Jane Jacobs in the Sky: Predicting Urban Vitality with Open Satellite Data

SANJA ŠĆEPANOVIĆ and SAGAR JOGLEKAR, Nokia Bell Labs, U.K.

STEPHEN LAW, University College London & The Alan Turing Institute, U.K.

DANIELE QUERCIA, Nokia Bell Labs, U.K.

The presence of people in an urban area throughout the day – often called 'urban vitality' – is one of the qualities world-class cities aspire to the most, yet it is one of the hardest to achieve. Back in the 1970s, Jane Jacobs theorized urban vitality and found that there are four conditions required for the promotion of life in cities: diversity of land use, small block sizes, the mix of economic activities, and concentration of people. To build proxies for those four conditions and ultimately test Jane Jacobs's theory at scale, researchers have had to collect both private and public data from a variety of sources, and that took decades. Here we propose the use of one single source of data, which happens to be publicly available: Sentinel-2 satellite imagery. In particular, since the first two conditions (diversity of land use and small block sizes) are visible to the naked eye from satellite imagery, we tested whether we could automatically extract them with a state-of-the-art deep-learning framework and whether, in the end, the extracted features could predict vitality. In six Italian cities for which we had call data records, we found that our framework is able to explain on average 55% of the variance in urban vitality extracted from those records.

CCS Concepts: • **Applied computing** → *Computers in other domains*; *Law, social and behavioral sciences*; • **Computing methodologies** → **Image representations**; • **Information systems** → **Geographic information systems**.

Additional Key Words and Phrases: satellite imagery, urban vitality, AI, computer vision, Sentinel

1 INTRODUCTION

A livable city is a place full of life, and that life is created by city dwellers [37, 42]. In her 1961 book "The Death and Life of Great American Cities", writer and activist Jane Jacobs identified the four "generators" of urban vitality (i.e., pedestrian activity throughout the day): diversity of land use, small block sizes, a mix of economic activities, and the concentration of people. Without them, a city will die. With them, it will thrive. Her ideas exerted a tremendous impact on the thinking of architects, developers, urban planners, and community activists. Despite their importance and popularity, for a long time, Jacobs's theories could not be tested at scale. Upon collection of both private and public data from a variety of sources, the hypothesized effectiveness of the four generators has

Authors' addresses: Sanja Šćepanović, sanja.scepanovic@nokia-bell-labs.com; Sagar Joglekar, sagar.joglekar@nokia-bell-labs.com, Nokia Bell Labs, Cambridge, U.K.; Stephen Law, slaw@turing.ac.uk, University College London & The Alan Turing Institute, London, U.K.; Daniele Quercia, Nokia Bell Labs, Cambridge, U.K..
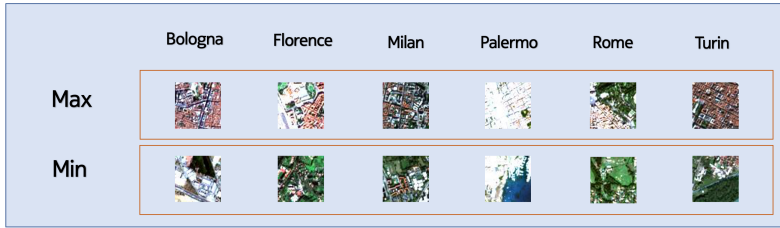
Fig. 1. Examples of satellite views for the areas with the highest and the lowest vitality levels in each of the six Italian cities. Despite having comparable vitality levels, the shown areas present very different urban forms.

been only recently tested in the city of Seoul, Korea [56], and across six Italian cities [14]. Yet, all these studies meticulously collected data from a variety of sources, and that took years (it took a decade in Seoul's case), considerably limiting research advancements.

To tackle the issue of data collection, we investigated whether it is possible to estimate some of the four generators of urban vitality (and vitality itself) from one single source: Sentinel-2 satellite imagery [58]. Since two of the four generators—that is, diversity of land use, and small block sizes—are visible to the naked eye, we hypothesized that they could be identified from satellite imagery. We also studied whether vitality could be directly predicted from such imagery. For both prediction tasks—predicting vitality generators first and then vitality indirectly from them, or predicting vitality directly—it is not obvious that neither of them could be performed upon satellite data. To see why, consider Figure 1: representative satellite views from the six Italian cities exhibit remarkably different visual features (such as colors and diversity of urban and natural forms) even if they come from areas of comparable vitality levels.

We built a processing framework that consists of three modules, each of which: *i)* extract small images from satellite imagery (which we call *imagelets*); *ii)* extract visual features from these imagelets with state-of-the-art deep learning methods; and *iii)* combine these features into district-level feature vectors. In so doing, we were able to make two main contributions:

C1 We built two techniques based on a deep learning framework, which, upon publicly available Sentinel-2 satellite imagery, extract representative *feature vectors for urban districts* (Section 3).

C2 We showed that, from these district feature vectors, we could predict proxies for *land use* and *small block size*, as well as predict *urban vitality* directly (Section 4). In a 5-fold cross-validation experiment, we found that satellite features explain, on average across six Italian cities, more than 55% of the variance (in terms of the adjusted coefficient of determination, R^2_{adj}) for urban vitality. Moreover, to ascertain how well the model generalizes, we performed a "leave-one-city" out validation (where the model is trained on a set of cities and tested on an unseen city). In this experiment, we found that the model can explain up to 50% of variance for Milan and 61% in Florence, suggesting the ability of our method to generalize to unseen cities. However, we also found that this performance drops for some cities - up to 25% for Rome. Upon a thorough investigation, we identified the potential reasons for the performance drop in Rome, which include the city's specific natural, cultural, and historical contexts. Accordingly, we discuss recommendations for future work (Section 5).

This work is at the intersection of two research areas—social computing and urban analytics. The combination of the two makes it possible to use emerging sources of data such as satellite

imagery to answer questions typical of the social sciences in a computational way, opening up new opportunities for future work (Section 5).

The project webpage is available at https://goodcitylife.org/vitality.

2 RELATED WORK

Three main lines of research are related to this study. First, social computing research on quantifying urban activity. Second, research that tested Jane Jacob's four generators of urban vitality at scale. Third, research on automatically inferring subtle urban properties from satellite imagery that are not necessarily visible to the naked eye (e.g., crime rates, property prices).

2.1 Measuring Urban Activity from Social Media Data

In the research community of social computing, researchers have worked on quantifying urban activity from different perspectives. They have, for example, crowd-sourced people's perceptions of beauty, quietness, and happiness of urban locations [46], and, by then combining these crowd-sourcing results with geo-referenced social media content (e.g., Flickr photos), they built path recommender systems that suggest beautiful, quiet, or happy urban routes in a city [44, 46, 60]. Le Falher et al. [35] used Foursquare check-ins to classify city neighborhoods in terms of prevalent activity types in them (e.g., dining, cultural, shopping). De Choudhury et al. [13] identified the movements of individual tourists based on their picture uploading patterns and found that average time spent at a location and location popularity were two useful predictors of interesting routes.

Researchers used openly available social media data to study urban dynamics. For example, Araujo et al. [3] recently used the Facebook marketing API to demographically profile an urban area. Redi et al. [47] predicted the ambiance of neighborhoods from the Flickr images in those neighborhoods. To augment deprivation indexes that have traditionally been derived from census data, and have been expensive to obtain, Venerandi et al. [61] used Foursquare and OpenStreetMap data. The premise behind all this work is that user-generated content on location-based services translates into rich behavioural traces [13, 19, 63]. More generally, in the social computing community, a need to empower cities, neighborhoods, and local communities with new technologies has long been recognized [12]. In a similar vein, our work develops a framework that uses openly available data to quantify the key aspect of urban vitality.

2.2 Verifying Empirically Jane Jacobs's Insights

The Household Travel Survey captures walking and driving activity in the whole city of Seoul and is conducted once every 5 years. Sung et al. [56] used the 2010 survey to operationalize urban vitality and showed that the four generators of urban vitality did hold across dongs, the small administrative areas of Seoul. In a similar setup, in six Italian cities, De Nadai et al. [14] have found that the four generators do apply to the Italian context as well. Instead of a costly multi-year survey, the researchers used mobile phone Internet density as a proxy for urban activity. Finally, beyond urban vitality, researchers have widely tapped into Jacobs's idea of *natural surveillance* to build predictive models for the crime: pedestrian activity acts as "natural surveillance", and that reduces crime rates [6, 7, 59].

2.3 Inferring Urban Variables from Satellite Imagery

Satellite images capture an overall structure of an area, which encodes different urban processes, as recent studies on satellite images have shown. For example, Han et al. [26] developed a deep learning architecture to map high-resolution satellite images across country districts into fixed-size district-level feature vectors. The authors then showed that such features can be used to predict socio-economic properties of districts, such as population density, population age distribution,

household count and size, and income per capita. Albert et al. [1] studied the patterns of land use in urban neighborhoods using high-resolution satellite imagery from Google Maps. To label the imagery, the authors resorted to the Urban Atlas, which provides land classification into 20 land use classes across 300 European cities. For a given location, they trained ResNet [27] and VGG-16 [52], two deep convolutional neural networks, to predict the most likely class for the 224×224 images. The achieved classification accuracy was between 0.7 and 0.8. Moreover, Albert et al. showed that the models can be trained on the imagery from one city to predict the classes in a completely different city. Law et al. [34] studied whether satellite/aerial imagery can enhance the prediction of housing prices. The authors obtained the high-resolution satellite/aerial imagery from Bing around the properties of interest and showed that the standard hedonic price approach for predicting housing prices can be improved by incorporating features extracted from the satellite imagery in a semi-interpretable manner.

Albert et al. [2] studied the distribution of key urban macroeconomic variables – population, luminosity (a proxy for energy access and use), and building density – from satellite data. Across all the cities with at least 10K inhabitants around the world, the authors demonstrated the strong link between the spatial distribution of lighting levels (which was previously shown to be a proxy for energy access and wealth levels [29]) and concentration of population. The work by Albert et al. [2] is similar to ours, in that, the authors studied population density and built infrastructure, both of which are proxies for urban vitality. However, the work by Albert et al. differs from ours in two main ways: *i)* there is no operationalization of the theoretical concept of urban vitality as this was not the focus; and *ii)* the authors studied 25K cities at a coarse-level (with TerraSar and LandScan data at a spatial resolution of $750m - 1km/px$), while in this present work, we were able to work at a much finer-level of urban detail (with Sentinel-2 data at a spatial resolution of $10m/px$).

