all locations. As such, we sought images between 1 April 2020 to 30 September 2020 and aimed to extract cloud free images for each postcode within this date range. A distribution of these images for all live postcodes are shown in Fig. 3; spanning the period 7th April 2020 through to 21st September 2020. From the available swaths, only those where the ESA metadata described 10% or less cloud cover filter were considered; which is a score ascribed for the entire swath.

Given that the source imagery was at 10 m resolution, a 16 × 16 (160 m × 160 m) square buffer was created for each postcode, thus forming each tensor. A tensor requires a square and symmetric matrix so it was not possible to utilise other buffer shapes (e.g. circle) or non-uniform geometry. This specific size was chosen as the closest to the median unit postcode area for Great Britain, which was derived by Voronoi polygons. As such, when we refer to the "context" of a postcode, empirically this is defined as a square buffer of 160 m around each postcode centroid. At a conceptual level, the scale at which context is experienced might be argued as better aligned to a street geography rather than a buffer zone. However, the utility of considering a wider area brings benefit in that this captures factors beyond a postcode

[2] https://geoportal.statistics.gov.uk/datasets/ons-postcode-directory-may-2021

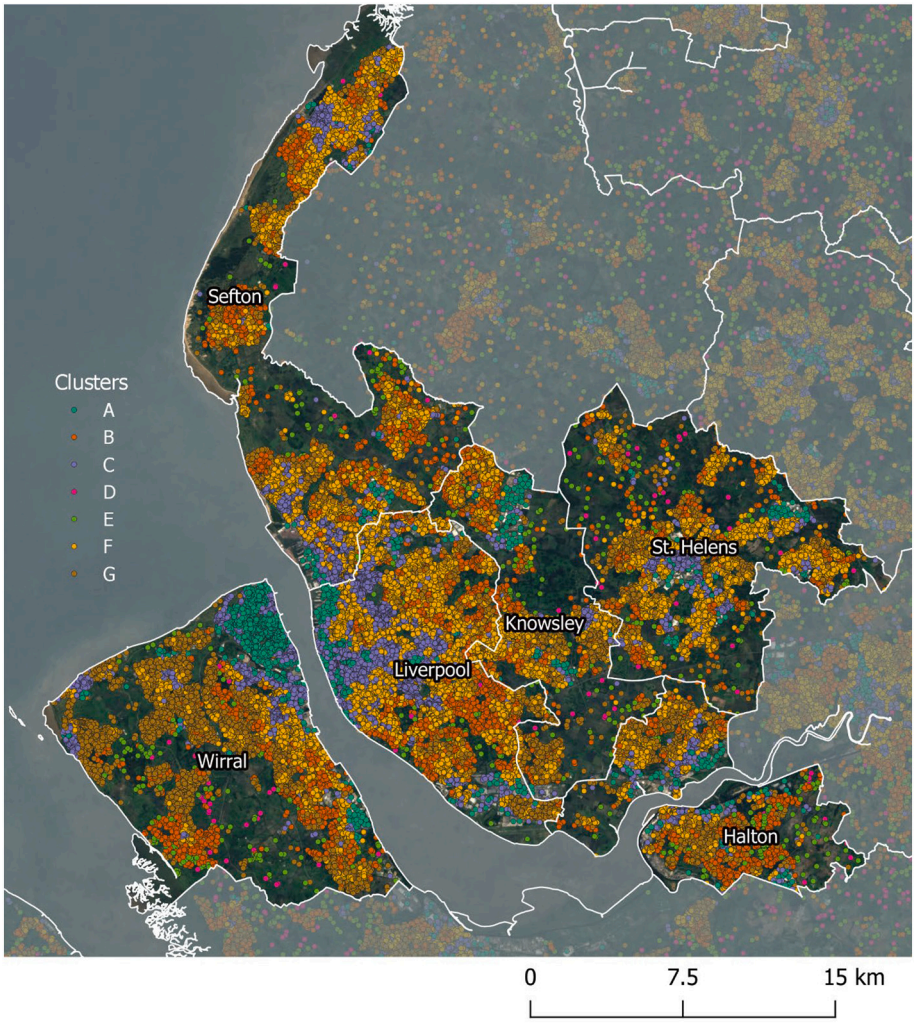

Fig. 9. The Liverpool City Region contextual Postcode geography.

Table 1
Ancillary data used to describe the clusters.

Measure	Geography	Source	Year/ Month
Limited Companies Postcode Count	Postcode	Valuation Office Agency; Companies House Register	2021/ 02
Council Tax Bands #	LSOA	Valuation Office Agency; Housing Stock Statistics	2019/ 09
Property Type #	LSOA	Valuation Office Agency; Housing Stock Statistics	2019/ 09
Year Built Bands #	LSOA	Valuation Office Agency; Housing Stock Statistics	2019/ 09
Postcode Population	Postcode	Office for National Statistics /Statistics Scotland; 2011 Census	2011/ 03
Postcode NDVI	Postcode	European Space Agency; Sentinel 2	2021/ 02
House Type #	Postcode	Ministry of Housing, Communities & Local Government; Energy Performance Certificate	2020/ 21
Road Features	Network	OS Open Roads	2020/ 21

Notes: # - These data were only available in England and Wales.

centroid that may also contribute to context. For example, a postcode surrounded by greenspace may be perceived very differently to one that is characterised by noisy roads or an industrial area. From the available swaths, buffer images that were closest to the swath pixel level aggregate median were extracted; alongside additional checks for damaged pixels.

2.2. Model design and training

The overarching aim of this analysis was to create a set of representative measures that describe the characteristics of the area surrounding every GB postcode as recorded within the associated input tensor set (multiple tensors) from the satellite image. To this end, we applied a Convolutional Neural Network (CNN) to derive a set of discriminative features that retain the information contained in the satellite data but effectively compress its dimensionality. Convolutional Neural Networks, the workhorse of deep learning (Lecun, Bengio, & Hinton, 2015), are a family of machine learning techniques that has seen wide use in a range of recent urban spatial problems related to the identification of features or the characterisation of different types of imagery (Ibrahim et al., 2020). Applications to imagery data are quite diffuse but include: the characterisation of street imagery (Comber, Arribas-Bel, Singleton, & Dolega, 2020; Law, Seresinhe, Shen, & Gutierrez-Roig, 2020; Liu, Silva, Wu, & Wang, 2017; Middel, Lukasczyk, Zakrzewski, Arnold, & Maciejewski, 2019), measuring the evolution of urban features (Saeedimoghaddam & Stepinski, 2020), regional movement (Yang & Gidófalvi, 2020), and detectingvarious kinds of urban objects (Alizadeh Kharazi & Behzadan, 2021; Branson et al., 2018; Campbell, Both, Sun, & Chayn)., 2019; Kang, Körner, Wang,

