Department of Land Economy

Environment, Law & Economics

UNIVERSITY OF CAMBRIDGE

Real Estate Research Centre

Working Paper Series

No. 2019-06

Title: Machine Learning, Building Vintage and Property Values

Authors: Thies Lindenthal[a], Erik B Johnson[b]

Affiliation: University of Cambridge[a], University of Alabama[b]

Contact corresponding author: Thies Lindenthal, htl24@cam.ac.uk

Machine Learning, Building Vintage and Property Values

Thies Lindenthal (University of Cambridge) & Erik B. Johnson (University of Alabama)

20 January 2019

Abstract

This paper makes three contributions: First, it introduces an algorithm that collects pictures of individual buildings from Google Street View. Second, it trains a deep convolutional neural network (CNN) to classify residential buildings into architectural styles, taking into account spatial dependencies of building vintages. Third, it investigates whether architectural styles influence house prices. For re-sales, the architectural style is a significant determinant of transaction prices while no such effect is found for new buildings. This indicates that any premia are for vintage-related quality characteristics and not for a home's beauty.

Introduction

If new developments were only pleasing to the eye, so Britain's *Building Better, Building Beautiful Commission*, nimbyism would cease and housing supply could finally reach the levels demanded by a growing and more affluent population (The Economist 2018). The inaugural chair of this new governmental commission boldly suggests that we should *"build, as our Georgian and Victorian forebears built* [. . .] *All objections to new building would slip away in the sheer relief of the public"* (Scruton 2018). Even Prince Charles, in similar spirit, put forward ten principles for urban growth and architecture that emphasize tradition and aesthetics (HRH The Prince of Wales 2014). When governments and princes occupy themselves with beauty, economists surely may take a closer look as well. In this paper, we aim to identify the architectural styles of residential buildings using computer vision techniques and to empirically search for any transaction price premia associated with these styles. If aesthetics were as strong a force in the built environment as claimed, we should indeed find a significant effect on property prices.

Corresponding author: Lindenthal (htl24@cam.ac.uk). Replication files are available from the author's Github repository https://github.com/thies. Paul E. Glade and Lukas Heckmann-Umhau are thanked for exceptional research assistance. Mike Langen, Colin Lizieri, Franz Fuerst, Carolin Schmidt, Stanimira Milcheva, D'Maris Coffman and participants at a UCL/Bartlett seminar offered excellent feedback and critique on earlier versions of this paper.

Recent work illustrates how street level imagery and machine learning classification is an efficient and powerful combination for measuring previously unobserved characteristics of the urban environment. Naik et al. (2017) describe how neighborhood demographics may impact the physical appearance of neighborhoods. Gebru et al. (2017) use classified vehicle make and model to predict income, race, education, and voting patterns at the precinct level. Glaeser et al. (2018) predict income in New York City. Naik, Raskar, and Hidalgo (2016) create a neighborhood safety based Streetscore which is shown to be highly correlated with neighborhood population density and household income. De Nadai et al. (2016) find that greenery and street facing windows contribute to a positive appearance of safety while Liu et al. (2017) evaluate the quality and upkeep of the built environment along Beijing's streets.

E. Glaeser, Kincaid, and Naik (2018) push the level of observation from the block, street, or street-section level to the individual *building level*. Utilizing images of buildings' exteriors collected from Google Street View, and to a lesser degree interior images from Zillow, they find that looks matter, at least in Boston: A one standard deviation improvement of a building's exterior is associated with and additional USD 70,000 in home value. Very intuitively, the link between good looks and value is bi-directional: The appearance of buildings that went through foreclosure deteriorated significantly (E. Glaeser, Kincaid, and Naik 2018).

Zooming in at individual buildings significantly increases the benefits of using mass collected street level imagery in economic research: property characteristics previously deemed "unobservable" can be directly observed in an accurate, objective, automatic and therefore cost effective way. Deriving additional variables from unconventional data sources like 3D airborne laser scanning (or in our case Google Street View) is fundamental since this provides "essential determinants influencing real estate prices [which] are constantly missing and are not accessible in official and mass appraiser databases" (Helbich et al. 2013).

In this short paper we suggest a method to collect a large number of images of individual UK buildings from Google Street View, classify the depicted buildings using deep convolutional neural networks, combine the derived information with sales price data and, ultimately, estimate marginal prices for the estimated building characteristics. Our work demonstrates that utilizing pretrained convolutional neural networks to detect complex characteristics such as housing vintages with relatively high accuracy comes at low computational costs and is feasible with only modestly-sized training data. No supercomputers are needed, a modern laptop will do.

https://www.google.co.uk/maps

Building vintage and property values

The aesthetic properties of different architectural eras make for an easy cocktail party conversation topic – and they also influence market prices. Using architectural assessments by human experts, Buitelaar and Schilder (2017) find a sizeable premium of 5% for new buildings in the Netherlands that refer to traditional styles and a staggering 15% premium for new buildings that closely follow traditional shapes, facade composition and details (Buitelaar and Schilder 2017). Their study carefully disentangles the architectural style from other unobserved characteristics such as building quality, differences in location or year of construction. This is crucial, as earlier work has established that age and vintage variables tend to be highly correlated. Coulson and McMillen (2008), for instance, suggest a non-parametric estimator and establish a U-shaped age function and distinct price discounts for postwar and contemporary vintages (vis-à-vis more historic vintages). Francke and Minne (2017) investigates the depreciation of residential real estate in the Netherlands and decompose land versus structure values and single out the effect of "physical deterioration, functional obsolescence and and vintage effects" (Francke and Minne 2017). They find that buildings from the 1930s carry a strong price premium.

A large scale assessment of buildings' exteriors will allow for an analysis of the externalities of architecture. Buildings hardly ever stand in isolation and Ahlfeldt and Mastro (2012) investigate the influence a building's architecture exerts on its surroundings. They observe a positive price effect for residential buildings in the direct proximity of iconic homes by Frank Lloyd Wright in Oak Park, Illinois. A building's exterior does not need to be an architectural masterpiece to co-determine the value of other houses close by. Homogeneity of building shapes within street segments does increase influence property values. A similarly shaped neighboring building is value enhancing while proximity to a wildly different neighboring shape, everything else remaining equal, is detrimental to property values (Lindenthal 2017a).