Wang et al. [62] studied whether the commercial activeness of an urban area, proxied with the number of online reviews for places in the area, can be predicted from satellite data (and street views). The work by Wang et al. is similar to ours, in that, a pipeline was used that extract features from satellite images and feeds those features into a regression method. However, there are two main differences compared to our present study: *i)* there is no investigation of urban vitality but only one of its components, i.e., commercial activeness; and *ii)* Wang et al. work used *commercial* satellite imagery, compared to the *openly available Sentinel-2 imagery* that we used. The commercial nature of the data might create hurdles in terms of inclusivity in the social computing community, especially for the developing world.

In summary, though previous work used satellite data to measure population density, commercial activeness, and other variables that influence urban vitality, no study directly focused on vitality. Furthermore, the majority of previous work used commercial satellite images such as those from Bing and Google Earth, instead of openly available Sentinel images, as we did. By basing this work on openly available satellite data, such as Sentinel, we show that we can systematically derive insights about the built environment (such as vitality) with a performance at par with methods based on other access controlled and expensive datasets. This is notwithstanding the differences in spatial resolution between Sentinel and other datasets. We foresee that this could further empower the community to perform follow-up studies, particularly around the developing parts of the world.

3   DATA AND METHODS

This section describes how each city's district was encoded into a *district feature vector*, which was then associated with the district's urban vitality value. First, we created *imagelet feature vectors* by extracting small image pieces (imagelets) from satellite imagery and parsing them using deep learning feature extractors (Section 3.1). Second, we computed the district's urban vitality by combing a variety of data sources, including Urban Atlas, Open Street Map, and mobile phone data
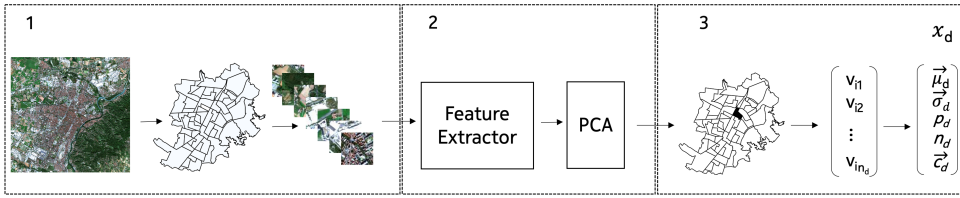
Fig. 2. Our framework for district feature extraction from satellite imagery. In step 1, the entire Sentinel-2 image of the city is split into imagelets (see Figure 3); in step 2, features are extracted from each imagelet using a deep learning feature extractor (see Figures 4 and 5), which is then followed by a PCA; and, in step 3, the imagelets i are grouped by their corresponding districts (Figure 6) and their feature vectors (v_i) processed to finally compute the representative district feature vectors (x_d).

(Section 3.2). Finally, we created the district's feature vector by combining the *imagelet feature vectors* corresponding to the imagelets composing the district, and associated the resulting *district feature vector* with the previously computed district's urban vitality value (Section 3.3).

3.1 Creating Imagelet Feature Vectors from Satellite Images

In 2014, the European Space Agency (ESA) has launched the first of the Sentinel satellites as part of the Copernicus[1] program [58], which aims at democratizing access to Earth Observation (EO) data. Thanks to this multi-billion investment, we can today freely access accurate and timely land satellite data from the radar (Sentinel-1, Sentinel-3), altimeter (Sentinel-3), and optical (Sentinel-2) sensors. Other satellites from the Sentinel series provide atmospheric and oceanic data. Goals behind programs such as Copernicus are quite ambitious: they include managing the environment, helping mitigate the effects of climate change, monitoring agriculture, and supporting urban development [51], to just name a few.

3.1.1 Sentinel-2 Imagery. For this study, we used Sentinel-2 optical imagery, which monitors the land surface conditions and serves for producing land-cover and land-change detection maps [17]. The optical sensors (Multi-Spectral Instrument [9]) sense 13 spectral bands (B1-B13) ranging in spatial resolution from $10m$ to $60m$. Out of the several processing levels in which the Sentinel-2 imagery is distributed, Level-1C and Level-2A are freely available to public. The Level-2A is the higher processing level and it includes atmospheric correction [39] on top of Level-1C products. Additionally, the imagery in this product is orthorectified (projected to the geodetic coordinates) and generated with an equal spatial resolution of $10m$ for the three so-called True Color Image (TCI) bands (i.e., B4, B3, B2). The TCI bands combination gives natural color representation of Sentinel-2 data [22] and is widely used in land cover studies [54]. Given its richness and that all these necessary pre-processing steps are already performed on it, we selected the Level-2A product for this study. We downloaded the Sentinel-2 Level-2A products for selected six Italian cities from 2018 (i.e., the earliest available, since this product became operational in 2018). All the images were re-projected to the *WGS 84 / UTM zone 32N* geodetic coordinate reference system (CRS), also sometimes denoted as *EPSG:32632*. This CRS is suitable for use in between 6°E and 12°E, northern hemisphere and between equator and 84°N – the area that includes Italy.

3.1.2 Pre-processing and Creating Imagelets. To create training data, for each city, we processed the Sentinel-2 imagery in three steps (Figure 3):

   (1) created geo-referenced raster image from the TCI bands (B4, B3, B2);

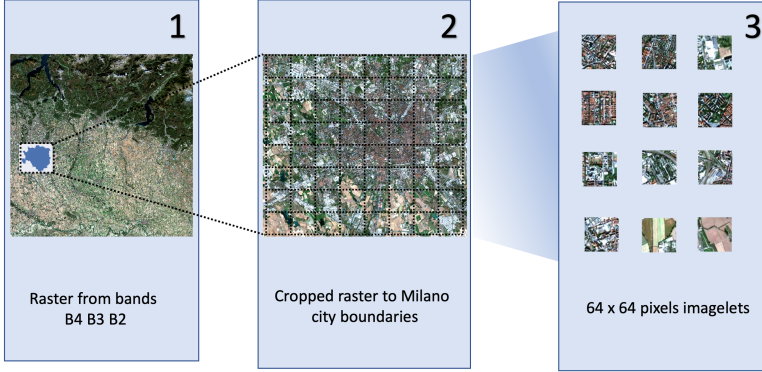[1]https://www.copernicus.eu/en/about-copernicus/copernicus-brief

Fig. 3. Process of creating training imagelets from the source Sentinel-2 imagery for the city of Milan.

(2) cropped the image based on the city boundaries (derived from an official shapefile); and

(3) split the cropped image into imagelets of size 64×64 pixels.

This imagelet size was selected so that the resulting imagelets are not too large compared to a district's size (Table 1) and not too small in terms of their final resolution, but they are also enough large to be successfully processed by our deep learning framework (Section 3.1.3). The pre-processing resulted in 9,115 imagelets in total across the six cities.

3.1.3 Deep Learning Feature Extractors. To extract visual features from the satellite imagelets, we applied a pipeline consisting of a standard convolutional neural network (CNN) [32] architecture followed by a principal component analysis (PCA). Given an image $i \in I$, where I is the set of imagelets, our goal was to retrieve a vector of uncorrelated and ordered visual components $v_i \in V$. To generate such a vector we used two techniques: the first was based on a pre-trained feature extractor (Figure 4), and the second consisted of training an unsupervised convolutional autoencoder (CAE) feature extractor (Figure 5).

Pre-trained CNN Feature Extractor. We adopted the $VGG16$ [53] CNN architecture (Figure 4) (weights pre-trained on ImageNet [48]): given an imagelet i, we extracted the output of the last but one fully connected layer (4,096 features) (z_i) and summarised these features z_i into a sparse set of uncorrelated components v_i^{CNN} using a PCA.

Convolutional Autoencoder (CAE) Feature Extractor. Since $VGG16$ architectures were not trained on "overhead view" satellite images, we chose an additional architecture which we could train on our imagelets. This architecture uses a convolutional autoencoder (CAE) [5, 41] (Figure 5), followed by PCA [33]. CAE is an unsupervised approach which comprises of two parameterized functions, a deterministic encoder $f_w(\cdot)$, and a deterministic decoder $G_u(\cdot)$. Convolutional layers can be stacked sequentially where the encoding layers reduce the dimension of the image i to a latent embedding z_i while the decoding layers expand the dimensions back to its reconstruction i'. The sequential architecture can be seen in Figure 5. Following [41], the parameters of the encoder $z_i = F_w(i)$ and the decoder $i' = G_u(z_i)$ are updated by minimizing the reconstruction losses between i and the reconstructed version $i' = G_u(F_w(i))$:

$$\mathcal{L}_{REC} = \frac{1}{n_{batch}} \sum_{j=0}^{j=n_{batch}} (i_j - G_u(F_w(i_j)))^2, \tag{1}$$

Table 1. Number of districts (N_c), their average size (in km^2) and population, and number of imagelets acquired per city.

City	#Districts (N_c)	Mean Area (in km^2)	Mean Population	#Imagelets
Milan	85	1.72	14,551	328
Bologna	23	3.34	15,918	170
Florence	21	2.89	16,633	135
Palermo	43	2.01	15,075	193
Turin	56	2.00	15,543	252
Rome	146	3.24	17,123	1068

Table 2. *CAE* Architecture with explicitly enumerated stack of *encoder* and *decoder* layers

LAYER	RESAMPLE	NORM	FILTER SIZE
Image	-	-	-
Conv	MaxPool	BatchNorm	5×5×16
Conv	MaxPool	BatchNorm	7×7×32
Conv	MaxPool	BatchNorm	7×7×32
Conv	-	BatchNorm	9×9×32
Conv	-	BatchNorm	9×9×32
Z			512×1
ConvTranspose	Upsample	-	9×9×32
ConvTranspose	Upsample	-	9×9×32
ConvTranspose	Upsample	-	7×7×32
ConvTranspose	Upsample	-	7×7×32
ConvTranspose	Upsample	-	5×5×16
Recon	-	-	-

where n_{batch} is the number of images in a batch, and $\mathcal{L}_{REC}$ is the reconstruction loss.

The convolutional autoencoder was trained with the satellite images using the ADAM [31] optimiser with a learning rate of 0.001 for 500 epochs with a batch size of 128. After extensive testing, we selected an embedding dimension of 512 since it gave us the best trade-off between compression and a lower reconstruction loss, where the loss stabilizes at 0.01 (Table 2 reports the implementation details of the architecture we ended up using). After training, we extracted a lower dimension embedding from the satellite imagelets using the trained encoder, which is then summarized into uncorrelated components v_i^{CAE} using *PCA*.

For both, *VGG*16 and CAE, we experimented with $n_{comp} = 12, ..., 64$ PCA components in our evaluation (Section 4), and found that using $n_{comp} = 16$ components yielded the best results when using all the data from our six cities. In the rest of the paper, for simplicity, we denote imagelet feature vectors with v_i, while we experimented with both v_i^{CNN} and v_i^{CAE}.