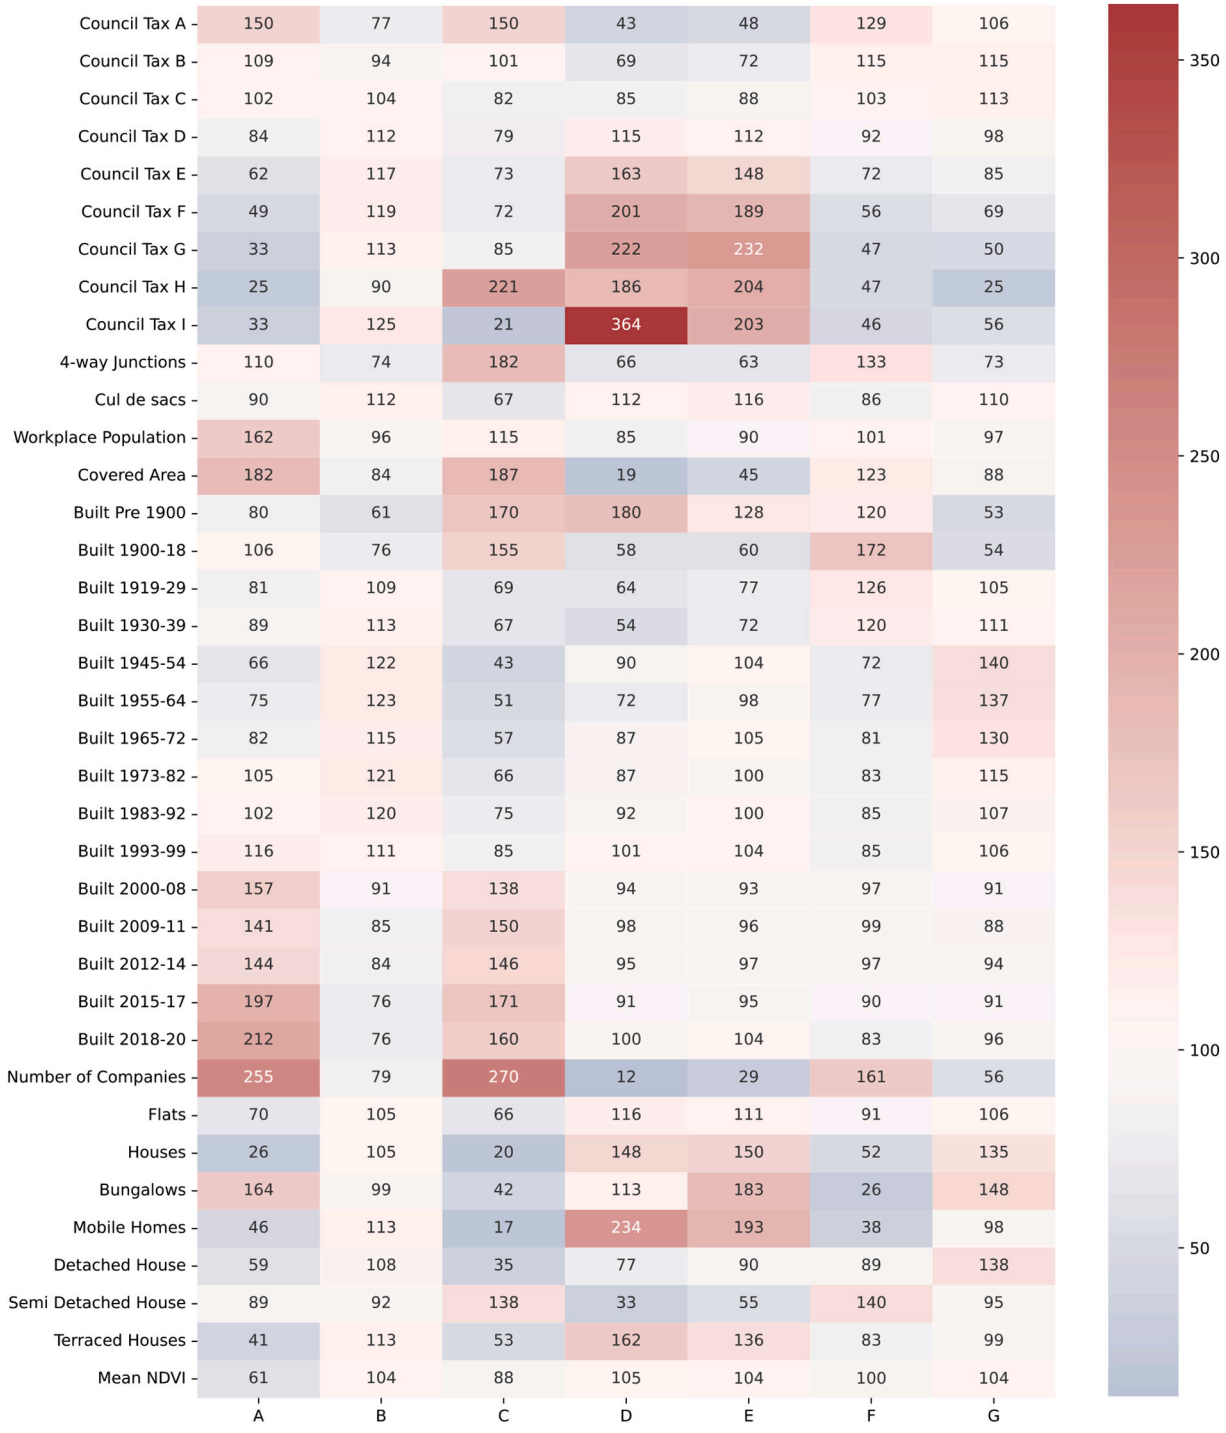

Fig. 10. Heatmap of index scores describing the clusters.

Taubenböck, & Zhu, 2018; Palmer et al., 2021; Zhao, Pang, & Wang, 2018).

Unlike many applications of CNNs, our approach here is unsupervised. This means that instead of training a network to predict a set of predefined labels or values, we use it to parse through the set of images and derive measures that aim to represent their latent structure. A convolutional autoencoder (CAE) is a type of CNN that performs this process. The CAE architecture is composed of two stages. The first is *encoding*, where the underlying structure of an image is learnt and stored in a latent vector of lower dimensionality to the original image. Once the latent representation is learnt, a second phase, the *decoding*, aims to reconstruct an image with similar characteristics to the input from the latent vector. For our application, the input tensor for each postcode comprised 16x16x4 (1024) pixel values which correspond to the 160 m (10 m resolution) square buffer around each postcode for the red, green, blue and near infrared bands of the satellite data.