Unfortunately, the traditional approaches chosen by e.g. Buitelaar and Schilder (2017) or Ahlfeldt and Mastro (2012) do not scale well. Each observation has to be classified into architectural styles by a human expert which is time-consuming and costly, imposing an upper limit on the number of observations and the level of detail captured for each observation in any given sample. Google Street View offers a solution to this nexus as it captures the images of almost all buildings in many cities around the world at high level of accuracy and detail. The challenge we address in this paper is to extract and utilize building level information from this ubiquitous sensor.

Methodology

Image collection

The first challenge we faced when trying to collect building images in the UK was fundamental: How can we identify the building of interest correctly? Simply looking up an address on Google Maps too often leads to imprecise 'street views' and not the building-level portraits needed. Fig. 1 presents a typical result from an address level search, showing a broad ensemble of buildings instead of singling out the building of interest (in this case, the partly captured terraced house at the very left margin).

– Insert Figure 1 about here –

Previous work sourcing imagery from Google Street View has mostly focused on neighborhood or precinct characteristics where exact spatial assignment of objects is not required and street-level images suffice. For major US cities, the accuracy of building image search results is higher, which should help studies such as E. Glaeser, Kincaid, and Naik (2018).

Self-evidently, the images we use for classifying architectural style must be focused on the specific building of interest. For the UK, and many other countries, the Google Street View API returns the coordinates of the nearest camera snapshot for a given location but fails to provide an accurate orientation and zoom-level of the camera needed to capture the front of the building. To estimate the correct camera orientation, we combine Street View metadata queries and Ordnance Survey maps to find the optimal camera orientation to center images on the front door of the building of interest.

– Insert Figure 2 about here –

Specifically, we first look up the nearest Google Street View panorama point (green dot in Figure 2) based on the centroid (red dot) coordinates of a given building obtained from Ordnance Survey maps. We then perform a viewshed analysis and identify which exterior walls are visible from the panorama point, ignoring any wall segments where the direct line of sight from the panorama point is obstructed by other buildings. We can then estimate the camera bearing (green line) and zoom factor, based on the angle of the most outer of the blue lines of sight.

Obstructed lines of sight due to greenery, fences, garden walls or large vehicles cannot be detected from the Ordnance Survey maps. We therefore use a first stage image classification procedure to identify if the building image is obstructed and we proceed the second-closest location if the line of sight is obstructed.

Code necessary to replicate this analysis is available from the author's GitHub.com code repositories: https://github.com/thies.

By combining land registry maps with the snapshot location, we are able to automatically build a set of building frontage images for approximately 46,000 properties in Cambridge (UK). The Ordnance Survey maps do not provide any use type classifications so we collect images on all buildings that have a ground plate between 40 and 250 m^2. This method, however, could easily be adapted to photograph buildings in locations for which high-resolution building level maps are not available by using LIDAR based building outlines instead, which would be in the spirit of Glaeser et al. (2018).

From pixels to computer vision

We download all pictures at the highest resolution offered by Google's Street View API (640x640 color pixels). We then obtain 2048-dimensional feature vectors for each picture, using the *Inception-v3* deep convolutional neural network (Szegedy et al. 2015). *Inception-v3* has been trained for the ImageNet Large Scale Visual Recognition Challenge (ILSVRC) which evaluates image classification and object detection algorithms for a wide range of objects. The pretrained classifications would allow us to identify pets, vehicles or people on the pictures – assessing architectural finesse, however, is beyond the canned classifiers' capabilities.

E. Glaeser, Kincaid, and Naik (2018) rely on an different ILSVRC competitor, *Resnet-101* (He et al. 2016). They reduce the extracted feature vectors to lower dimensionality (1024 to 100 dimensions) based on principal component analysis (PCA). We follow a different strategy and double the dimensionality by including the feature vector of the building's nearest neighbor. Doubling up allows us to model spatial dependencies in building styles, similar to spatially correlated land cover classifications in Ghimire, Rogan, and Miller (2010).

Image classification

Cambridge's houses can be classified into seven broad eras:

- *Georgian* (c1714–1837) houses feature key characteristics such as sash windows, fan lights above doors, the use of stucco on facades, often wrought work grilles, railings etc.
- In the *Early Victorian* era (c1837–c1870s), a growing taste for individualized embellishment led to the development of elaborate features such as carved barge boards or finials. The development of sheet glass led to sash windows becoming more affordable, and, increasingly, wider.
- In the *Late Victorian* era (c1870s–1901), bay windows became more and more widespread, and

https://developers.google.com/maps/documentation/streetview/intro

We are grateful to colleagues from the Architecture Department at the University of Cambridge to provide these general vintage descriptions.

increasingly substantial. Stylistic movements such as the Queen Anne revival style contributed richly ornamental details to the formal repertoire employed by designers. Stained glass became more popular. *Edwardian* architecture (1901-1910) tends to be less ornate than late Victorian architecture.

- The *Interwar* period (1918–1939) saw the cost of building construction fall, amidst a drive to provide better housing for the working classes. New housing types were being favored.
- The *Postwar* (1950-1980) era continued on this path, with an embrace of high-rise as well as low rise housing. Facades vary greatly between brick, tiling, pebbledash and render.
- The cut-off year for our *Contemporary* era to begin is 1980. *Faux Victorians* are contemporary buildings trying to emulate Victorian architecture.

Two final-year architectural students classified a large sub-sample of approximately 25,000 images from our data set of Cambridge houses. This is a much larger sample than needed in our case. In our case, each category requires less than 250 samples to reach almost fully diminished training accuracy for additional observations. We greatly exceed this number for the purpose of this paper so that we can compare the out of sample convolutional neural network predictions to the groundtruth as assigned by the experts. This allows us to examine the power and size of the assignment tests. In addition having both human and machine classification for a large sample of the data allows for a robustness checks on the machine comparisons.

We create stratified training samples of 600 buildings from each category but *Georgian*, for which we only sample 330 examples. This leaves a sufficiently sized out-of-sample verification data set.