3.2 Computing Six Vitality Proxies and Urban Vitality for a District

We built upon previous work by De Nadai et al. [14]. The authors empirically validated Jane Jacobs theory across six cities in Italy. They showed that the four generators of urban vitality defined by Jacobs, i.e., *land use, small blocks, diversity of economic activity*, and *concentration of people*, indeed predicted urban vitality.

From our satellite imagery, we can estimate neither diversity of economic activity (it is not a property of the urban form) nor concentration of people (as our satellite images are open data but medium-resolution). By contrast, land use and small blocks are the two vitality generators out of the four that could be potentially visible from satellites. To test the extent to which they are so, we
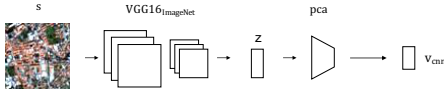
Fig. 4. Schematic of our pre-trained CNN feature extractor. We extracted a set of features z from satellite image using a pretrained $VGG16_{imagenet}$ [53]. We then summarised these features into a set of components v^{CNN} with a PCA.
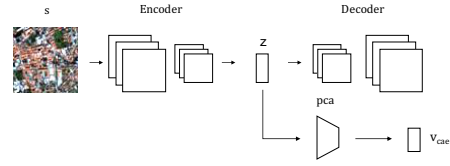
Fig. 5. Schematic of our Convolutional Auto Encoder (CAE) feature extractor. We trained a CAE, where the encoder compresses the satellite image and a decoder reconstructs it. After training, we extracted a set of features z from the encoder and summarise it into a set of components v^{CAE} with a PCA.

Table 3. The six vitality proxies plus urban vitality itself together with their distribution, mean (μ), and standard deviation (σ).

Land Use		μ	σ	Small Blocks		μ	σ
Land use mix		0.733	0.201	Block size		9.618	0.459
Building height		0.689	0.216	Intersection density		10^{-4}	10^{-4}
Small parks		0.004	0.003	Anisotropicity		0.385	0.042
Vitality		μ	σ				
Activity density		0.006	0.005				

operationalized them with *six vitality proxies* (Table 3) computed for the six cities (Milan, Florence, Bologna, Turin, Palermo, and Rome).

Before describing these six proxies, we have to define and motivate our choice for spatial unit of analysis. We took *district* (or *area di censimento*) as the unit of analysis, as was done by previous work [14]. Districts are census areas, which consist of neighboring blocks (i.e., sections delimited by street segments) grouped based on socio-economic conditions. In the Italian context, districts are comparable in terms of population (having between 13K and 18K inhabitants) and size (covering an average area of $2.47km^2$) (Table 1). The average population density of a district for the 6 cities in our dataset is 10K per km^2. These parameters correspond to those of the spatial areas that Jacobs discussed in her book [28].

3.2.1 Land Use. Jacobs suggested that a mix of primary uses in a district promotes vitality [28]. Examples of primary use categories are residential buildings, office spaces, industrial, entertainment, education, recreation, and cultural facilities. One of the most comprehensive datasets on the urban land use categories is the Urban Atlas[2] produced for most of the cities in Europe including our test cities. It provides information on 20 land use classes (such as green urban areas, sports and leisure facilities, and urban fabric). The Urban Atlas is built from high resolution satellite imagery in combination with ancillary data, such as Google Earth, Open Street Map, and manually-collected in-situ data.

Land use mix. From the Urban Atlas, we first computed *land use mix* in each district. According to Manaugh and Kreider [40], land use mix can be calculated as the entropy of three main categories

---

[2]http://www.eea.europa.eu/data-and-maps/data/urban-atlas

of urban land uses: i) *residential*, ii) *commercial, industrial, institutional* and *governmental*, and iii) *recreational, parks*, and *water*. We used the formula:

$$\text{land use mix}_d = -\sum_{j=1}^{n_{LUC}} \frac{P_{d,j} \log(P_{d,j})}{\log(n_{LUC})} \tag{2}$$

where $P_{d,j}$ is the percentage of area with land use j in district d, and $n_{LUC} = 3$ is the number of land use categories.

Small parks. By attracting people to spend time or just walk there, small parks promote vitality [4, 28]. For each district d, we calculated the average distance of all its blocks from the nearest small park (with the area < 1 km^2):

$$\text{small parks}_d = \left(\frac{1}{|B_d|} \sum_{j \in B_d} \text{dist}(j, \text{closest}(j, SM))\right)^{-1} \tag{3}$$

where B_d is the set of blocks in district d, closest(j, Y) is a function that finds the geographically closest element in set Y from block j's centroid, SM is the set of small parks, and dist(a, b) is the geographic distance between two elements' centroids a and b.

Building height. Jacobs posited that lower building height encourages the opening of restaurants, stores, and other services, which, in turn, all promote pedestrian activity [28]. Sung et al. [56] used the average number of floors per building in a district as a proxy for *building height*, which was computed as:

$$\text{building height}_d = \frac{\sum_{hc} b_{hc,d} \cdot f_{hc}}{\sum_{hc} b_{hc,d}} \tag{4}$$

where $b_{hc,d}$ is the number of buildings that are in height category hc in district d, and f_{hc} is the number of floors corresponding to height category hc. To collate data on building floors, we used the Italian National Institute for Statistics (ISTAT)[3].

3.2.2 Small Blocks. Another of our six vitality proxies is small blocks, which have been found to increase walkability and opportunities of cross-use.

Block sizes. A district's *block size* variable can be defined as the average size of all the blocks B_d within the district d:

$$\text{block sizes}_d = \frac{1}{|B_d|} \sum_{j \in B_d} \text{area}_j \tag{5}$$

Intersection density. Another variable in the category of small blocks contributing to urban vitality by increasing random contacts is *intersection density*:

$$\text{intersection density}_d = \frac{|\text{intersections}_d|}{\text{area}_d} \tag{6}$$

Anisotropicity. Anisotropy refers to the irregularity of a geometric shape concerning its orientation and spacing. A block can be of relatively small size, but, if it exhibits an anisotropic shape (i.e., one of its sides is considerably longer than the other), then that decreases the chance of being in contact with one of its sides. We assigned *anisotropicity* to each district d as the average anisotropicity [36] of its blocks B_d:

$$\text{anisotropicity}_d = \frac{1}{|B_d|} \sum_{j \in B_d} \Phi_j \tag{7}$$

---

[3]http://www.istat.it/it/archivio/104317

where Φ_j is the ratio between the area of the block j and the area of its circumscribed circle C_j:

$$\Phi_j = \frac{\text{area}_j}{\text{area}_{C_j}}. \tag{8}$$

3.2.3   Urban Vitality. We used mobile Internet activity as a proxy for urban vitality. Unlike mobile phone data on calls and SMS activity, this type of data senses peoples' presence even when they are not actively using their phones, enabling better tracking of their mobility in an area. Telecom Italia Mobile, i.e., the largest mobile operator in Italy (34% of total mobile phone user base), provided its data from February to October 2014.

Activity density. To estimate the number of Internet connections that fell into each district, we represented the space with a set of Voronoi polygons [18] based on the radio stations' positions. To estimate the Internet activity in each district d at time t, we calculated the number of Internet connections over all polygons p's that fell into district d:

$$S_d = \sum_v R_p \frac{A_{p \cap d}}{A_p - A_{p \cap W}} \tag{9}$$

where p is a polygon, and R_p is the number of Internet connections in p throughout a typical business day. The count of Internet connections is weighted by $\frac{A_{p \cap d}}{A_p - A_{p \cap W}}$, which is the proportion $\frac{A_{p \cap d}}{A_p}$ of p's area that falls into district d ($A_{p \cap d}$ is p's area that falls into district d, and A_p is p's total area). From p's total area we removed sea areas denoted by W (i.e., we removed $A_{p \cap W}$).

3.3   Combining District Feature Vectors with Urban Vitality

Having prepared imagelet-level features and district-level labels, we combined the two datasets as follows. First, we assigned imagelets to districts with the following procedure. If an imagelet was overlapping with several districts, we assigned it to the district with which it has the largest overlap. An illustration for assigning imagelets in Turin is shown in Figure 6. Out of the 9,115 original imagelets, most of them fell outside of the city's boundaries. With this procedure, we were able to assign a different percent of imagelets for each city, in total 2,146 imagelets (the breakdown per city is shown in Table 1). Notice, however, that we still used all the 9K imagelets to train the autoencoder feature extractor *CAE* (described in Section 3.1). These imagelets constituted the basis for our train and test set.

Second, using the imagelet features, for each city and each of its districts, we derived a representative district feature vector in a way similar to the previous work [26]. For each district $d \in 1, ..., N$, we took feature vectors $v_{i_1}, v_{i_2}, ...v_{i_{n_d}} \in R^{n_{comp}}$ (where n_{comp} is the number of PCA components discussed in Section 3.1) of its imagelets $i_1, i_2, .., i_{n_d}$, and aggregated them into a single vector $x_d = (\vec{\mu}_d, \vec{\sigma}_d, p_d, n_d, \vec{c_d})$, where $\vec{\mu}_d$ is the vector of size n_{comp} obtained by taking the mean across the n_{comp} components of the district imagelet feature vectors, $\vec{\sigma}_d$ is the vector also of size n_{comp} obtained in a similar way by taking the standard deviation across the components of imagelet vectors, p_d is a number obtained by taking the the average Pearson's correlation between pairs of imagelet vectors, n_d is the total number of imagelets in d, and $\vec{c_d} = (c_d^{lon}, c_d^{lat})$ is a vector of size two representing the district centroid (i.e., the lon/lat of the centroid point). The first four groups of elements capture and describe district satellite data by central tendency ($\vec{\mu}_d$), dispersion ($\vec{\sigma}_d$), association (p_d), and size (n_d). Compared to [26], we additionally enriched our visual feature set with geolocation features, i.e., district centroids ($\vec{c_d}$) as previous research has shown that their inclusion encodes spatial positioning of the imagelets, and improves performance [10, 57]. Hence, if we denote as x_d^k the k^{th} element of the feature vector x_d, then:
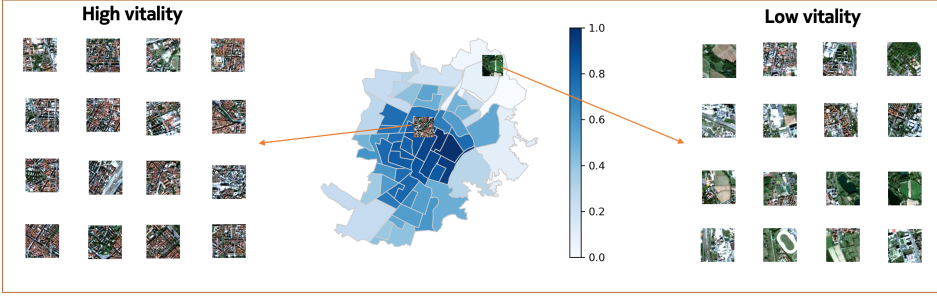
Fig. 6. Our approach of associating imagelets with districts and their vitality levels (which are computed from mobile phone activity density and are our labels) in the city of Turin. A darker color denotes higher values of urban vitality. The two imagelets shown on the map will be associated with different vitality levels (i.e., the imagelet in the city center with high vitality, while the one on the outskirts with low). Examples of imagelets associated with high versus low vitality levels are also shown. The high-vitality imagelets are characterized by a denser urban fabric, smaller blocks, and more brown-reddish colors (likely representing the rooftops of low-rise buildings). On the other hand, the low-vitality imagelets feature large homogeneous areas (e.g., stadium, larger parks, and fields), a low-density urban fabric, and white colors (likely representing industrial buildings).