For the later purpose of creating a classification, we only required the latent representation of the data: this *encoding* phase is summarised in the left part of Fig. 4 where the dimensionality reduction at each convolution is presented (f_e). However, as validation of this encoding, the decoding phase (f_d) was also implemented under the assumption that more effective encoding would return more accurate decoding. As such, when building our model, 20% of the original source images were held back as a validation sample, and the model converged to 88.74%

Table 2
Cluster descriptions.

Label	Description
A: Inner City and Town Centres	Predominantly found covering many central areas of cities and towns; this Group characterises many highstreets and small commercial premises supporting large working day populations. Coalescing with commercial land uses, residential accommodation is mostly flats. This Group is also observed in some less central commercial parks / zones; and dense terraced residential houses proximal to these areas with very limited greenspace.
B: Spacious Residential Properties	Areas characterised by this Group are predominantly residential and often contain properties at a low density, with gardens, wide streets and often surrounded by greenspace. Typically, this Group is found within larger developments on suburban fringes and in villages, and mostly containing property that have been built since the 1940s, although less prevalently since the 2000s.
C: High Density Urban Core	This cluster tends to be found within the built-up urban core of many large towns and cities but also characterises areas of out-of-town business, retail or industrial zones. Coalescing with structures supporting a large workday population, residential accommodation is predominantly in flats, but also includes areas of dense traditional terraced houses.
D: Rural Residences	Properties within this group tend to be residential and surrounded by greenspace. These areas are characterised by clusters of larger adjoining or traditional cottages, with ages dating from before 1900.
E: Rural Properties and Farms	This Group are mostly found within rural or more isolated locations and are typically characterised by larger residential properties or farms. As might be expected, properties tend to be surrounded by greenspace; however, many plots also encompass large paved agricultural yards.
F: Inner Suburban Townhouses and Large Terraces	Properties within these inner suburban locations tend to comprise semi-detached properties or larger terraces; with some properties sub divided into flats. Many of the property within these areas are pre-World War 2, some of which have gardens / outdoor space.
G: Outer Suburbs	This Group is mostly found towards the periphery of urban areas, consisting of mature residential neighbourhoods of mainly semi-detached properties. Many of the properties were built after World War 2 and until the 1970s and 80s.

validation accuracy after 100 training epochs. The decoded portion of the model was then discarded, and the output of the encoding phase of the CAE was retained. This comprised a set of learned latent representations of dimensions 8x8x1 for each area surrounding a postcode. We then use these as the input dataset for the next phase of the analysis with a record for every postcode and 64 separate attributes describing the context of their surrounding 160 m area.

Our hypothesis in this work was that the CAE would derive a useful characterisation of salient attributes of the context for each postcode; akin to how a principle component analysis (PCA) might be applied in the analysis of tabular data(Liu, Singleton, & Arribas-Bel, 2019) but relaxing the linearity of PCA to allow for complex, non-linear representations of the input images. In practical terms, each of the 64 dimensions output from this process equated to an abstract aspect of the learned structure. In isolation, these may form a useful input to predictive models aligned to some specific outcome. In our context however, given their very detailed nature, we require further summary to derive a useful aggregate indicator of comparative and generalisable urban context.

3. Results: representing salient context through cluster analysis

As introduced earlier, tabular approaches to the representation of urban context will often take a set of input measures for different social or built environment characteristics and apply a clustering algorithm to derive discrete labels describing salient characteristics. Following previous work in the creation of geodemographic and built environment classifications (Alexiou et al., 2016; Gale et al., 2016; Spielman & Singleton, 2015b), we applied a k-means clustering algorithm to the latent vectors produced by the CAE. K-means searched the feature space of its input data (1,710,715 × 64) to find those postcodes sharing similar characteristics. K-means was initiated using k random locations (seeds), to which all postcodes were matched to their nearest seed. Cluster mean values across the input attributes were then calculated and postcodes reassigned to their nearest cluster centroid. This process was repeated until no further assignments were made. Because k-means is initiated with a random seed, the results are stochastic; and in order to obtain an optimal fit, each k-means model was run 10,000 times, retaining the model that best fit the data, as measured by a total within sum of squares statistic. Despite the relative computational efficiency of k-means, the very high dimensionality of this problem, alongside the requirement for multiple runs required parallel computation, and as such was also completed within a GPU framework.

A general challenge for k-means is assigning an appropriate number of clusters (k) that represent yet uncovered structure contained within the input data, under the assumption that there is no theoretical rational for selecting a particular value of k. To explore a suitable value for k we implemented a Clustergram (Schonlau, 2002). This presents different potential k values by plotting the weighted mean of the first component of a principal component analysis (PCA) for each individual cluster. Given that the first component of a PCA provides measures the majority of variance contained within the input data, then separation of the clusters along the Y axis gives an indicatior of their difference. Such charts also highlight how clusters are spit by moving to different values of k; with the width of the lines representing the overall size of the cluster for each solution. A Clustergram is shown in Fig. 5 and was created to explore potential values of k in the range 2 to 19. From this analysis a k = 7 solution was chosen which demonstrated clear separation between the clusters.

A k = 7 model was then run for 1000 iterations of k means to explore the saliency of this solution. After mapping the clusters, it became apparent how sensitive the CAE model was at picking up differences between the input images. On balance this was sub optimal both in terms of the utility of the classification created, but also as a representation of how context might be experienced. An example of the issue is illustrated in Fig. 6 for an area of South Liverpool, UK, which plots a set of residential addresses (unique property reference numbers: UPRN) ascribed with a colour corresponding to their associated postcode cluster (multiple UPRN per postcode). Many of the clusters are not spatially contiguous; and although UPRN are shown here to better delineate street morphology; the input postcodes and their associated buffers are coarser. Additionally, the satellite imagery presented alongside these data to provide context here is very high resolution (25 cm) relative to the open data used to create the classification (10 m).

Mitigation of such artefacts were implemented by drawing strength from the spatial structure of the underlying postcode geography. A smoothed average for each of the 64 input measures was calculated for every postcode, comprising values from those postcodes found within each 160 m buffer. Empirically, this created a new set of spatially lagged input measures following the distribution of postcode geography, and conceptually produces a more continuous rather than discretely delineated representation of context. After fitting an additional Clustergram on the spatially smoothed input data, 7 clusters were again identified as suitable partitioning, with the output mapped for the same example area in Fig. 6 B. This highlights how the additional smoothing process returned more contiguous clusters, without some of the artifact effects