We train various versions of softmax classifiers with seven categories, ultimately settling for a parsimonious setup combining

- an input layer the size of the feature vectors,
- one dense layer (relu) half the size of the input layer, one subsequent dropout (rate 0.5) layer,
- one dense layer (relu) a quarter the size of the input layer, one subsequent dropout (rate 0.5) layer,
- and the final dense output layer with softmax activation.

All classifiers are implemented using the Keras/Tensorflow APIs. The computational burden of this rather shallow model design is modest.

For 100 randomly drawn, stratified data sets, we train two models: The first is based on the building specific feature vectors (2048 dimensions) only while the second also incorporates the feature vector of the nearest

Keras: https://keras.io/, Tensorflow: https://github.com/tensorflow/tensorflow neighbor (combined 4096 dimensions). In cases where trees, fences or large vehicles obstruct the view or where the front of a building cannot be classified with high certainty, the image of the nearest neigbor provides a "second opinion" that is spatially correlated with the observation of interest. Finally we classify all buildings and evaluate each model's out-of-sample performance in terms of recall, precision and F_1-scores.

Building era and property price

Merging the estimated vintage classification with sales data we estimate a hedonic regression equation that establishes marginal prices for the building vintage (similar to Moorhouse and Smith 1994; Asabere, Hachey, and Grubaugh 1989; Vandell and Lane 1989; Fuerst, McAllister, and Murray 2011; Plaut and Uzulena 2006), among other characteristics:

$$ln(Price_i) = \alpha + \beta \mathbf{X_i} + \delta \mathbf{Vintage}_i + \eta \mathbf{VintNeigh}_i + \gamma \mathbf{Y}_i + \lambda \mathbf{Loc}_i + \epsilon_i \tag{1}$$

Here, the natural logarithm of sales prices is explained by a linear combination of hedonic attributes described in vector $\boldsymbol{X}$, vectors of year $\boldsymbol{Y}$ and neighborhood $\boldsymbol{Loc}$ dummy variables and the building's estimated $\boldsymbol{Vintage}$ and the prevailing vintage of other buildings in the direct proximity ($VintNeigh$). The intercept is denoted by α while $\boldsymbol{\beta}$, $\boldsymbol{\delta}$, $\boldsymbol{\eta}$, γ and λ are vectors of regression coefficients. ϵ is the IID error term. Heteroscedastic robust standard errors will be reported.

Are buildings from different eras imperfect substitutes catering to multiple groups of households with distinct vintage preferences? Buitelaar and Schilder (2017) indicate that any premium for a architectural style must stem from either differences in construction prices (which they do not find in their Dutch sample) or from supply constraints, as new construction potentially does not capture the demand for traditional styles. For Cambridge, new supply will inevitably be of either Contemporary or Faux Victorian style as historic vintage buildings, by definition, are not supplied any more. Estimating Eq. 1 for a subset of newly constructed buildings will show whether construction prices or supply constraints for new homes built according to different architectural styles persist – if too few vernacular buildings were built, prices should reflect such a shortage.

For repeated sales observations of home i, both observed and unobserved quality characteristics will cancel out, as long as they are not time-varying:

$$ln(Price_{i,t_2}) - ln(Price_{i,t_1}) = \gamma(\mathbf{Y}_{i,t_2} - \mathbf{Y}_{i,t_1}) + \delta(\mathbf{A}_{i,t_2} - \mathbf{A}_{i,t_1}) + \eta \mathbf{Vintage}_i + \lambda \mathbf{Loc}_i + \epsilon_i \tag{2}$$

The change in general market prices is captured by the difference of the year dummies Y and the estimated vector of coefficients γ can be directly translated in a constant quality market price index (e.g. $Index = exp(\gamma) * 100$)).

The estimated building vintages **Vintage** do not change between transactions and would therefore cancel out, but we add them back to explicitly investigate return differences across vintages. Any differences in the estimated coefficients η indicate distinct price trends for vintages, after controlling for location **Loc**.

To non-parametrically capture the decay of buildings between t_1 and t_2 we add dummies for building ages at both the purchase and the sale date. The coefficients δ translate into the age related decay net of capital expenditures and upkeep.

Data

Residential real estate transactions are public data in the UK, collected and published in by the Land Registry (Land Registry 2017). The records include the date of transaction, price paid, street address, a classification of the property type (flat, detached, semi-detached, or terraced house), the estate type (freehold or leasehold) and an indicator for newly build properties. We select transactions from Cambridge proper which were recorded between January 1995 and October 2018, excluding any leaseholds, flats and properties classified as type "other", and sales with prices below £50,000 or in excess of £2,000,000. Table 3 presents summary statistics for sample.

The Ordnance Survey *AddressBase* (Ordnance Survey 2017a) links street addresses to building outlines on Ordnance Survey maps (Ordnance Survey 2017b), which allows us to calculate the buildings floor plate area (in m^2) and to estimate the building's volume from digital elevation models (Environment Agency 2015), as suggested by Lindenthal (2017b). We control for the location of each building by calculating the distance to the city center proxied by Great St. Mary's Church, and non-parametrically by using 69 indicator variables for each of the smallest census tracts (*Lower Super Output Areas*, LSOA) subdividing Cambridge (Office for National Statistics 2017b). LSOAs typically have 1,000–3,000 residents and 400–1,200 households of comparable economic and socio-demographic characteristics (Office for National Statistics 2017a).

Overall, we collect pictures of more than 48,000 distinct buildings from the Google Street View API. A subset of 25,000 picture has been classified by trained architects into seven vintages and will serve as *ground truth* for our study. Two additional categories identify buildings hidden by *greenery* or uninformative pictures featuring *plain walls*. The classified images are available at the authors' websites.

Results

Almost 16,000 of the 25,000 buildings for which we have the architects' ground truth could be matched to the 23,768 sales transactions that have been recorded for Cambridge between 1995 and 2018. Thus, we observe both the machine *estimated vintage* and architect defined *ground truth* classification for the majority of buildings in the data set.