The first n_{comp} elements are the average values of the 1^{st}, 2^{nd}, …, n_{comp}-th component (respectively) of all the district's imagelet vectors,

The elements from position $(n_{comp} + 1)$ to position $(2 \cdot n_{comp})$ are the standard deviation values of the 1^{st}, 2^{nd}, …, n_{comp}-th component of all the district's imagelet vectors,

The element at position $(2n_{comp} + 1)$ is the Pearson's correlation of the imagelet vectors,

The element at position $(2n_{comp} + 2)$ is the count of imagelets in the district, and

The last two elements are the longitude and latitude coordinates of the district's centroid.

Our way of defining the district feature vectors has the nice property that they are of equal length (i.e., $x_d \in R^m$, where $m = 2n_{comp} + 1 + 1 + 2$), independently of the district sizes or the number of imagelets.

Last, for each of the collected six proxies (Table 3), we took the corresponding value y_d for the district d and associated this value with the feature vector x_d of the district. This combined dataset constituted our training and test set.

4 EVALUATION

Having prepared the training and test data, we set up our two main evaluation experiments. The main goals of these experiments were to test whether we could estimate two classes of quantities from satellite imagery: *1)* the six proxies in Table 3 (Section 4.4), and *2)* urban vitality itself (Section 4.5).

4.1 Setup

Using the produced district feature vectors, for each of the six proxies in Table 3, we created a regression task whose dependent variable is initially each of the six vitality proxies. The distributions of different variables, including activity density used as a proxy for urban vitality, are shown in Table 3. Since activity density and the variables in the *land use* category were skewed, we log-transformed them using the natural logarithm. Additionally, since the scales of the variables varied considerably,

we standardized and normalized them before running regression models so that we could later compare their relative importance.

4.2   Regression methods

Given the district feature vectors $x_d \in R^m$ (calculated as in Section 3.3) and their corresponding labels $y_d \in R$ (calculated as in Section 3.2), we aimed at predicting $\hat{y}_d$, using linear regression models of the form

$$\hat{y}_d = w_0 + w_1 x_d^1 + ... + w_m x_d^m, \quad d = 1, ..., N, \tag{10}$$

where the coefficients w_j are learned by the model, m is district feature vector size, and N is the total number of districts.

In addition to **Ordinary Least Squares regression**, which minimizes $||y_d - X_d w||_2^2$, where X_d denotes the matrix of all district vectors x_d; we also experimented with **ElasticNet regression**, which minimizes $\frac{1}{2N}||y_d - X_d w||_2^2 + \alpha\rho||w||_1 + \frac{\alpha(1-\rho)}{2}||w||_2^2$, where w is the vector of weights $w = (w_1, ..., w_m)$, and α and ρ are the parameters that control the L_1 (absolute value of coefficients magnitude) or L_2 (squared value of coefficients magnitude) penalisation. We chose ElasticNet because it encourages sparsity, i.e., automatically eliminates insignificant variables, and in that way deals better with the high-dimensional satellite feature vectors [38]. We performed a grid search across the parameter space while predicting vitality, and we found these parameters to work best $\alpha = 0.01$, and $\rho = 0.1$.

Moreover, we also experimented with **Support Vector Regression (SVR)** [8] method. A grid search on the parameter space yielded RBF kernel with this set of optimal hyper parameters $C = 1.0$, $degree = 3$, $\gamma = scale$, and the constant $\epsilon = 0.0001$.

Finally, we also tested an ensemble of decision trees with extreme gradient boosting [21], i.e., **XGBoost regression**. On small structured datasets, such as the one we dealt with, decision tree based algorithms are considered to be among the best performing methods. Upon parameter grid search, we used 350 weak learners, learning rate of 0.01, the *huber* loss, and the maximum tree depth of 3.

4.3   Metrics

To assess how well each of the regression methods predicted our urban variables from satellite features, we resorted to two of the standard metrics used in regression methods: coefficient of determination (R^2), its adjusted version (R_{adj}^2), and mean absolute error (MAE). Simply, R^2 measures the proportion of the variance in the target variable (urban variables) that is predictable from the input variables (image features), while its adjusted version R_{adj}^2 accounts for the number of input variables and training data size, hence preventing a spurious increase in R^2 that is only due to the introduction of new input variables. MAE measures the errors between true and predicted values (set of l value pairs):

$$MAE = \frac{1}{l} \sum_{j=1}^{l} |y_j - \hat{y}_j|. \tag{11}$$

4.4   Predicting the Six Vitality Proxies

We applied the selected regression methods to estimate each of the six proxies in Table 3, and then used the estimated values to predict urban vitality with a two-stage regression. For each of the methods, we conducted a 5-fold cross-validation 100 times (repeated k-folds) and averaged the scores. This is done to obtain stable results given a high variance from our low sample size.
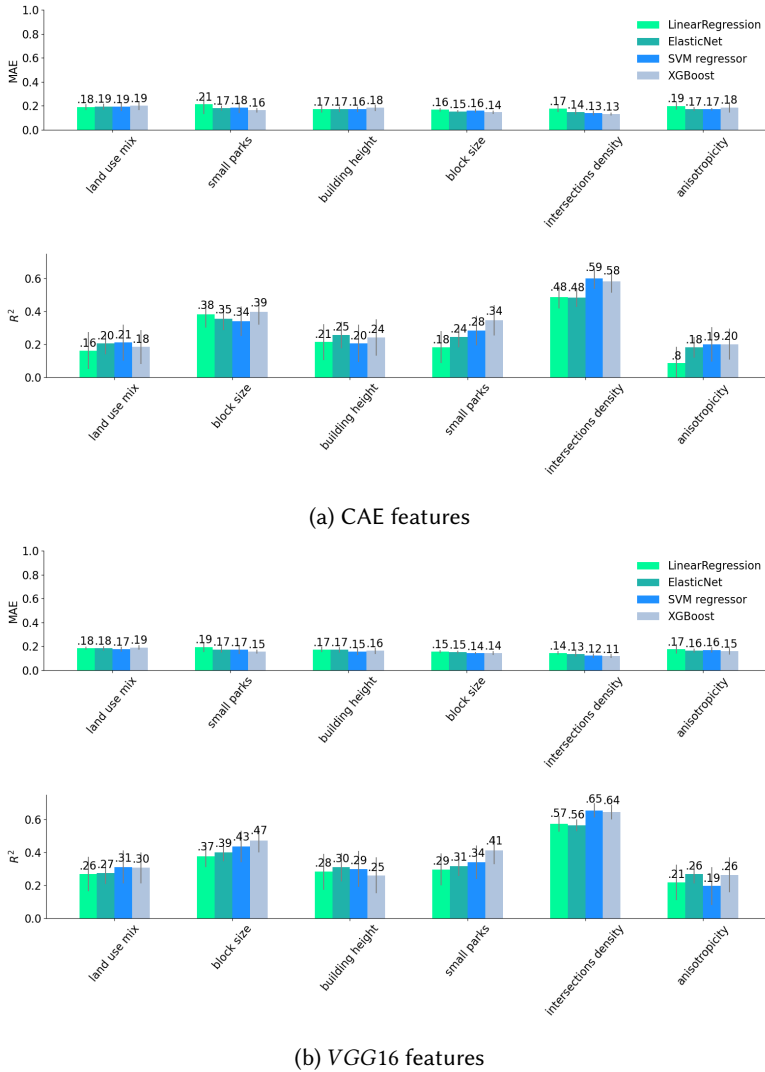
(a) CAE features

(b) *VGG*16 features

Fig. 7. Regression models predicting the six vitality proxies in Table 3 from satellite-derived features: (a) CAE-based, (b) *VGG*16-based features. This is the first stage of our two-step regression model, where we predict the features first, before predicting vitality. Average results of 5-fold cross-validation ran 100 times are shown: (Top) Mean absolute error (MAE) scores; (Bottom) Coefficients of determination (R^2). Error bars represent standard deviation across the runs.

The results of the first stage are shown in Figure 7. The starting observation is that the different features (*VGG*16 vs CAE-based) yield comparatively similar patterns across the variables, i.e., they both predict *intersection density* and *block size* best. However, the *VGG*16 features performed better overall for all the variables, despite the fact that the CAE was trained on satellite imagery and *VGG*16 on regular imagery. It is likely that this is due a relatively small size of our dataset (9K imagelets, in total) and that CAE features can become competitive if a larger dataset for training would be used. Hence, from now on, we discuss the results for *VGG*16 features. Out of the six

Table 4. Regression results for predicting vitality from the previously inferred six vitality proxies. 5-fold cross-validation results are shown. This is the second stage of our two-step regression model.

Method	R2
LinearRegression	.336 ± .063
ElasticNet	.366 ± .070
SVM regressor	.357 ± .064
XGBoost	.228 ± .097

Table 5. ElasticNet coefficients for predicting vitality from the six proxies estimated from satellite imagery.

Variable	coef
intersections density	0.497
small parks	0.124
anisotropicity	0.063
land use mix	0.019
block size	-0.118
building height	-0.152

proxies, *anisotropicity* is the hardest to predict. SVR and XGBoost are two methods that are the best performing overall. Another variable that is challenging to predict is *building height* (the R^2 score of SVR is .30 and of XGBoost .24). For the four other variables, SVR was able to predict them with R^2 score of .31 for *land use mix*, .43 for *block size*, .33 for *small parks*, and up to .65 for *intersection density*. They are easier to predict not least because they can be seen in the images by the naked eye.