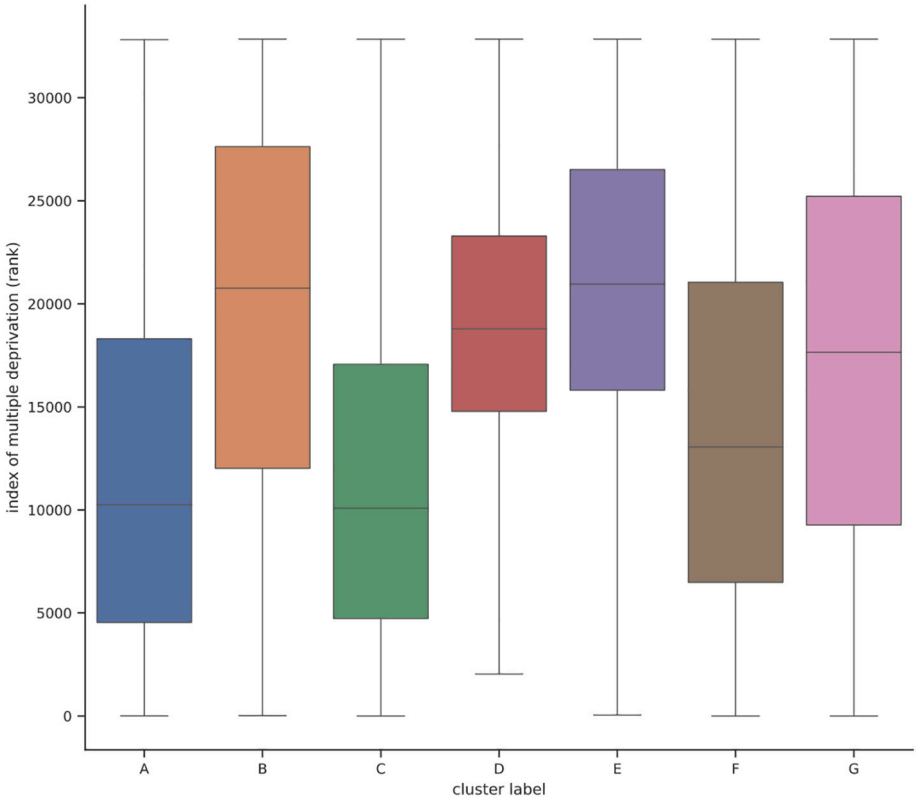

Fig. 11. Distribution of the 2019 English Index of Multiple Deprivation rank by cluster.

previously observed. More generally, better overall performance could be measured when the fit of the two different clustering solutions to their respective data inputs were calculated through a total within sum of squares statistic. The input attributes for each postcode were compared to their assigned cluster mean, and the squared difference between these scores derived. A lower within sum of squares (WSS) score equates to a postcode that has a better fit to the assigned cluster, and can be summed to give an overall measure of fit for the presented solution. This returned a total WSS score of 2,049,026 for the raw CAE input and 739,162 for the model with spatial smoothing; representing a large improvement, and overall much more compact clusters. Following these results; a final version of the smoothed model was returned by running k means 10,000 times, extracting the result with the lowest total within sum of squares.

Before examining the characteristics of the identified clusters their internal robustness was assessed for systematic differences in cluster quality. For the final optimised cluster run, WSS were calculated for each postcode and presented at the cluster level and for postcode districts.[3] Fig. 7 presents a box and whisker plot of the WSS scores by their associated cluster. As noted earlier, the addition of smoothing resulted in lower total WSS, however in the optimised result there are differences in fit between clusters, with A, D and E all showing greater error. Such errors will be contextualised in the next section after describing the cluster characteristics, however by mapping the errors (Fig. 8) we find clear spatial variations, with lower error exhibited in urban versus rural areas. There are two likely reasons why this effect is observed: the first is that there are more postcodes found within urban areas, and as such the CAE is learning structure from a larger amount of data derived within these locations; and secondly, given that spatial smoothing improved the overall classification fit, and given that the density of postcodes is higher in urban versus rural areas, this is also likely to have had greater impact within these contexts.

4. Case study: a national classification of context

Although our final classification covers the full extent of Great Britain; to illustrate the overarching geography of the identified patterns; a subset of the full output is shown in Fig. 9 for the Liverpool City Region, which is in the Northwest of England and comprises six Local Authority areas. The postcode clusters are labelled A-G, and the colours delineate these different classes. At this scale differentiation between primarily residential and commercial areas (Clusters A & C) are visible; and differences within urban and more rural parts of the city region can be observed, as also aligned with the changing density of points.

Because the CAE measures forming input to the classification were abstract derivations from imagery, these were of limited utility when describing salient features of the clusters. As such, a range of ancillary data were appended to the postcodes; the sources of which are summarised in Table 1. In addition to assisting with description of each cluster's salient characteristics, this process had two additional benefits. The first is external validation, in that the clusters should differentiate logically between attributes of the built environment; and secondly, as an indicator of those types of contextual features that are well discriminated through the process of applying the CAE to satellite data of the specified resolution.

For the attributes extracted from the identified ancillary data we computed index scores that compare either the proportion or rate within a cluster to the national average. These scores are standardized so that 100 equates to the national average, 50 would be a half and 200 double. A tabular representation of the scores, shaded by their relative propensity are shown in Fig. 10. These scores in addition to maps of the clusters were used to generate Group labels and descriptions that are presented in Table 2.

[3] A postcode district represents the first part of the postcode before the space – for example, for the full postcode - PO7 8NJ, the postcode district would be PO7.