When classifying a building picture, our models return a vector of scores between 0 to 1 that jointly sum up to 1. We select the vintage with the highest score as the best representation of the vintage of the building. By not excluding observations where multiple scores are vying for the top rank we retain as many observations as possible but risk a higher misclassification rate.

– Insert Table 1 about here –

The match rate for estimated vintages and the corresponding expert classifications are surprisingly strong. When relying on feature vectors of the buildings and their closest neighbors, on average 61 to 74 percent of true classifications are matched with correct estimated labels (1). The recall rates are especially high for older vintages: For *Georgians*, it is 74%, *Early Victorian* 74%, *Late Victorian/Edwardian* 72%, *Interwar* 74%, respectively. For more recent *Postwar* buildings, the rate drops to 61%, *Contemporary* 62% and *Faux Victorian* 61%.

The precision rate, however, for the relatively rare categories of *Georgian* and *Faux Victorian* is low, with more buildings classified as e.g. *Georgian* belonging to a different ground truth class than buildings from the true *Georgians* stock. The F_1-scores combine the recall and precision metrics.

A visual inspection of classifications is also reassuring. Figure 3 displays the pictures which carry the highest out-of-sample scores for each vintage. The model is clearly able to differentiate based on subtle cues, even when only parts of the facade are captured on the picture.

– Insert Figure 3 about here –

The majority of misclassification is concentrated in temporally adjacent eras: 14% of *Late Victorian/Edwardian* buildings are labelled as *Early Victorian* or 21% of *Postwar* buildings are erroneously regarded as stemming from the *Interwar* period. Hardly ever is a contemporary building mistaken for a historic home.

Overall, the 'spatial' softmax classification incorporating the nearest neighbor's feature vector outperforms the base classifier. Fig. **??** displays estimated density functions from the F_1-scores from the 100 trained models for the spatial model (solid lines) and the standard models (dotted lines). Fig. 4 confirms this finding:

pairwise differences of the F_1-scores from the 'spatial' and the base classifier are clearly positive. In sum, for spatially correlated dependent variables we recommend to additionally utilize feature vectors from other observations close by.

– Insert Figure **??** and Figure 4 about here –

Hedonic regression estimates

We estimate three different versions of the hedonic regression specified in Eq. (1) and reported the coefficient estimates in Table 4. For all models, the hedonic control variables show the expected signs: Negative coefficients for the relative distance to the city center, discounts for terraced homes and semi-detached homes relative to detached houses and positive elasticities for building floor plate and building volumes. Model 1 and 3 show a positive marginal prices for new buildings compared to second-hand homes – Model 2 is estimated on new homes only. Year and neighborhood dummies control for time effects and local amenities but their coefficients are not reported due to space constraints.

– Insert Table 3 about here –

The first column in Table 4 presents the estimated regression coefficients for dummy variables on building vintages that have been estimated with a 'spatial' classifier. The base vintage *Contemporary* is more expensive than almost all other vintages which show negative coefficients that are significantly different from 0. A clear pecking order appears: *Late Victorian/Edwardian* buildings are the dearest (+0.03), followed by *Contemporary* (base), *Georgian* and *Faux Victorian* (−0.03), *Early Victorian* (−0.08), *Interwar* (−0.09) and *Postwar* (−0.12).

Model 2 estimates the same equation with a subset of recent buildings (built since 1995). This model basically contrasts buildings which might *look* different - but which are *new buildings* at their core. Differences in e.g. materials, floor plans green space and gardens are minimal. For this subsample, we do not find *any* effect of the estimated styles on price. After controlling for location and building characteristics and quality, buyers show no willingness to pay a premium for any particular style.

Extending the model with variable capturing the most prominent style of buildings on the same street and within 100 m (**VintNeigh**), we find the price premia on estimated vintages basically unchanged. The coefficients on *Georgian* and *Faux Victorian* styles cease to be statistically different from *Contemporary* buildings, though. All coefficients on neighboring styles are negative compared to the base case. This

The combination of location dummies and the distance to the city center measure controls for proximity to the city center *within* each neighborhood.

suggest that, everything else equal, buildings close to ensembles of *Contemporary* buildings achieve higher transaction prices than buildings in more historic areas. Not surprisingly, *Postwar* buildings are the least popular neighbors. The difference between *Faux Victorian* neighbors carry a negative coefficient of -0.05 – again, we fail to find evidence in favor of building according to vernacular architecture.

A market wide price index calculated from time dummy variables in Eq. (2) reveal a strong upward trend in Cambridge's residential property market (Figure 5). Between 1995 and 2017, the transaction prices for constant quality homes across all styles have risen to 588% of 1995's value. Distinct difference in time trends per vintage become discernible (Table **??**) with buildings from historic vintages outperforming more contemporary homes. *Late Victorian/Edwardian* and *Interwar* homes gained most value ($exp(+0.05)$), followed by *Early Victorian.* Prices for *Contemporary* buildings grew slower than prices for vernacular architecture. That difference in price dynamics is statistically not significant, however.

– Insert Figure 5 and Figure **??** about here –

Model 2 in Table **??** explores the net decay of new buildings in the first 25 years after completion. Overall, we find values of new buildings to grow slower during their first year than prices for older buildings. The underperformance is significant: The negative -0.13 coefficient for the *Age 20-24* variable indicates that return is 12 percentage points lower ($exp(0.13) - 1 = -12.2$). In the first years, new buildings age badly, apparently. Fig. **??** displays the relative decay of buildings in their first 25 years. The dashed line visualizes the decay for *Contemporary*-style new buildings compared to new buildings in all other styles (solid line) as estimated in Model 3, Table **??**. The difference between the two decay rates is statistically not significant. The lack of significance for the interaction terms for *New cont. Age 15–19* and *New cont. Age 20–24* is due to low number of observations in these categories.