In the second stage, we studied whether we can predict vitality indirectly from the previously estimated six proxies. In particular, we considered all the selected variables, took their predicted vectors and applied the same regression models to predict vitality. We found (results in Table 4) that the best performing method in this setup was the form of linear regression, i.e., ElasticNet ($R^2 = .37 \pm .07$).

Finally, we studied the relative importance of the six variables in predicting vitality. By looking at the coefficients found by ElasticNet for each of those variables (Table 5), we found that the most predictive variable positively influencing vitality is intersection density (coefficient .497), which is in line with the previously reported results by De Nadai et al. [14]. Also *small parks* is positively influencing vitality, while *block sizes* negatively influences vitality. The *building height*, which is defined through the average number of floors in buildings, matters as well: the higher it is, the less chances for creating services, such as restaurants and stores at the ground floor – and, hence, the lower the vitality.

4.5 Predicting Vitality

In our second experiment, we evaluated the predictive power of the four regression models in estimating vitality directly from the visual features of the satellite imagery. First, for each of the models, we conducted a 5-fold cross-validation experiment (Figure 8). In the previous set of the experiments, SVR and XGBoost regressor were the best models and shown a similar performance. In this experiment, SVR ($R^2_{adj} = .55 \pm .03$, $MAE = .12 \pm .00$) slightly outperformed XGBoost ($R^2_{adj} = .54 \pm .04$, $MAE = .13 \pm .01$).

In the final experiment, we studied the generalizability of the urban vitality prediction, i.e., we asked whether the models can be trained on some cities to predict vitality in another (unseen) city. For this purpose, we conducted leave-one-city-out validation experiments and used SVR, as
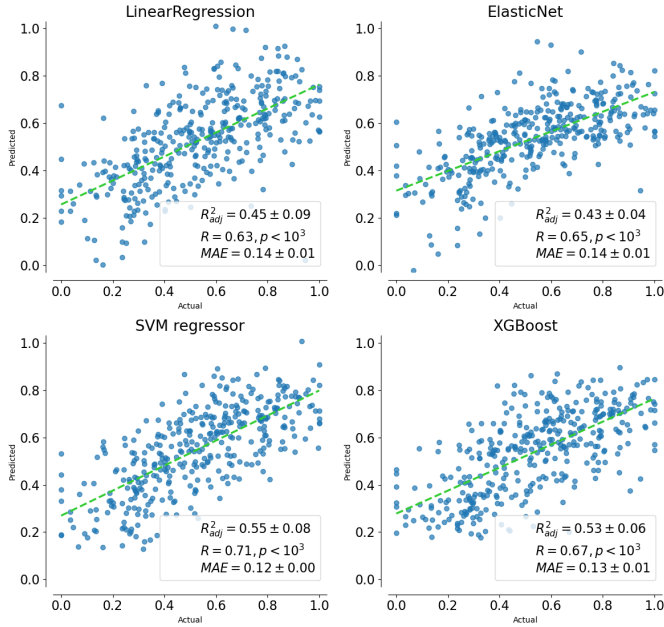
Fig. 8. Regression models that predict vitality: 5-fold cross-validation results. Each plot represents a different model evaluated in terms of: adjusted coefficient of determination (R^2_{adj}), Pearson correlation coefficient (R) ($p < .001$), and mean absolute error (MAE). To aid visual interpretation, a green line representing the best linear fit between true and predicted values is shown.

it was the best performing model in the previous experiments. Specifically, for each of the six cities, we trained the model on the data from five cities and tested it on the remaining city. The results are shown in Figure 9. The model's predictive power varies: in Milan and Florence, it is able to explain, respectively, 50% and 61% of the variance in vitality, but, in Palermo and Rome, the explained variance is, respectively, 32%, and 25%. To illustrate the ability of satellite-derived features in predicting vitality, in Figure 10, we show maps with true and predicted vitality for the six cities.

4.6 Explaining Vitality and its Proxies

Our analyses up to this point suggested that satellite data can predict urban vitality, and it can do so with varying degrees of accuracy. Our best performing method (SVR), tested across all the cities, achieved an R^2_{adj} of .55, and its scores on individual cities ranged from .61 for Florence to .25 for Rome. In this section, we study the factors potentially affecting prediction accuracy (Subsection 4.6.1), and the association between models' inference and the presence of PoIs in the area (Subsection 4.6.2).

4.6.1 Factors affecting vitality prediction. The city with the lowest prediction score in our experiments is Rome (R^2_{adj} = .25). That could be partly because it has the smallest training set. To see why, consider that Rome has the *largest number of districts* (see Table 1) and, in our leave-one-city-out evaluation, Rome's training set consists of areas in the remaining cities, while the test set consists
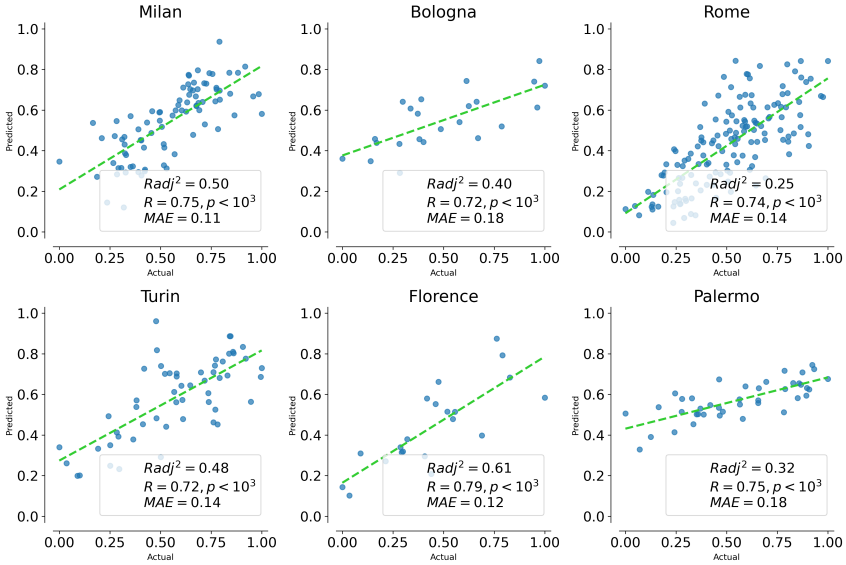
Fig. 9. Leave-one-city-out prediction results using SVR. The model is trained on the data from 5 cities and tested on the last city. Each plot represents the model for a different city (the one on which it was tested) evaluated in terms of: adjusted coefficient of determination (R_{adj}^2), Pearson correlation coefficient (R) ($p < .0001$), and mean absolute error (MAE). To aid visual interpretation, a green line representing the best linear fit between true and predicted values is shown.
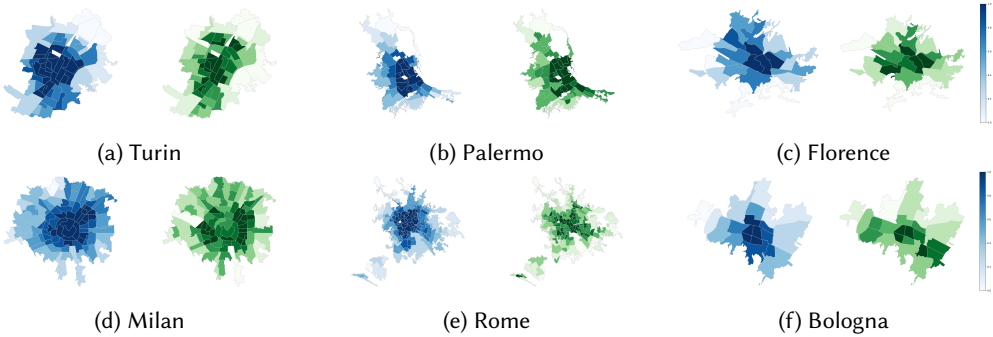

Fig. 10. Maps of the true (blue) and predicted (green) urban vitality levels. Across the cities, the predictions follow an overall pattern of higher vitality in the centers, and lower vitality in the suburbs. However, inconsistencies exist for specific cities. Notably, in Milan, predictions for the districts in the very center of the city (containing two large parks) are high but not maximum. On the other hand, our model overestimates vitality for Bologna's East districts, which, again, happen to feature small parks and relatively small blocks.

of areas in Rome. As a result, Rome has the *smallest training set* compared to the other cities. So one might speculate that the limited training data could influence the adjusted R^2 score, i.e., R_{adj}^2. Indeed, Rome's unadjusted R^2 is .39, which is significantly higher than its R_{adj}^2. To study whether this is the case for the other cities, we set out to determine whether prediction performance (R_{adj}^2)
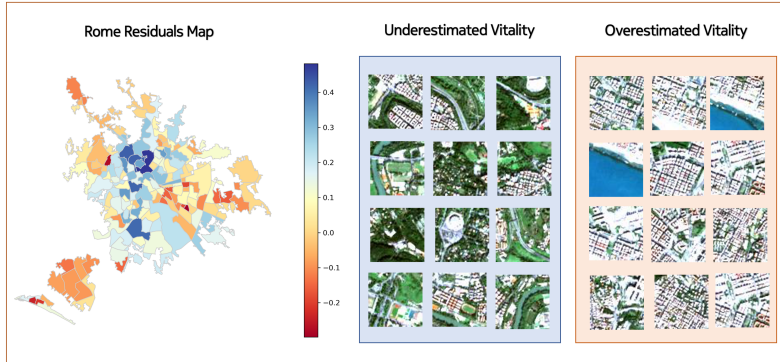
Fig. 11. The map of Rome's residuals and its satellite views for which our model overestimated and underestimated vitality level. The map shows residuals, i.e., differences between the true and predicted vitality levels: blue/red means that our model estimated lower/higher vitality levels compared to the ground truth. Example imagelets reveal that our model underestimated vitality levels in areas with large parks, rivers, highways, and stadiums, while it underestimated vitality levels in areas with high density of buildings and near the sea.

depends on the size of the training data. To obtain a city's training data size (denoted as TS), we summed the number of districts N_c from the remaining five cities (see Table 1). There was no significant correlation between TS and R^2_{adj} across cities, meaning that the prediction score was generally not influenced by the size of the training data.