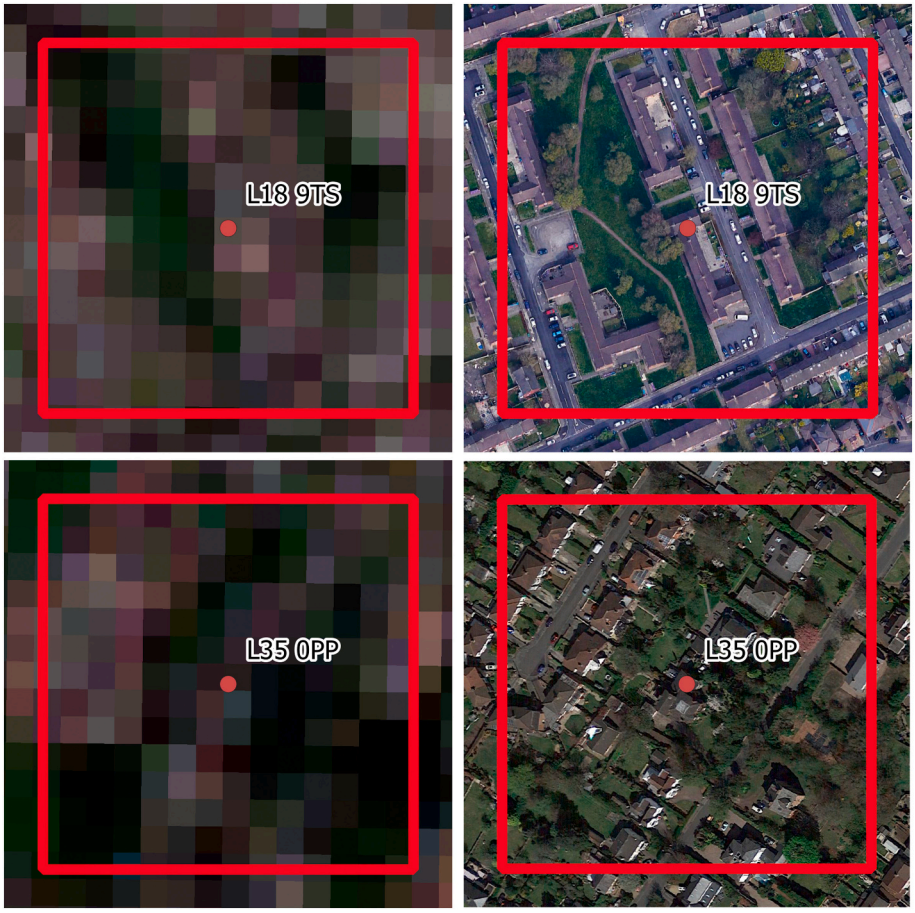

Fig. 12. Two Postcodes that were both asigned into Cluster B however from different deprivation profiles. The 160 m contextual buffers are shown over ESA Sentinel 2 (left) and Google satellite imagery (right). Modified Copernicus Sentinel data 2021.
Map data ©2021 Google, ©2021 Bluesky, Bluesky, Infoterra Ltd. & COWI A/S, Getmapping plc, Infoterra Ltd. & Bluesky, Maxar Technologies, The Geo-Information Group, Map data ©2021.

The results present a range of salient attributes found within each cluster; illustrating the potential utility of CAE when building classifications aiming to differentiate context. Out of all the Groups presented, perhaps the least well differentiated were locations defined as "D: Rural Residences" and "E: Rural Properties and Farms". These had quite similar profiles based on the ascribed attributes, however there were differences when mapped; notably "D: Rural Residences" tended to contain buildings that were more clustered; and "E: Rural Properties and Farms" also included more farm locations. For simplicity we have left these clusters as presented, but a more parsimonious classification might be returned if such clusters were manually merged.

A further observation on the overall shape of the classification is that aspects of morphometric structure are differentiated by the identified clusters; notably 4-way junctions and cul-de-sacs. Although such measures likely correlate to other features of context identified by the CAE, it is none the less encouraging to see that these differences were observed; and we would argue that this indicates potential for higher resolution imagery revealing richer measures of street density / complexity.

Finally, in framing the utility of the created classification, and the general approach of using CAE for the derivation of contextual measures, it is also important to explore their limits. The output of the CAE is framed here as our best attempt at capturing the characteristics of postcode context as represented by the input image taken from the specified satellite data. As outlined earlier in this section, the approach maps reasonably successfully onto observed structures that differentiate places. However, in more general tabular approaches to the analysis of urban context such as geodemographic classification, fuller descriptions tend to be provided through the inclusion of a wider array of variables describing both the physical structure of places and the socio-economic characteristics of their residents. Although outside the scope and objectives of this work, elsewhere we have demonstrated where remotely sensed data have shown utility in differentiating patterns or urban socioeconomic structure (Arribas-Bel et al., 2017). Although such work takes an alternative methodological approach, it is not clear that our windowed approach as presented here is well suited to such purposes; and as such we would argue for caution in drawing associations beyond the scope of what can be identified from the data. For example, Fig. 11 plots the rank distribution of postcodes within England by the 2019 Index of Multiple Deprivation (1 = the most deprived). Although there are some differences, such as more rural clusters tending to be less deprived than those in urban areas, there is much overlap. Despite the objective of the approach used to generate the IMD being divergent from the work presented in this paper, if the same comparative analysis was completed with more traditional geodemographic classification, it would be possible to see greater differentiation in such patterns, as these typically include socio-economic measures.

Such lack of differentiation on this particular domain may very likely be a constraint of the data resolution used for the classification. This is illustrated in Fig. 12 by examining two postcodes found with values very close to the mean of Cluster 1 within the Liverpool City Region: L18 9TS is more deprived and found in the 3rd decile of the IMD, whereas L35 0PP is more affluent and ranked in the 8th decile. The built structure of these areas are quite different, which can be seen on the right column images, presenting the higher resolution remotely sensed data. L18 9TS are local authority flats within an open greenspace plot, and L35 0PP are Semi-Detached properties with private gardens. These are quite different areas in terms of their socio-economic circumstances, but their grouping into a single cluster is more understandable (as more visibly similar) when true colour composites generated from the RGB source data are compared (left column). We draw this comparison to illustrate

the limits of what the presented work can show, and also the likely potential for greater differentiation of social attributes if higher resolution data were available and where these have association with characteristics of the built environment.

5. Conclusions and directions for future research

In this paper we have introduced a new and novel method to extract local contextual measures from Sentinel 2 satellite data following the distribution of postcodes in Great Britain. Our methodology implemented a CAE to produce a data-driven encoding of the context of 160 m buffers around each unit postcode. This equated to 1,710,715 different image inputs and the massive computational demands necessary to fit the CAE model to these data required extensive use of GPU architecture. Measures were extracted after the encoding phase of the CAE to reduce the source data into 64 separate features representing different encodings of the context of each image. The utility of these measures was then demonstrated, through their input into a cluster analysis with the objective of identifying those areas sharing similar salient contextual characteristics. The final output of the model was 7 clusters, which were mapped and described through a range of independent ancillary data. Our produced representations complement alternative morphometric, tabular and LULC approaches to the representation of urban form and context.

Even with input data at the resolution of Sentinel 2 (10 m / pixel), the autoencoder process through CNN was illustrated to be extremely sensitive; requiring tspatial smoothing of the raw outputs to limit unhelpful heterogeneity. Some of this variability related to how context was conceptualised in this study: notably through a 160 m square buffer around each postcode. This zone was fixed for all postcodes as a requirement of the CNN for a Tensorset comprising of tensors of uniform size and shape. Furthermore, although the zone chosen for this study might be contracted to reduce the contextual extent, through experimentation we found that this had quite negative effects on the model performance as there were much less data upon which the CNN could learn structure. We argue that this presents an interesting case of the Modifiable Areal Unit Problem as the success of the analytical technique itself are directly related to the zoning and scale of the input data used, not just the representations presented. Although outside the scope of this study or the potential of those input data utilised, it would be possible with higher resolution data to explore these effects systematically. Moreover, the CAE demonstrated that it could learn structure from the data that are presented, which in this application was a very large corpus of images for the national extent. As with many machine learning algorithms, CAEs are able to learn more accurate representations of underlying structure when presented with larger volumes of data, assuming all other things are constant, such as the quality of the data. As such, in future work, we argue that there is also a need to explore the extent to which overall model performance might be impacted by a more constrained geographic extent and reduced data input, and how such effects might also interact when also presented with different resolutions of input data. Such issues would also warrant investigation in association with a need for spatial weighting as was applied here, or whether other types of weighing schemes or approaches impact model performance.