Conclusion

The contributions of this paper are threefold: First, it introduced an algorithm that collects pictures of *individual* buildings from Google Street View. Earlier work has not achieved this level of detail and was, at least in the UK, limited to street sections only. A large-scale application of automatic classification of individual buildings' characteristics using Google Street View has potential not only in the UK. The image collection and classification method can easily be ported to other study areas which have existing Street View data and either LIDAR based building outlines or high resolution satellite images. In a follow-up project, we

These estimates of decay are not too far off from decay schedules estimated for commercial real estate in the US by @Bokhari2016. Any comparison of decay rates across sectors and countries has clear limits, obviously.

are working on an improved classification workflow in which a customized object detection model is able to recognize specific building outlines without the need for additional maps or other data.

Second, we developed a new database of 25,000 building pictures that have been classified by architecture experts into relevant architectural styles or vintages. We subsequently trained a neural network classifier to automatically classify all residential buildings of a mid-sized English city into architectural styles or vintages. The suggested classifier is trained on feature vectors of buildings and their nearest neighbors to exploit spatial correlation in observed classifications.

Third, we test whether revival or vernacular architecture leads to price premia over contemporary architecture. We do not find evidence for a preference of residents for specific architecture after controlling for quality and location of buildings. This applies to both direct price effects but also indirect effects of a buildings appearance on the value of neighboring homes. Further work will extend the spatial scope of the study area to the UK in general.

Tables and Figures

Figure 1: ‘Street View’, not ‘Building View’ – Identification of buildings remains a challenge in the UK

Notes: For the UK, the Google Street View API returns the coordinates of the nearest camera snapshot for a given location but fails to provide an accurate orientation and zoom-level of the camera needed to capture the front of the building exactly. In this typical example, the building of interest is only partially shown at the very left margin of this result. *Image source*: Google Street View.

Figure 2: Image Collection on Google Street View: Camera Direction and Zoom

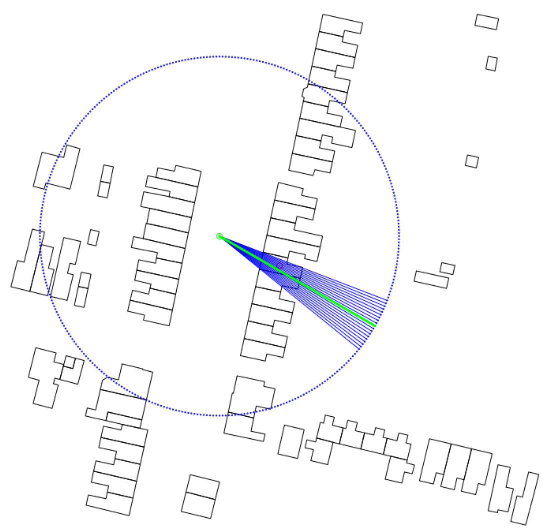

Notes: We first look up the nearest Google Street View panorama point (green dot) based on the centroid (red dot) coordinates of a given building obtained from Ordnance Survey maps. A viewshed analysis identifies which exterior walls are visible from the panorama point, ignoring any wall segments where the direct line of sight from the panorama point is obstructed by other buildings. The camera bearing (green line) and zoom factor are based on the angle of the most outer lines of sight (blue lines).

Table 1: Confusion matrix for classification based on images of property and its nearest neighbor

Panel A: Model based on individual building and nearest neighbor

Predicted	*Ground truth*						
	Georgian	Early Vic.	Late V./Edw.	Interwar	Postwar	Cont.	Revival
Georgian	284	109	77	38	14	23	22
Early Vic.	50	1755	427	86	86	56	34
Late V./Edw.	10	172	3260	213	56	21	29
Interwar	10	46	254	5884	997	56	54
Postwar	1	17	48	514	3914	74	40
Cont.	3	50	63	101	333	855	69
Revival	3	29	45	139	145	98	501
Georgian	79%	5%	2%	1%	0%	2%	3%
Early Vic.	14%	81%	10%	1%	2%	5%	5%
Late V./Edw.	3%	8%	78%	3%	1%	2%	4%
Interwar	3%	2%	6%	84%	18%	5%	7%
Postwar	0%	1%	1%	7%	71%	6%	5%
Cont.	1%	2%	2%	1%	6%	72%	9%
Revival	1%	1%	1%	2%	3%	8%	67%
Recall	0.79	0.81	0.78	0.84	0.71	0.72	0.67
Precision	0.50	0.70	0.87	0.81	0.85	0.58	0.52
F_1-score	0.61	0.75	0.82	0.82	0.77	0.64	0.59

Panel B: Model based on individual buildings only

Predicted	*Ground truth*						
	Georgian	Early Vic.	Late V./Edw.	Interwar	Postwar	Cont.	Revival
Georgian	267	121	110	46	33	36	17
Early.Vic.	50	1673	438	112	107	63	52
Late V./Edw.	19	183	3099	272	86	33	44
Interwar	12	71	335	5786	1419	84	82
Postwar	2	24	56	462	3364	75	38
Cont.	5	63	82	141	366	778	79
Revival	6	43	54	156	170	114	437
Georgian	74%	6%	3%	1%	1%	3%	2%
Early.Vic.	14%	77%	10%	2%	2%	5%	7%
Late V./Edw.	5%	8%	74%	4%	2%	3%	6%
Interwar	3%	3%	8%	83%	26%	7%	11%
Postwar	1%	1%	1%	7%	61%	6%	5%
Cont.	1%	3%	2%	2%	7%	66%	11%
Revival	2%	2%	1%	2%	3%	10%	58%
Recall	0.74	0.77	0.74	0.83	0.61	0.66	0.58
Precision	0.42	0.67	0.83	0.74	0.84	0.51	0.45
F_1-score	0.54	0.72	0.78	0.78	0.70	0.58	0.51

Notes: This table presents the distribution of out-of-sample classifications versus ground truth. The top panel presents counts, while the mid panel shows the share per true category. *Recall* is the share of buildings from a ground truth category being predicted correctly (diagonal in mid panel) and *Precision* is the share of buildings predicted to belong to a category that are indeed from that category. The F_1-score is the harmonious mean of Precision and Recall: F_1*-score = 2 Recall * Precision / (Recall + Precision).*

Figure 3: Examples of Estimated Building Vintages

Georg.	*Early Vic.*	*Vic./Edw.*	*Interwar*	*Postwar*	*Contemp.*	*Faux Vic.*

Notes: For a brief discussion of the building vintages see page 5.