Prediction residuals for Rome. To look for alternative explanations, we further focused on Rome. Figure 11 shows the map of its residuals, i.e., the differences between true and predicted vitality levels. The areas in blue/red are those for which our model underestimated/overestimated. We found that our model tended to *underestimate* in: *i)* areas with large parks and green spaces, rivers, highways, and stadiums (i.e., urban features that Jane Jacobs named *border vacuums*); and *ii)* areas with ancient landmarks in the city center. Jane defined border vacuums as *"perimeters of a single massive or stretched-out use of territory [...] that exerts an active influence on pedestrian activity"*. According to her theory, border vacuums can, indeed, depending on the context, be associated with both high vitality (*"if some visual or motion penetration is allowed through it, [...] it then becomes a seam rather than a barrier, a line of exchange along which two areas are sewn together"*) and low vitality (if the border vacuum cuts off activity in public spaces). Our analyses reveal that border vacuums are associated with high vitality in central parts of Rome, and with low vitality in the other cities (for example, Figure 6 shows border vacuums in low-vitality areas of Turin). Also, the overall percentage of greenery in Rome is higher than that of any of the other cities[4]. By inspecting the underestimated central locations on Google Maps, we also found that they included ancient landmarks popular among tourists (e.g., Largo di Torre Argentina, Foro Romano). These landmarks are not urban features typically found in the other cities (especially not in Turin where the size of its Roman district is quite limited), and, based on satellite images, might appear to be similar to ruins. That might be why the model trained on the cities other than Rome did not learn to label those locations in Rome as being of high vitality, despite being popular among tourists.

---

[4]The percentage of greenery is 40% in Rome, compared to less than 13% in Milan. http://www.worldcitiescultureforum.com/data/of-public-green-space-parks-and-gardens
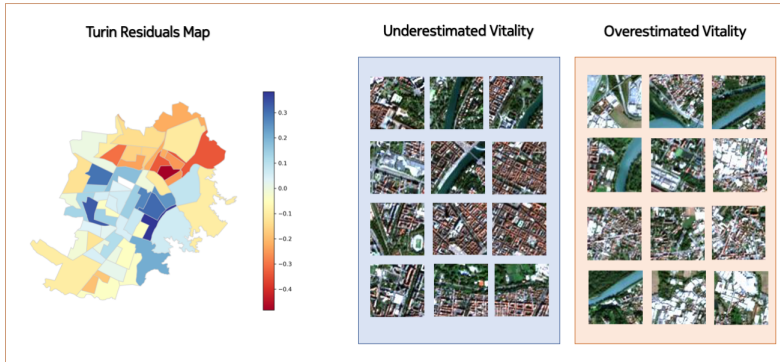
Fig. 12. The map of Turin's residuals and its satellite views for which our model overestimated and underestimated vitality level. The map shows residuals, i.e., differences between the true and predicted vitality levels: blue/red means that our model estimated lower/higher vitality levels compared to the ground truth. Example imagelets reveal that our model overestimated vitality levels in areas with large industrial buildings and close to the natural park, while it underestimated vitality levels in areas among the three popular Turin city parts (Porta Nuova, Parco del Valentino, and Giardini Reali Superiors).

On the other hand, we found that our model tended to *overestimate* vitality levels in areas with small blocks and dense housing predominantly located in the south-west district near the sea (the City of Fiumicino), which hosts the busiest airport in Italy. As one expects, given noise pollution, this district does not enjoy a rich outdoor life (including the presence of retail shops) as more central districts do, impacting one of the four Jacobs' dimensions, that of mixed economic activities.

Prediction residuals for Turin. To extend our prediction factors analysis to a different geographical, cultural, and economic environment compared to Rome, and also to look at a city where our models performed better, next we inspected Turin. Turin is located at the north of the country, does not have an access to the seaside like Rome; it has just one third of Rome's population, but due to the smaller size, it is three times more densely populated. Turin features a beautiful historical city centre, and buildings, castles, public squares in architectural styles ranging from renaissance, to rococo, and art nouveau. The city also hosts headquarters of several well-known Italian car-maker companies, and two large football stadiums.

Figure 12 shows the map of residuals for predicted vitality levels in Turin. This time, too, the areas in blue/red are those for which our model underestimated/overestimated vitality. We found that in Turin, our model tended to *underestimate* in: *i)* areas in between the largest Railway Station in the city, Porta Nuova, to the west, and a popular public park, Parco del Valentino, to the east, and up to the elegant park with statues and kid's play equipment, Giardini Reali Superiors, to the north, and *ii)* in areas next to the smaller athletic stadium, Stadio Primo Nebiolo. In the area among the three popular Turin city parts (Porta Nuova, Parco del Valentino, and Giardini Reali Superiors) that are at a walking distance from each other, indeed, the pedestrian mobility is maximally high. However, from the satellite imagery showing a mix of buildings, rivers, and relatively-large parks, our models inferred it to be high but still not the maximum level. The area around the athletic stadium Stadio Primo Nebiolo, interestingly, is associated with high vitality, while, for instance, the

Juventus Stadium is found in a low-vitality area to the north-west. This, once again, confirms that border vacuums can have different interplay with vitality, and, in this specific case, that has caused our method to underestimate vitality around Stadio Primo Nebiolo.

On the other hand, we found that our model tended to *overestimate* vitality levels in the northeast areas of Turin, close to the river Po and the nature park, Riserva Naturale del Meisino e dell'Isolone Bertolla. To the north from the park, there are several large industrial buildings, hosting agricultural production, warehouses, and some cooperatives. These buildings, as can be seen in the satellite images for overestimated vitality in Figure 12, appear white, and are mixed with red buildings likely representing residential houses nearby, and with parks. We speculate that our models, in this case, took this diversity as a mix of old and new buildings close to small parks, overestimating vitality for them.

4.6.2 Association between the model's inference and PoIs presence. Our prediction pipelines consist of deep learning methods that extract visual features. One issue is that the extracted features are hard to interpret and explain. To understand the associations between these abstract image features and actual points of interest on the ground, we queried the OpenStreetMap (OSM) database and collected all the Points of Interest (PoI) across the 6 Italian cities. OSM classifies the PoIs into 6 broad categories of amenities[5]: *sustenance, education, transportation, financial, healthcare,* and *entertainment.* Since vitality is a measure of human traffic in an area, we focused on the categories of sustenance, transportation, and entertainment, which cover the majority of amenities contributing to a city's social life. Next, we mapped the locations of the PoIs onto the 2,146 imagelets described in Section 3.3. With this mapping, we were able to count in each imagelet the total number of PoIs belonging to either of the three PoI categories. For any given imagelet i containing a total number of $n_{sustenance}$, $n_{transportation}$, and $n_{entertainment}$ PoIs, we computed the PoI scores for each category:

$$S_{cat}^i = \log(1 + n_{cat}), \quad \forall cat \in \{sustenance, transportation, entertainment\} \tag{12}$$

The log operation allowed us to transform the long tailed distribution of the number of PoIs into a normal distribution. We added 1 to avoid undefined scores for imagelets with zero PoIs of any given category. We then used these three scores computed for all the imagelets to explain the association between the model's inference and the PoI categories on the ground.

We did this by splitting the imagelets into two groups, a separate split each time, for vitality and each of its proxies, based on their ground truth values. The first group contained all those imagelets whose ground truth values for vitality, or one if its proxies, were in the top tertile, which we called the *high* class. Similarly, the other group contained imagelets with values in the bottom tertile, which we called the *low* class. We then sampled 50% of the *high* and *low* groups into a training set of imagelets and kept the remaining ones as the test set. We then trained a set of binary classifiers to predict vitality and its proxies using the satellite-derived features on the training set. The models' performances were comparable to the original models', with AUC scores in a 5-fold cross validation setup equal to .63 ± .07 for small parks, .88 ± .06 for small blocks, and .93 ± .06 for vitality. We then predicted the binary classes for vitality and its proxies for the test set of imagelets using these trained models. Finally, to quantify the relationship between the predicted labels on the imagelets and the PoIs found on the ground, we binarized the predicted class labels, i.e., class $c_i = 1$, if imagelet i is predicted to be in the *high* class, and $c_i = 0$, if imagelet i is predicted to be in the *low* class. We also computed the category scores for any given imagelet i using equation (12). These scores are denoted by $S_{sustenance}^i$, $S_{transportation}^i$ and $S_{entertainment}^i$, and were normalized on the [0, 1] scale. With this setup, we then fit a logistic regression that predicts the probability of an

---

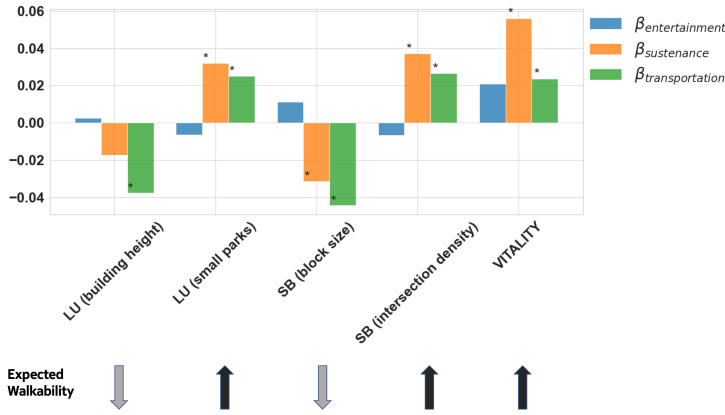[5]https://wiki.openstreetmap.org/wiki/Key:amenity

Fig. 13. The logistic regression (LR) coefficients for each OSM Places category in imagelets with high/low values (as output by our models based on $VGG16$-based features) for the six urban variables. The included variables are the five vitality proxies for which prediction on our satellite data with both types of features yielded R^2 above .25 (see Figure 7) and vitality. Significant LR coefficients (p-value $< .05$) are marked with a $*$. The up/down arrows denote the increasing/decreasing expected association between a specific urban variable (e.g., building height), and walkability/pedestrian presence in an area.

imagelet being in class 1, given the OSM PoI category scores.

$$Pr(c_i = 1) = \frac{1}{1 + e^{-(\alpha + \beta_1 \cdot S^i_{entertainment} + \beta_2 \cdot S^i_{sustenance} + \beta_3 \cdot S^i_{transportation})}} \tag{13}$$

Figure 13 shows the values for the logistic regression coefficients β_1, β_2, β_3 associated with the predictors $S_{sustenance}$, $S_{transportation}$, and $S_{entertainment}$, respectively. One can interpret the values of the β coefficients in terms of odds ratios. To get an upper bound of the predictive difference corresponding to a unit difference in each predictor, we used the 'divide by 4' rule [23] (p. 82). This says that, in a logistic regression, by dividing each predictor's coefficient by 4, one can draw conclusions about the increase in likelihood $Pr(c_i = 1)$ for each unit increase in the predictor. Simply put, we can quantify the percentage increase in the probability of the dependent variable to be in the positive class. Indeed, the results in Figure 13 reveal that satellite features capture meaningful insights in terms of places: all three place categories are positively predictive of vitality. Specifically, each 1% increase in the PoI score for the sustenance category increases the likelihood of the imagelet to belong in the *high vitality class* (+1.5%) and the *high intersection density* class ($\approx$ +1%). Similarly, a 1% increase in the PoIs score for the transportation category would decrease the corresponding imagelet's probability of belonging to the *high block size* class (-1%). PoIs in the sustenance and transportation categories are both predictive of *higher intersection density* and more likely presence of *small parks*. Also, as expected, these two categories are negatively associated with *high block size* (large blocks, which do not encourage walkability), and *high building height* (high-rise buildings). In summary, the results associated with PoIs from the sustenance and transportation categories meet expectations. For the entertainment category, however, this is only the case for vitality and for its proxy "block size" but not for the other proxies. That is partly because the OSM entertainment category[5] contains places that are typically found in dense city centers characterized by high vitality (e.g., fountain, cinema, theater) and places that act as border vacuums characterized by low vitality (e.g., casino, gambling, convention center).