Overall we are very encouraged by the results of this study, as the CAE outputs demonstrated clear utility for the differentiation of contexts when these were integrated into a cluster analysis. The output of the CAE themselves are however data driven and so need further ancillary data to make sense of these representations and ascribe their limits. The example of socio-economic differentiation was presented to illustrate this point; but also highlights potential for the outputs of the model presented here to be combined with other available data to create richer profiles of areas; perhaps as input to fuller and more comprehensive geodemographic classification.

Funding

This research was funded by the ESRC – grant numbers ES/L011840/1 and ES/T005238/1.

CRediT authorship contribution statement

Alex Singleton: Conceptualization, Funding acquisition, Project administration, Methodology, Validation, Resources, Supervision, Writing – original draft, Writing – review & editing. **Dani Arribas-Bel:** Writing – review & editing, Supervision. **John Murray:** Conceptualization, Data curation, Formal analysis, Software, Writing – review & editing. **Martin Fleischmann:** Software, Validation, Visualization, Writing – review & editing.

References

Alexiou, A., Singleton, A. D., & Longley, P. A. (2016). A classification of multidimensional open data for urban morphology. *Built Environment*. https://doi.org/10.2148/benv.42.3.382

Alizadeh Kharazi, B., & Behzadan, A. H. (2021). Flood depth mapping in street photos with image processing and deep neural networks. *Computers, Environment and Urban Systems, 88*, Article 101628. https://doi.org/10.1016/j.compenvurbsys.2021.101628

Araldi, A., & Fusco, G. (2019). From the street to the metropolitan region: Pedestrian perspective in urban fabric analysis. *46*(7), 1243–1263. https://doi.org/10.1177/2399808319832612

Arribas-Bel, D., Patino, J. E., & Duque, J. C. (2017). Remote sensing-based measurement of living environment deprivation: Improving classical approaches with machine learning. *PLoS One*. https://doi.org/10.1371/journal.pone.0176684

Balarabe, A. T., & Jordanov, I. (2021). *LULC image classification with convolutional neural network* (pp. 5985–5988). https://doi.org/10.1109/IGARSS47720.2021.9555015

Batty, M. (2013). *The new science of cities*. MIT Press.

Bennett, M. M., & Smith, L. C. (2017). Advances in using multitemporal night-time lights satellite imagery to detect, estimate, and monitor socioeconomic dynamics. In *Remote Sensing of Environment*. https://doi.org/10.1016/j.rse.2017.01.005

Bhosle, K., & Musande, V. (2019). Evaluation of deep learning CNN model for land use land cover classification and crop identification using hyperspectral remote sensing images. *Journal of the Indian Society of Remote Sensing, 47*(11), 1949–1958. https://doi.org/10.1007/S12524-019-01041-2/TABLES/8

Blaschke, T. (2010). Object based image analysis for remote sensing. *ISPRS Journal of Photogrammetry and Remote Sensing, 65*(1), 2–16. https://doi.org/10.1016/J.ISPRSJPRS.2009.06.004

Boeing, G. (2018). A multi-scale analysis of 27,000 urban street networks: every US city, town, urbanized area, and Zillow neighborhood. *47*(4), 590–608. https://doi.org/10.1177/2399808318784595

Branson, S., Wegner, J. D., Hall, D., Lang, N., Schindler, K., & Perona, P. (2018). From Google Maps to a fine-grained catalog of street trees. *ISPRS Journal of Photogrammetry and Remote Sensing, 135*, 13–30. https://doi.org/10.1016/j.isprsjprs.2017.11.008

Campbell, A., Both, A., Sun, Q., & (Chayn).. (2019). Detecting and mapping traffic signs from Google Street View images using deep learning and GIS. *Computers, Environment and Urban Systems, 77*, Article 101350. https://doi.org/10.1016/j.compenvurbsys.2019.101350

Chen, Y., Ming, D., & Lv, X. (2019a). Superpixel based land cover classification of VHR satellite image combining multi-scale CNN and scale parameter estimation. *Earth Science Informatics, 12*(3), 341–363. https://doi.org/10.1007/S12145-019-00383-2

Chen, Y., Ming, D., & Lv, X. (2019b). Superpixel based land cover classification of VHR satellite image combining multi-scale CNN and scale parameter estimation. *Earth Science Informatics, 12*(3), 341–363. https://doi.org/10.1007/S12145-019-00383-2/FIGURES/19

Cheng, G., Han, J., & Lu, X. (2017). Remote sensing image scene classification: Benchmark and state of the art. *Proceedings of the IEEE, 105*(10), 1865–1883. https://doi.org/10.1109/JPROC.2017.2675998

Comber, S., Arribas-Bel, D., Singleton, A., & Dolega, L. (2020). Using convolutional autoencoders to extract visual features of leisure and retail environments. *Landscape and Urban Planning, 202*, Article 103887. https://doi.org/10.1016/j.landurbplan.2020.103887

Dibble, J., Prelorendjos, A., Romice, O., Zanella, M., Strano, E., Pagel, M., & Porta, S. (2017). On the origin of spaces: Morphometric foundations of urban form evolution. *46*(4), 707–730. https://doi.org/10.1177/2399808317725075

Esch, T., Asamer, H., Bachofer, F., Balhar, J., Boettcher, M., Boissier, E., … Zeidler, J. N. (2020). Digital world meets urban planet–new prospects for evidence-based urban studies arising from joint exploitation of big earth data, information technology and shared knowledge. *International Journal of Digital Earth, 13*(1), 136–157. https://doi.org/10.1080/17538947.2018.1548655

Fan, R., Feng, R., Wang, L., Yan, J., & Zhang, X. (2020). Semi-MCNN: A semisupervised multi-CNN ensemble learning method for urban land cover classification using submeter HRRS images. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing, 13*, 4973–4987. https://doi.org/10.1109/JSTARS.2020.3019410

Feng, Q., Yang, J., Zhu, D., Liu, J., Guo, H., Bayartungalag, B., & Li, B. (2019). Integrating multitemporal Sentinel-1/2 data for coastal land cover classification using a multibranch convolutional neural network: A case of the Yellow River Delta. *Remote Sensing, 11*(9), 1006. https://doi.org/10.3390/RS11091006

Fleischmann, M., Feliciotti, A., Romice, O., & Porta, S. (2021). Methodological foundation of a numerical taxonomy of urban form. *0*(0), 1–17. https://doi.org/10.1177/23998083211059835