Figure 4: Difference of F_1-scores: spatial model vs. base model

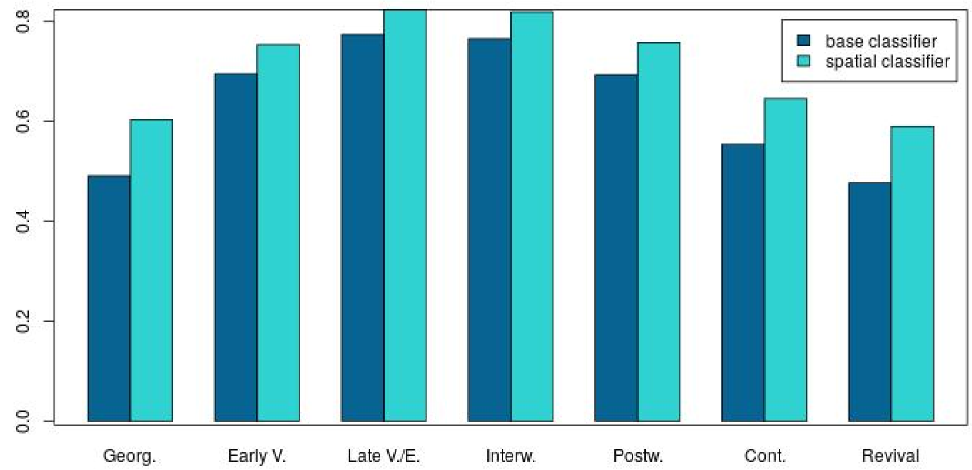

Notes: The barplot shows pairwise difference in F_1-scores per category for the classifier using both building level and neighbor information vs. a classifier using building level information only.

Table 2: Akaike information criterion (AIC) for reverse regression models

	No ML pred.	*Base pred.*	*Spatial pred.*
Georgian	1282	914	912
Early Victorian	10232	6209	5712
Late Vic./Edw.	16223	8863	8186
Interwar	18334	10009	8895
Postwar	16846	11228	9471
Contemporary	2885	1926	1743
Revival	3336	2255	2099

Notes: AIC values for 21 models in which...

Figure 5: Repeat Sales Price Index, Cambridge 1995-2018

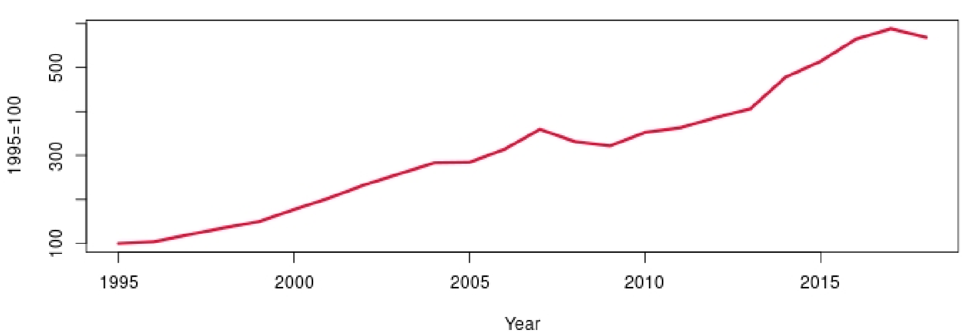

Notes: t.b.d.

Table 3: Summary statistics residential property transactions

Statistic	N	Mean	St. Dev.	Min	Pctl(25)	Pctl(75)	Max
Price	23,768	253,646.20	170,987.20	10,000	127,000	326,374.8	1,000,000
Year	23,768	2,004.89	6.59	1,995	1,999	2,010	2,018
ln(volume)	23,768	286.00	188.14	0.00	215.50	370.89	1,926.45
ln(area)	23,768	63.25	28.44	4.94	45.74	72.46	1,058.70
ln(dist. city center)	23,768	2,477.63	1,007.55	112.78	1,652.39	3,166.05	5,023.62
New	23,768	0.04	0.20	0	0	0	1
Type: detached	23,768	0.13	0.34	0	0	0	1
Type: semi-detached	23,768	0.35	0.48	0	0	1	1
Type: terraced	23,768	0.52	0.50	0	0	1	1
Georgian	23,768	0.03	0.16	0	0	0	1
Early Vic.	23,768	0.15	0.36	0	0	0	1
Late Vic./Edw.	23,768	0.21	0.40	0	0	0	1
Interwar	23,768	0.32	0.47	0	0	1	1
Postwar	23,768	0.23	0.42	0	0	0	1
Contemporary	23,768	0.03	0.17	0	0	0	1
Faux Vic.	23,768	0.04	0.20	0	0	0	1
Neigh: Georgian	23,768	0.02	0.12	0	0	0	1
Neigh: Early Vic.	23,768	0.14	0.35	0	0	0	1
Neigh: Late V./Edw.	23,768	0.22	0.41	0	0	0	1
Neigh: Interwar	23,768	0.32	0.47	0	0	1	1
Neigh: Postwar	23,768	0.25	0.43	0	0	1	1
Neigh: Contemporary	23,768	0.02	0.15	0	0	0	1
Neigh: Faux Vic.	23,768	0.03	0.18	0	0	0	1

Notes: Summary statistics for sample of 23,768 residential real estate transactions for the city of Cambridge (UK) between 1995 and 2018 where buildings could be matched with Google Street View images. The buildings floor plate area (in m^2) is based on OS maps and the buildings' volumes are estimated from digital elevation models, as suggested by Lindenthal (2017). We control for the location of each building by calculating the distance to the city center proxied by Great St. Mary's Church, and non-parametrically by using 69 location dummies (based on LSOA). The estimated scores for building vintages are inferred from the Google Street View classification model.