5 DISCUSSION AND CONCLUSION

We have proposed a deep-learning framework that extracts features from publicly available Sentinel-2 imagery and predicts not only the six proxies for urban vitality (i.e., indirect prediction) but also vitality itself (i.e., direct prediction). The framework turned out to work better for the direct prediction of vitality (55% of the variance is explained) rather than its indirect prediction via six intermediate interpretable urban proxies (36%). That can be attributed to three reasons: *i)* the six proxies upon which the indirect prediction relied do not exhaustively capture vitality (e.g., mix of economic activities and concentration of people are not captured); *ii)* the prediction errors suffer from the two-step process; and *iii)* the direct prediction relied on raw satellite features, which may capture aspects overlooked by the indirect prediction.

Practical Implications. The presented predictive ability might be beneficial for supporting urban planning interventions in the developed world, and for sustainable development initiatives in the rapidly growing developing world. Specifically, we identified four main practical implications.

> *Satellite Data in City Dashboards.* Existing city dashboards and digital services could be enhanced with satellite data that may well fill the gap between the high frequency nature of mobile phone data and the low frequency nature of census data. We showed that some structural features of the urban environment that are important for vitality can be inferred from satellite data. Given that this data is publicly available, at a medium-level spatial resolution (10m) and continuously updated (every 5-7 days), it can be used for monitoring changes in the urban environment, to support urban planners/designers and policymakers in their decisions and planning. This is particularly relevant for developing countries, which frequently lack access to other types of urban data.

> *Guidelines for Urban Measurement from Satellite Data.* Our results also suggest that having highly diverse geographic data is critical when making inference on unseen cities, especially if these new cities are sufficiently different in terms of history or culture from the cities in the training data.

> *Facilitating Rapid Urbanization.* Jacobs's emphasis on diversification is particularly relevant to today's globalization: her theories shed light on urban inequality in Africa [43], slum clearance in preparation for mega-events [25], as well as efficient growth, such as in Taiwan [20]. Our approach can provide a dynamic view of neighborhood structure and help to track issues relevant to globalization.

> *Digital Earth.* Extending our methods to track vitality across the globe can contribute to the interactive view of our planet that enables a shared understanding of the relationships between the physical environment and society [11, 24].

Theoretical Implication. Replicating the experiments we did in Italy in other countries will require to obtain mobile phone data that can serve as a proxy for vitality in those countries. We can then learn visual features capturing vitality, which will likely uncover subtleties in how vitality is expressed across different natural and cultural environments.

Limitations and Future Work. While the openness of the Sentinel-2 dataset is an advantage, its limited spatial resolution is a limitation. If higher spatial resolution data would be publicly available, our models' performance will likely improve.

We also relied on mobile Internet data to quantify a proxy for vitality, which could be another limitation. However, since the market share of the provider (Telecom Italia) is 34%, which is the largest in the country, we believe that such a proxy offers a close approximation of the actual footfall in the specific context of our study. More generally though, in today's highly connected world, mobile activity is one of the best proxies for measuring presence of people [16, 45, 50, 55], and has been successfully used to estimate urban vitality [14, 30].

Another limitation with the data has to do with the temporal lag between mobile phone data and satellite data, which, in our case, amounts to three years. It is possible that some aspects of urban vitality could have changed during this time period.

Moreover, since this study is only conducted in the Italian context, its generalizability could be questioned. However, the fact that the model trained on data coming from five cities was able to predict the remaining city's vitality levels speaks to the generalizability of our approach, at least in the Italian, if not European, context. Rome posed a challenge though. Being ancient, Rome's city centre exhibits a unique mix of the old and the new, while our proxies for vitality are grounded on modern urban morphologies. A similar type of conflict between modern and ancient urban forms was also found for Barcelona [15], in which nearby districts tend to enjoy a highly diverse historical mix. For our models to better generalize to unseen cities, greater diversity in terms of urban, historical, cultural, and natural contexts in the training set might be required.

Another limitation is that our models are primarily capturing the *potential* for - rather than actual - vitality, and, in certain - albeit rare - circumstances, their predictions might not match the actual vitality levels. For example, if we took the satellite images of Milan during the full COVID-19 epidemic lock-down (during which pedestrian activity was heavily limited), our models would still predict levels of vitality similar to those of ordinary times. This limitation is, however, shared with the majority of previous work, and could be partly fixed whenever higher-resolution satellite images (in which pedestrian activity is observable) would be publicly available.

Another line of future work is to integrate multi-modal data, given that vitality is, by definition, a multi-faceted concept. By fusing satellite data with other freely available datasets that, for example, capture mix of commercial activities and concentration of people, one is expected to improve prediction accuracy. Finally, to obtain explainable predictions without compromising accuracy, explainable AI methods should be further researched [49].

REFERENCES

[1] Adrian Albert, Jasleen Kaur, and Marta C Gonzalez. 2017. Using convolutional networks and satellite imagery to identify patterns in urban environments at a large scale. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining.* ACM, 1357–1366.

[2] Adrian Albert, Emanuele Strano, Jasleen Kaur, and Marta Gonzalez. [n.d.]. The dark side of the Earth: benchmarking lighting access for all cities on Earth and the CityNet dataset. In *Proceedings of the Conference on Knowledge Discovery in Databases (KDD) Workshop on Urban Computing (UrbComp).* ACM.

[3] Matheus Araujo, Yelena Mejova, Michaël Aupetit, and Ingmar Weber. 2018. Visualizing geo-demographic urban data. In *Companion of the 2018 ACM Conference on Computer Supported Cooperative Work and Social Computing.* 45–48.

[4] Federica Banchiero, Ivan Blečić, Valeria Saiu, and Giuseppe A Trunfio. 2020. Neighbourhood Park Vitality Potential: from Jane Jacobs's Theory to Evaluation Model. *Sustainability* 12, 15 (2020), 5881.

[5] Yoshua Bengio, Pascal Lamblin, Dan Popovici, and Hugo Larochelle. 2007. Greedy Layer-Wise Training of Deep Networks. In *Advances in Neural Information Processing Systems 19*, B. Schölkopf, J. C. Platt, and T. Hoffman (Eds.). MIT Press, 153–160. http://papers.nips.cc/paper/3048-greedy-layer-wise-training-of-deep-networks.pdf

[6] Andrey Bogomolov, Bruno Lepri, Jacopo Staiano, Emmanuel Letouzé, Nuria Oliver, Fabio Pianesi, and Alex Pentland. 2015. Moves on the Street: Classifying Crime Hotspots Using Aggregated Anonymized Data on People Dynamics. *Big Data* 3, 3 (2015), 148–158.

[7] A. Bogomolov, B. Lepri, J. Staiano, N. Oliver, A. Pianesi, and A. Pentland. 2014. Once upon a crime: Towards crime prediction from demographics and mobile data. In *Proceedings of the 16th International Conference on Multimodal Interaction.* 427–434.

[8] Chih-Chung Chang and Chih-Jen Lin. 2011. LIBSVM: A library for support vector machines. *ACM transactions on intelligent systems and technology (TIST)* 2, 3 (2011), 1–27.

[9] Vincent Chorvalli, Vincent Cazaubiel, Stefan Bursch, Mario Welsch, Heinz Sontag, Philippe Martimort, Umberto Del Bello, Omar Sy, Paolo Laberinti, and François Spoto. 2010. Design and development of the Sentinel-2 Multi Spectral Instrument and satellite system. In *Sensors, Systems, and Next-Generation Satellites XIV*, Vol. 7826. International Society for Optics and Photonics, 78260J.

[10] Grace Chu, Brian Potetz, Weijun Wang, Andrew Howard, Yang Song, Fernando Brucher, Thomas Leung, and Hartwig Adam. 2019. Geo-Aware Networks for Fine-Grained Recognition. *2019 IEEE/CVF International Conference on Computer Vision Workshop (ICCVW)* (Oct 2019). https://doi.org/10.1109/iccvw.2019.00033

[11] Max Craglia, Kees de Bie, Davina Jackson, Martino Pesaresi, Gábor Remetey-Fülöpp, Changlin Wang, Alessandro Annoni, Ling Bian, Fred Campbell, Manfred Ehlers, et al. 2012. Digital Earth 2020: towards the vision for the next decade. *International Journal of Digital Earth* 5, 1 (2012), 4–21.

[12] Elizabeth Daly, Sheena Erete, Rosta Farzan, Gary Hsieh, Cliff Lampe, Claudia Lopez, Andres Monroy-Hernandez, Daniele Quercia, Raz Schwartz, and Amy Voida. 2015. Supporting cities, neighborhoods, and local communities with information and communication technologies. In *Proceedings of the 18th ACM Conference Companion on Computer Supported Cooperative Work & Social Computing*. 277–281.

[13] Munmun De Choudhury, Moran Feldman, Sihem Amer-Yahia, Nadav Golbandi, Ronny Lempel, and Cong Yu. 2010. Automatic construction of travel itineraries using social breadcrumbs. In *Proceedings of the 21st ACM conference on Hypertext and hypermedia*. 35–44.

[14] Marco De Nadai, Jacopo Staiano, Roberto Larcher, Nicu Sebe, Daniele Quercia, and Bruno Lepri. 2016. The death and life of great Italian cities: a mobile phone data perspective. In *Proceedings of the 25th international conference on World Wide Web*. 413–423.

[15] Xavier Delclòs-Alió and Carme Miralles-Guasch. 2018. Looking at Barcelona through Jane Jacobs's eyes: Mapping the basic conditions for urban vitality in a Mediterranean conurbation. *Land Use Policy* 75 (2018), 505–517.