Forrest, R., & Kearns, A. (2001). Social cohesion, social capital and the neighbourhood. *Urban Studies, 38*(12), 2125–2143. https://doi.org/10.1080/00420980120087081

Gaetano, R., Ienco, D., Ose, K., & Cresson, R. (2018). A two-branch CNN architecture for land cover classification of PAN and MS imagery. *Remote Sensing, 10*(11), 1746. https://doi.org/10.3390/RS10111746

Gale, C. G., Singleton, A. D., Bates, A. G., & Longley, P. A. (2016). Creating the 2011 area classification for output areas (2011 OAC). *Journal of Spatial Information Science, 12*. https://doi.org/10.5311/JOSIS.2016.12.232

Galster, G. (2001). On the nature of neighbourhood. *Urban Studies, 38*(12), 2111–2124. https://doi.org/10.1080/00420980120087072

Galster, G. C. (2013). The mechanism(s) of neighbourhood effects: Theory, evidence, and policy implications. In *, Vol. 9789400723092. Neighbourhood effects research: New perspectives* (pp. 23–56). Netherlands: Springer. https://doi.org/10.1007/978-94-007-2309-2_2.

van Ham, M., Manley, D., Bailey, N., Simpson, L., & Maclennan, D. (2012). Neighbourhood effects research: New perspectives. In *, Vol. 9789400723092. Neighbourhood effects research: New perspectives* (pp. 1–21). Netherlands: Springer. https://doi.org/10.1007/978-94-007-2309-2_1.

Helber, P., Bischke, B., Dengel, A., & Borth, D. (2018). Introducing eurosat: A novel dataset and deep learning benchmark for land use and land cover classification. In *International Geoscience and Remote Sensing Symposium (IGARSS), 2018-July* (pp. 204–207). https://doi.org/10.1109/IGARSS.2018.8519248

Heryadi, Y., Miranda, E., Heryadi, Y., & Miranda, E. (2019). Land cover classification based on Sentinel-2 satellite imagery using convolutional neural network model: A case study in Semarang area, Indonesia. *Studies in Computational Intelligence, 830*, 191–206. https://doi.org/10.1007/978-3-030-14132-5_15

Hijazi, I., Li, X., Koenig, R., Schmit, G., el Meouche, R., Lv, Z., & Abune'meh, M. (2016). Measuring the homogeneity of urban fabric using 2D geometry data. *44*(6), 1097–1121. https://doi.org/10.1177/0265813516659070

Hillier, B., Leaman, A., Stansall, P., & Bedford, M. (2016). Space syntax. *3*(1990). https://doi.org/10.1068/B030147, 056–01.

Ibrahim, M. R., Haworth, J., & Cheng, T. (2020). Understanding cities with machine eyes: A review of deep computer vision in urban analytics. *Cities, 96*, Article 102481. https://doi.org/10.1016/j.cities.2019.102481

Imhoff, L., Lawrence, W. T., Stutzer, D. C., & Elvidge, C. D. (1997). A technique for using composite DMSP/OLS "City Lights" satellite data to map urban area. *Remote Sensing of Environment, 61*(3), 361–370. https://doi.org/10.1016/S0034-4257(97)00046-1

Jean, N., Burke, M., Xie, M., Davis, W. M., Lobell, D. B., & Ermon, S. (2016). Combining satellite imagery and machine learning to predict poverty. *Science*. https://doi.org/10.1126/science.aaf7894

Kang, J., Körner, M., Wang, Y., Taubenböck, H., & Zhu, X. X. (2018). Building instance classification using street view images. *ISPRS Journal of Photogrammetry and Remote Sensing, 145*, 44–59. https://doi.org/10.1016/j.isprsjprs.2018.02.006

Kareem, R. S. A., Ramanjineyulu, A. G., Rajan, R., Setiawan, R., Sharma, D. K., Gupta, M. K., … Sengan, S. (2021). Multilabel land cover aerial image classification using convolutional neural networks. *Arabian Journal of Geosciences, 14*(17), 1–18. https://doi.org/10.1007/S12517-021-07791-Z/FIGURES/9

Karney, C. F. F. (2011). Transverse Mercator with an accuracy of a few nanometers. *Journal of Geodesy*. https://doi.org/10.1007/s00190-011-0445-3

Kotaridis, I., & Lazaridou, M. (2021). Remote sensing image segmentation advances: A meta-analysis. *ISPRS Journal of Photogrammetry and Remote Sensing, 173*, 309–322. https://doi.org/10.1016/J.ISPRSJPRS.2021.01.020

Law, S., Seresinhe, C. I., Shen, Y., & Gutierrez-Roig, M. (2020). Street-frontage-net: Urban image classification using deep convolutional neural networks. *International Journal of Geographical Information Science, 34*(4), 681–707. https://doi.org/10.1080/13658816.2018.1555832

Lecun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature, 521*(7553), 436–444. https://doi.org/10.1038/nature14539

Liao, J., Cao, J., Wang, K., & Xu, Z. (2022). Land cover classification from very high spatial resolution images via multiscale object-driven CNNs and automatic annotation. *16*(1), 014513. https://doi.org/10.1117/1.JRS.16.014513

Liu, L., Silva, E. A., Wu, C., & Wang, H. (2017). A machine learning-based method for the large-scale evaluation of the qualities of the urban environment. *Computers, Environment and Urban Systems, 65*, 113–125. https://doi.org/10.1016/j.compenvurbsys.2017.06.003

Liu, Y., Singleton, A. D., & Arribas-Bel, D. (2019). A principal component analysis (PCA)-based framework for automated variable selection in geodemographic classification. *Geo-Spatial Information Science, 22*(4), 251–264. https://doi.org/10.1080/10095020.2019.1621549

Lloyd, C. T., Chamberlain, H., Kerr, D., Yetman, G., Pistolesi, L., Stevens, F. R., … Tatem, A. J. (2019). Global spatio-temporally harmonised datasets for producing high-resolution gridded population distribution datasets. *Big Earth Data*. https://doi.org/10.1080/20964471.2019.1625151

Mahabir, R., Croitoru, A., Crooks, A., Agouris, P., & Stefanidis, A. (2018). A critical review of high and very high-resolution remote sensing approaches for detecting and mapping slums: Trends, challenges and emerging opportunities. *Urban Science*. https://doi.org/10.3390/urbansci2010008

Mesev, T. V., Longley, P. A., Batty, M., & Xie, Y. (1995). Morphology from imagery: Detecting and measuring the density of urban land use. *Environment & Planning A, 27*(5), 759–780. https://doi.org/10.1068/a270759

Middel, A., Lukasczyk, J., Zakrzewski, S., Arnold, M., & Maciejewski, R. (2019). Urban form and composition of street canyons: A human-centric big data and deep learning approach. *Landscape and Urban Planning, 183*, 122–132. https://doi.org/10.1016/j.landurbplan.2018.12.001