Table 4: Hedonic Regression Estimates

	Dependent variable: ln(price)					
	(1)	(2)	(3)	(4)	(5)	(6)
Constant	13.45***	13.20***	15.19***	13.10***	13.27***	13.47***
	(0.17)	(0.20)	(1.70)	(2.20)	(0.16)	(0.16)
ln(dist. city center)	−0.50***	−0.45***	−0.69***	−0.41	−0.46***	−0.49***
	(0.02)	(0.03)	(0.22)	(0.31)	(0.02)	(0.02)
Type: semi-detached	−0.14***	−0.12***	−0.19***	−0.10*	−0.14***	−0.14***
	(0.01)	(0.01)	(0.04)	(0.05)	(0.01)	(0.01)
Type: terraced	−0.20***	−0.19***	−0.16***	−0.12**	−0.21***	−0.21***
	(0.01)	(0.01)	(0.04)	(0.05)	(0.01)	(0.01)
ln(area)	0.40***	0.40***	0.38***	0.43***	0.40***	0.40***
	(0.01)	(0.01)	(0.07)	(0.08)	(0.01)	(0.01)
ln(volume)	0.01***	0.01***	0.004	0.002	0.01***	0.01***
	(0.001)	(0.001)	(0.01)	(0.01)	(0.001)	(0.001)
New	0.18***	0.09***			0.15***	0.15***
	(0.01)	(0.02)			(0.01)	(0.01)
	Base: Contemporary					
Georgian	0.06***	−0.04	0.21***	−0.22	0.06***	−0.02
	(0.02)	(0.03)	(0.06)	(0.28)	(0.02)	(0.03)
Early Vic.	−0.01	−0.13***	0.02	−0.28*	0.01	−0.12***
	(0.01)	(0.02)	(0.06)	(0.17)	(0.01)	(0.02)
Late Vic./Edw.	0.11***	0.01	0.02	−0.17	0.09***	0.04**
	(0.01)	(0.02)	(0.05)	(0.14)	(0.01)	(0.02)
Interwar	−0.02**	−0.15***	0.03	−0.22**	−0.003	−0.13***
	(0.01)	(0.02)	(0.04)	(0.10)	(0.01)	(0.02)
Postwar	−0.06***	−0.22***	0.05	−0.14	−0.03***	−0.19***
	(0.01)	(0.02)	(0.03)	(0.11)	(0.01)	(0.02)
Revival	0.05***	0.04**	0.09***	−0.14	0.06***	0.03
	(0.01)	(0.02)	(0.03)	(0.11)	(0.01)	(0.02)
	Base: Contemporary					
Neigh: Georgian					−0.05	
					(0.03)	
Neigh: Early Vic.					−0.12***	
					(0.02)	
Neigh: Late V./Edw.					−0.03*	
					(0.02)	
Neigh: Interwar					−0.10***	
					(0.01)	
Neigh: Postwar					−0.13***	
					(0.01)	
Neigh: Revival					−0.02	
					(0.02)	
Neigh. interact.	No	No	No	No	No	See next table
Year dummies	Yes	Yes	Yes	Yes	Yes	Yes
Neigh. dummies	Yes	Yes	Yes	Yes	Yes	Yes
Observations	23,415	15,642	930	463	23,256	23,415
Adjusted R^2	0.88	0.88	0.88	0.90	0.88	0.88

Notes: *p<0.1; **p<0.05; ***p<0.01. Standard errors are robust (White's estimator). Models (1) and (3) are estimated on sales of new and existing houses, while Model (2) is estimated on 930 sales of newly completed buildings only. The vintage of direct neighborhoods is defined as the most frequent detected style on the same street and within 100m. The coefficients for

Table 5: Coefficients for interaction terms: Building style and Neigh. style

Neigh.	*Building*						
	Georgian	Early Vic.	Late V./Edw.	Interwar	Postwar	Cont.	Revival
Georgian	–	-0.12*** (0.04)	0.09* (0.05)	-0.07 (0.04)	-0.12** (0.05)	-0.15 (0.41)	0.08* (0.04)
Early Vic.	0.01 (0.07)	–	0.06*** (0.01)	0.02 (0.02)	0.01 (0.02)	0.12** (0.06)	0.04 (0.03)
Late V./E.	0.15* (0.08)	-0.12*** (0.01)	–	-0.09*** (0.02)	-0.13*** (0.03)	-0.06 (0.08)	-0.14** (0.06)
Interwar	-0.01 (0.03)	-0.04* (0.02)	0.05*** (0.01)	–	-0.04*** (0.01)	0.02 (0.03)	0.10*** (0.03)
Postwar	-0.22*** (0.03)	0.03 (0.03)	0.11*** (0.03)	0.01 (0.01)	–	0.13*** (0.02)	0.09*** (0.02)
Cont.	0.22*** (0.03)	-0.12*** (0.03)	-0.11** (0.05)	-0.15*** (0.03)	-0.17*** (0.02)	–	-0.09** (0.05)
Revival	-0.09 (0.06)	-0.13*** (0.03)	-0.001 (0.03)	-0.13*** (0.02)	-0.14*** (0.02)	-0.03 (0.03)	–

Notes: *p<0.1; **p<0.05; ***p<0.01. Standard errors are robust (White's estimator). This table features the coefficients for interaction terms of building style and neighborhood styles only, while Table 4, Column 6, presents all other coefficients for this model.

Bibliography

Ahlfeldt, Gabriel, and Alexandra Mastro. 2012. "Valuing Iconic Design: Frank Lloyd Wright Architecture in Oak Park, Illinois." *Housing Studies* 27 (8): 1079–99. https://doi.org/10.1080/02673037.2012.728575.

Asabere, Paul K., George Hachey, and Steven Grubaugh. 1989. "Architecture , Historic Zoning , and the Value of Homes." *Journal of Real Estate Finance and Economics* 2 (3): 181–95. https://doi.org/10.1007/BF00152347.

Buitelaar, Edwin, and Frans Schilder. 2017. "The Economics of Style: Measuring the Price Effect of Neo-Traditional Architecture in Housing." *Real Estate Economics* 45 (1): 7–27. https://doi.org/10.1111/1540-6229.12137.

Coulson, N. Edward, and Daniel P. McMillen. 2008. "Estimating time, age and vintage effects in housing prices." *Journal of Housing Economics* 17 (2): 138–51. https://doi.org/10.1016/j.jhe.2008.03.002.