[16] Pierre Deville, Catherine Linard, Samuel Martin, Marius Gilbert, Forrest R Stevens, Andrea E Gaughan, Vincent D Blondel, and Andrew J Tatem. 2014. Dynamic population mapping using mobile phone data. *Proceedings of the National Academy of Sciences* 111, 45 (2014), 15888–15893.

[17] Matthias Drusch, Umberto Del Bello, Sébastien Carlier, Olivier Colin, Veronica Fernandez, Ferran Gascon, Bianca Hoersch, Claudia Isola, Paolo Laberinti, Philippe Martimort, et al. 2012. Sentinel-2: ESA's optical high-resolution mission for GMES operational services. *Remote sensing of Environment* 120 (2012), 25–36.

[18] Qiang Du, Vance Faber, and Max Gunzburger. 1999. Centroidal Voronoi tessellations: Applications and algorithms. *SIAM review* 41, 4 (1999), 637–676.

[19] Abdallah El Ali, Sicco NA Van Sas, and Frank Nack. 2013. Photographer paths: sequence alignment of geotagged photos for exploration-based route planning. In *Proceedings of the 2013 conference on Computer supported cooperative work*. 985–994.

[20] David Ellerman*. 2004. Jane Jacobs on development. *Oxford Development Studies* 32, 4 (2004), 507–521.

[21] Jerome H Friedman. 2001. Greedy function approximation: a gradient boosting machine. *Annals of statistics* (2001), 1189–1232.

[22] A Gatti and A Bertolini. 2013. Sentinel-2 products specification document. *Available online (accessed February 23, 2015) https://earth. esa. int/documents/247904/685211/Sentinel-2+ Products+ Specification+ Document* (2013).

[23] A Gelman and J Hill. 2007. Analytical methods for social research.

[24] Al Gore. 1998. The digital earth: understanding our planet in the 21st century. *Australian surveyor* 43, 2 (1998), 89–91.

[25] Solomon J Greene. 2003. Staged cities: Mega-events, slum clearance, and global capital. *Yale Hum. Rts. & Dev. LJ* 6 (2003), 161.

[26] Sungwon Han, Donghyun Ahn, Hyunji Cha, Jeasurk Yang, Sungwon Park, and Meeyoung Cha. 2019. Lightweight and Robust Representation of Economic Scales from Satellite Imagery. *arXiv preprint arXiv:1912.08197* (2019).

[27] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 770–778.

[28] Jane Jacobs. 2016. *The death and life of great American cities*. Vintage.

[29] Neal Jean, Marshall Burke, Michael Xie, W Matthew Davis, David B Lobell, and Stefano Ermon. 2016. Combining satellite imagery and machine learning to predict poverty. *Science* 353, 6301 (2016), 790–794.

[30] Young-Long Kim. 2018. Seoul's Wi-Fi hotspots: Wi-Fi access points as an indicator of urban vitality. *Computers, Environment and Urban Systems* 72 (2018), 13–24.

[31] Diederik P. Kingma and Jimmy Ba. 2017. Adam: A Method for Stochastic Optimization. arXiv:cs.LG/1412.6980

[32] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. 2017. Imagenet classification with deep convolutional neural networks. *Commun. ACM* 60, 6 (2017), 84–90.

[33] Stephen Law and Mateo Neira. 2019. An Unsupervised Approach to Geographical Knowledge Discovery Using Street Level and Street Network Images. In *Proceedings of the 3rd ACM SIGSPATIAL International Workshop on AI for Geographic Knowledge Discovery* (Chicago, IL, USA) *(GeoAI 2019)*. Association for Computing Machinery, New York, NY, USA, 56–65. https://doi.org/10.1145/3356471.3365238

[34] Stephen Law, Brooks Paige, and Chris Russell. 2019. Take a look around: using street view and satellite images to estimate house prices. *ACM Transactions on Intelligent Systems and Technology (TIST)* 10, 5 (2019), 1–19.

[35] Géraud Le Falher, Aristides Gionis, and Michael Mathioudakis. 2015. Where is the Soho of Rome? Measures and algorithms for finding similar neighborhoods in cities. In *Ninth International AAAI Conference on Web and Social Media*.

[36] Rémi Louf and Marc Barthelemy. 2014. A typology of street patterns. *Journal of The Royal Society Interface* 11, 101 (2014).

[37] Kevin Lynch. 1984. *Good city form*. MIT press.

[38] Adyasha Maharana and Elaine Okanyene Nsoesie. 2018. Use of deep learning to examine the association of the built environment with prevalence of neighborhood adult obesity. *JAMA network open* 1, 4 (2018), e181535–e181535.

[39] Magdalena Main-Knorn, Bringfried Pflug, Jerome Louis, Vincent Debaecker, Uwe Müller-Wilm, and Ferran Gascon. 2017. Sen2Cor for sentinel-2. In *Image and Signal Processing for Remote Sensing XXIII*, Vol. 10427. International Society for Optics and Photonics, 1042704.

[40] Kevin Manaugh and Tyler Kreider. 2013. What is mixed use? Presenting an interaction method for measuring land use mix. *Journal of Transport and Land Use* 6, 1 (2013), 63–72.

[41] Jonathan Masci, Ueli Meier, Dan Cireşan, and Jürgen Schmidhuber. 2011. Stacked Convolutional Auto-encoders for Hierarchical Feature Extraction. In *Proceedings of the 21th International Conference on Artificial Neural Networks - Volume Part I* (Espoo, Finland) *(ICANN'11)*. Springer-Verlag, Berlin, Heidelberg, 52–59. http://dl.acm.org/citation.cfm?id=2029556.2029563

[42] John Montgomery. 1998. Making a city: Urbanity, vitality and urban design. *Journal of urban design* 3, 1 (1998), 93–116.

[43] Franklin Obeng-Odoom. 2015. The social, spatial, and economic roots of urban inequality in Africa: Contextualizing Jane Jacobs and Henry George. *American Journal of Economics and Sociology* 74, 3 (2015), 550–586.

[44] Fernando S Peregrino, David Tomás, Paul Clough, and Fernando Llopis. 2012. Mapping routes of sentiments. In *Spanish Conference on Information Retrieval*.

[45] Santi Phithakkitnukoon, Zbigniew Smoreda, and Patrick Olivier. 2012. Socio-geography of human mobility: A study using longitudinal mobile phone data. *PloS one* 7, 6 (2012), e39253.

[46] Daniele Quercia, Rossano Schifanella, and Luca Maria Aiello. 2014. The shortest path to happiness: Recommending beautiful, quiet, and happy routes in the city. In *Proceedings of the 25th ACM conference on Hypertext and social media*. 116–125.

[47] Miriam Redi, Luca Maria Aiello, Rossano Schifanella, and Daniele Quercia. 2018. The spirit of the city: Using social media to capture neighborhood ambiance. *Proceedings of the ACM on Human-Computer Interaction* 2, CSCW (2018), 1–18.

[48] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. 2014. ImageNet Large Scale Visual Recognition Challenge. arXiv:cs.CV/1409.0575

[49] Wojciech Samek and Klaus-Robert Müller. 2019. Towards explainable artificial intelligence. In *Explainable AI: Interpreting, Explaining and Visualizing Deep Learning*. Springer, 5–22.

[50] Sanja Šćepanović, Igor Mishkovski, Pan Hui, Jukka K Nurminen, and Antti Ylä-Jääski. 2015. Mobile phone call data as a regional socio-economic proxy indicator. *PloS one* 10, 4 (2015), e0124160.

[51] Randy Showstack. 2014. Sentinel satellites initiate new era in earth observation. *Eos, Transactions American Geophysical Union* 95, 26 (2014), 239–240.

[52] Karen Simonyan and Andrew Zisserman. 2014. Very Deep Convolutional Networks for Large-Scale Image Recognition. *CoRR* abs/1409.1556 (2014).

[53] Karen Simonyan and Andrew Zisserman. 2014. Very Deep Convolutional Networks for Large-Scale Image Recognition. arXiv:cs.CV/1409.1556

[54] Blaž Sovdat, Miha Kadunc, Matej Batič, and Grega Milčinski. 2019. Natural color representation of Sentinel-2 data. *Remote sensing of environment* 225 (2019), 392–402.

[55] John Steenbruggen, Maria Teresa Borzacchiello, Peter Nijkamp, and Henk Scholten. 2013. Mobile phone data from GSM networks for traffic parameter and urban spatial pattern assessment: a review of applications and opportunities. *GeoJournal* 78, 2 (2013), 223–243.

[56] Hyungun Sung, Sugie Lee, and SangHyun Cheon. 2015. Operationalizing jane jacobs's urban design theory: Empirical verification from the great city of seoul, korea. *Journal of Planning Education and Research* 35, 2 (2015), 117–130.

[57] Kevin Tang, Manohar Paluri, Li Fei-Fei, Rob Fergus, and Lubomir Bourdev. 2015. Improving Image Classification with Location Context. *2015 IEEE International Conference on Computer Vision (ICCV)* (Dec 2015). https://doi.org/10.1109/iccv.2015.121

[58] R. Torres, P. Snoeij, D. Geudtner, D. Bibby, M. Davidson, E. Attema, P. Potin, B. Rommen, N. Floury, M. Brown, I.N. Traver, P. Deghaye, B. Duesmann, B. Rosich, N. Miranda, C. Bruno, M. L'Abbate, R. Croci, A. Pietropaolo, M. Huchler, and F. Rostan. 2012. GMES Sentinel-1 mission. *Remote Sensing of Environment* 120 (2012), 9–24. https://doi.org/10.1016/j.rse.2011.05.028

[59] M. Traunmueller, G. Quattrone, and L. Capra. 2014. Mining mobile phone data to investigate urban crime theories at scale. In *International Conference on Social Informatics*. 396–411.

[60] Steven Van Canneyt, Steven Schockaert, Olivier Van Laere, and Bart Dhoedt. 2011. Time-dependent recommendation of tourist attractions using Flickr. In *23rd benelux conference on artificial intelligence (bnaic 2011)*.

[61] Alessandro Venerandi, Giovanni Quattrone, Licia Capra, Daniele Quercia, and Diego Saez-Trumper. 2015. Measuring urban deprivation from user generated content. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*. 254–264.

[62] Wenshan Wang, Su Yang, Zhiyuan He, Minjie Wang, Jiulong Zhang, and Weishan Zhang. 2018. Urban perception of commercial activeness from satellite images and streetscapes. In *Companion Proceedings of the The Web Conference 2018*. 647–654.

[63] Hyoseok Yoon, Yu Zheng, Xing Xie, and Woontack Woo. 2012. Social itinerary recommendation from user-generated digital trails. *Personal and Ubiquitous Computing* 16, 5 (2012), 469–484.