Morton, D., Marston, C. G., O'Neil, A. W., & Rowland, C. S. (2020). *Land cover map 2017 (land parcels, N. Ireland)*. NERC Environmental Information Data Centre. https://doi.org/10.5285/efb98222-5b9a-4d56-990d-5ab85eaf187e

Nogueira, K., Penatti, O. A. B., & dos Santos, J. A. (2017). Towards better exploiting convolutional neural networks for remote sensing scene classification. *Pattern Recognition, 61*, 539–556. https://doi.org/10.1016/J.PATCOG.2016.07.001

Palmer, G., Green, M., Boyland, E., Vasconcelos, Y. S. R., Savani, R., & Singleton, A. (2021). A deep learning approach to identify unhealthy advertisements in street view images. *Scientific Reports, 11*(1), 4884. https://doi.org/10.1038/s41598-021-84572-4

Pokhriyal, N., & Jacques, D. C. (2017). Combining disparate data sources for improved poverty prediction and mapping. *Proceedings of the National Academy of Sciences of the United States of America*. https://doi.org/10.1073/pnas.1700319114

Proville, J., Zavala-Araiza, D., & Wagner, G. (2017). Night-time lights: A global, long term look at links to socio-economic trends. *PLoS One*. https://doi.org/10.1371/journal.pone.0174610

Saeedimoghaddam, M., & Stepinski, T. F. (2020). Automatic extraction of road intersection points from USGS historical map series using deep convolutional neural networks. *International Journal of Geographical Information Science, 34*(5), 947–968. https://doi.org/10.1080/13658816.2019.1696968

Salvadori, L., Badas, M. G., di Bernardino, A., Querzoli, G., & Ferrari, S. (2021). A street graph-based morphometric characterization of two large urban areas. *Sustainability, 13*(3), 1025. https://doi.org/10.3390/SU13031025

Sampson, R. J. (2019). Neighbourhood effects and beyond: Explaining the paradoxes of inequality in the changing American metropolis. *Urban Studies, 56*(1), 3–32. https://doi.org/10.1177/0042098018795363

Schonlau, M. (2002). The clustergram: A graph for visualizing hierarchical and nonhierarchical cluster analyses. *The Stata Journal: Promoting Communications on Statistics and Stata*. https://doi.org/10.1177/1536867x0200200405

Shendryk, Y., Rist, Y., Ticehurst, C., & Thorburn, P. (2019). Deep learning for multi-modal classification of cloud, shadow and land cover scenes in PlanetScope and Sentinel-2 imagery. *ISPRS Journal of Photogrammetry and Remote Sensing, 157*, 124–136. https://doi.org/10.1016/J.ISPRSJPRS.2019.08.018

Singleton, A. D., Alexiou, A., & Savani, R. (2020). Mapping the geodemographics of digital inequality in Great Britain: An integration of machine learning into small area estimation. *Computers, Environment and Urban Systems*. https://doi.org/10.1016/j.compenvurbsys.2020.101486

Singleton, A. D., & Arribas-Bel, D. (2019). Geographic data science. *Geographical Analysis*. https://doi.org/10.1111/gean.12194

Singleton, A. D., & Spielman, S. E. (2014). The past, present, and future of geodemographic research in the United States and United Kingdom. *The Professional Geographer, 66*(4), 558–567.

Spielman, S. E., & Singleton, A. D. (2015a). Studying neighborhoods using uncertain data from the American community survey: A contextual approach. *Annals of the Association of American Geographers, 105*(5), 1003–1025.

Spielman, S. E., & Singleton, A. D. (2015b). Studying neighborhoods using uncertain data from the American community survey: A contextual approach. *Annals of the Association of American Geographers, 105*(5), 1003–1025. https://doi.org/10.1080/00045608.2015.1052335

Steele, J. E., Sundsøy, P. R., Pezzulo, C., Alegana, V. A., Bird, T. J., Blumenstock, J., … Bengtsson, L. (2017). Mapping poverty using mobile phone and satellite data. *Journal of the Royal Society, Interface*. https://doi.org/10.1098/rsif.2016.0690

Vanderhaegen, S., & Canters, F. (2017). Mapping urban form and function at city block level using spatial metrics. *Landscape and Urban Planning, 167*, 399–409. https://doi.org/10.1016/J.LANDURBPLAN.2017.05.023

Venerandi, A., Zanella, M., Romice, O., Dibble, J., & Porta, S. (2016). Form and urban change – An urban morphometric study of five gentrified neighbourhoods in London. *44*(6), 1056–1076. https://doi.org/10.1177/0265813516658031

Webber, R., & Burrows, R. (2018). *The predictive postcode: The geodemographic classification of British society* (SAGE).

Wei, R., Rey, S., & Knaap, E. (2021). Efficient regionalization for spatially explicit neighborhood delineation. *International Journal of Geographical Information Science, 35*(1), 135–151. https://doi.org/10.1080/13658816.2020.1759806

Weng, Q., Mao, Z., Lin, J., & Guo, W. (2017). Land-use classification via extreme learning classifier based on deep convolutional features. *IEEE Geoscience and Remote Sensing Letters, 14*(5), 704–708. https://doi.org/10.1109/LGRS.2017.2672643

Weng, Q., Quattrochi, D., & Gamba, P. E. (2018). *Urban Remote Sensing*. CRC Press.

Wurm, M., & Taubenböck, H. (2018). Detecting social groups from space – Assessment of remote sensing-based mapped morphological slums using income data. *Remote Sensing Letters*. https://doi.org/10.1080/2150704X.2017.1384586

Yang, C., & Gidófalvi, G. (2020). Detecting regional dominant movement patterns in trajectory data with a convolutional neural network. *International Journal of Geographical Information Science, 34*(5), 996–1021. https://doi.org/10.1080/13658816.2019.1700510

Zhang, C., Sargent, I., Pan, X., Li, H., Gardiner, A., Hare, J., & Atkinson, P. M. (2018). An object-based convolutional neural network (OCNN) for urban land use classification.

Remote Sensing of Environment, 216, 57–70. https://doi.org/10.1016/J.RSE.2018.06.034

Zhao, R., Pang, M., & Wang, J. (2018). Classifying airborne LiDAR point clouds via deep features learned by a multi-scale convolutional neural network. *International Journal of Geographical Information Science, 32*(5), 960–979. https://doi.org/10.1080/13658816.2018.1431840

Zhu, Z., Zhou, Y., Seto, K. C., Stokes, E. C., Deng, C., Pickett, S. T. A., & Taubenböck, H. (2019). Understanding an urbanizing planet: Strategic directions for remote sensing. *Remote Sensing of Environment, 228*, 164–182. https://doi.org/10.1016/j.rse.2019.04.020