De Nadai, Marco, Radu L. Vieriu, Gloria Zen, Stefan Dragicevic, Nikhil Naik, Michele Caraviello, Cesar A. Hidalgo, Nicu Sebe, and Bruno Lepri. 2016. "Are Safer Looking Neighborhoods More Lively? A Multimodal Investigation into Urban Life," August. http://arxiv.org/abs/1608.00462.

Environment Agency. 2015. "LIDAR Composite DSM - 1m." https://data.gov.uk/dataset/lidar-composite-dsm-1m1.

Francke, Marc K., and Alex M. van de Minne. 2017. "Land, Structure and Depreciation." *Real Estate Economics* 45 (2): 415–51. https://doi.org/10.1111/1540-6229.12146.

Fuerst, Franz, Patrick McAllister, and Claudia B Murray. 2011. "Designer buildings: estimating the economic value of 'signature' architecture." *Environment and Planning A* 43 (1): 166–84. https://doi.org/10.1068/a43270.

Gebru, Timnit, Jonathan Krause, Yilun Wang, Duyun Chen, Jia Deng, Erez Lieberman Aiden, and Li Fei-Fei. 2017. "Using Deep Learning and Google Street View to Estimate the Demographic Makeup of the US," 1–41. http://arxiv.org/abs/1702.06683.

Ghimire, B., J. Rogan, and J. Miller. 2010. "Contextual land-cover classification: Incorporating spatial dependence in land-cover classification models using random forests and the Getis statistic." *Remote Sensing Letters* 1 (1): 45–54. https://doi.org/10.1080/01431160903252327.

Glaeser, Edward, Michael Scott Kincaid, and Nikhil Naik. 2018. "Computer Vision and Real Estate: Do Looks Matter and Do Incentives Determine Looks?" http://www.nber.org/papers/w25174.

Glaeser, Edward L., Scott Duke Kominers, Michael Luca, and Nikhil Naik. 2018. "Big Data and Big Cities: The Promises and Limitations of Improved Measures of Urban Life." *Economic Inquiry* 56 (1): 114–37. https://doi.org/10.1111/ecin.12364.

He, Kaiming, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. "Deep Residual Learning for Image Recognition." *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 770–78.

Helbich, Marco, Andreas Jochem, Werner Mücke, and Bernhard Höfle. 2013. "Boosting the predictive accuracy of urban hedonic house price models through airborne laser scanning." *Computers, Environment and Urban Systems* 39. Elsevier Ltd: 81–92. https://doi.org/10.1016/j.compenvurbsys.2013.01.001.

HRH The Prince of Wales. 2014. "Facing up to the future: Prince Charles on 21st century architecture." *Architectural Review*, December. https://www.architectural-review.com/essays/viewpoints/facing-up-to-the-future-prince-charles-on-21st-century-architecture/8674119.article.

Land Registry. 2017. "Price Paid Data." http://landregistry.data.gov.uk/app/ppd.

Lindenthal, Thies. 2017a. "Beauty in the Eye of the Home-Owner: Aesthetic Zoning and Residential Property Values." *Real Estate Economics*, 1–26. https://doi.org/10.1111/1540-6229.12204.

———. 2017b. *Estimating Supply Elasticities for Residential Real Estate in the United Kingdom.* Comprehens. Elsevier Inc. https://doi.org/https://doi.org/10.1016/B978-0-12-409548-9.09682-2.

Liu, Lun, Elisabete A. Silva, Chunyang Wu, and Hui Wang. 2017. "A machine learning-based method for the large-scale evaluation of the qualities of the urban environment." *Computers, Environment and Urban Systems* 65. The Authors: 113–25. https://doi.org/10.1016/j.compenvurbsys.2017.06.003.

Moorhouse, John C., and Margaret Supplee Smith. 1994. "The Market for Residential Architecture: 19th century Row Houses in Boston's South End." *Journal of Urban Economics* 35 (3): 267–77. https://doi.org/10.1006/juec.1994.1016.

Naik, Nikhil, Scott Duke Kominers, Ramesh Raskar, Edward L. Glaeser, and César A. Hidalgo. 2017. "Computer vision uncovers predictors of physical urban change." *Proceedings of the National Academy of Sciences* 114 (29): 7571–6. https://doi.org/10.1073/pnas.1619003114.

Naik, Nikhil, Ramesh Raskar, and César A. Hidalgo. 2016. "Cities Are Physical Too: Using Computer Vision to Measure the Quality and Impact of Urban Appearance." *American Economic Review* 106 (5): 128–32. https://doi.org/10.1257/aer.p20161030.

Office for National Statistics. 2017a. "A Beginner's Guide to UK Geography."

———. 2017b. “Open Geography Portal – 2011 Census Boundaries.” http://geoportal.statistics.gov.uk/datasets?q=COA{_}B

Ordnance Survey. 2017a. “AddressBase.” https://www.ordnancesurvey.co.uk/business-and-government/help-and-support/products/addressbase-premium.html.

———. 2017b. “OS Maps Local.” https://www.ordnancesurvey.co.uk/business-and-government/products/os-open-map-local.html.

Plaut, Steven, and Egita Uzulena. 2006. “Architectural Design and the Value of Housing in Riga, Latvia.” *International Real Estate Review* 9 (1): 112–31.

Scruton, Roger. 2018. “The Fabric of the City.” Colin Amery Memorial Lecture. Policy Exchange. https://policyexchange.org.uk/wp-content/uploads/2018/11/The-Fabric-of-the-City.pdf.

Szegedy, Christian, Vincent Vanhoucke, Sergey Ioffe, Jonathon Shlens, and Zbigniew Wojna. 2015. “Rethinking the Inception Architecture for Computer Vision.” https://doi.org/10.1109/CVPR.2016.308.

The Economist. 2018. “The line of beauty.” *The Economist*, November. London. https://www.economist.com/britain/2018/11/17/how-to-defeat-nimbyism-build-more-beautiful-houses.

Vandell, Kerry D., and Jonathan S. Lane. 1989. “The Economics of Architecture and Urban Design: Some Preliminary Findings.” *AREUEA Journal* 17 (2): 235–60.