Using Deep Learning To Infer House Prices From Google Street View Images

B. Chahal[1], M. Gutiérrez-Roig[2], T. Preis[2,3], H. S. Moat[2,3]

[1]EPSRC & MRC Centre for Doctoral Training in Mathematics for Real-World Systems, University of Warwick; [2]Warwick Business School, University of Warwick; [3]The Alan Turing Institute

Introduction

In the UK, especially London, house prices have been on the rise for decades. Therefore, obtaining a model that could predict house prices reasonably well is appealing to individuals and corporations alike. The model could aid them in making informed decisions regarding investments and relocation, for example. Intuitively, houses that are worth more are more aesthetically pleasing. However, including information on visuals in a computational model has been proven challenging due to a lack of large scale quantitative data. That was until recently. The availability of Google Street View images has increased over the past decade, and we wish to see if we can capture any new information from them.

Methodology

We ask two questions:
1. Can we infer current house prices from current Google Street View images?
2. Can we identify the areas which have seen relatively high increases in prices from them?

Figure 1 displays the methodology which consists of two parts. At the image level, we download images across London and extract values of features from Places Scene Recognition and SUN Attributes, both developed by MIT [1,2]. These values essentially tell us how likely a certain feature is to appear in a given image. At the MSOA level, where MSOAs are ONS-defined geographical areas, we average the feature values per MSOA and use them to predict median house prices in London from 2015-2016 and their relative change over the decade [3].

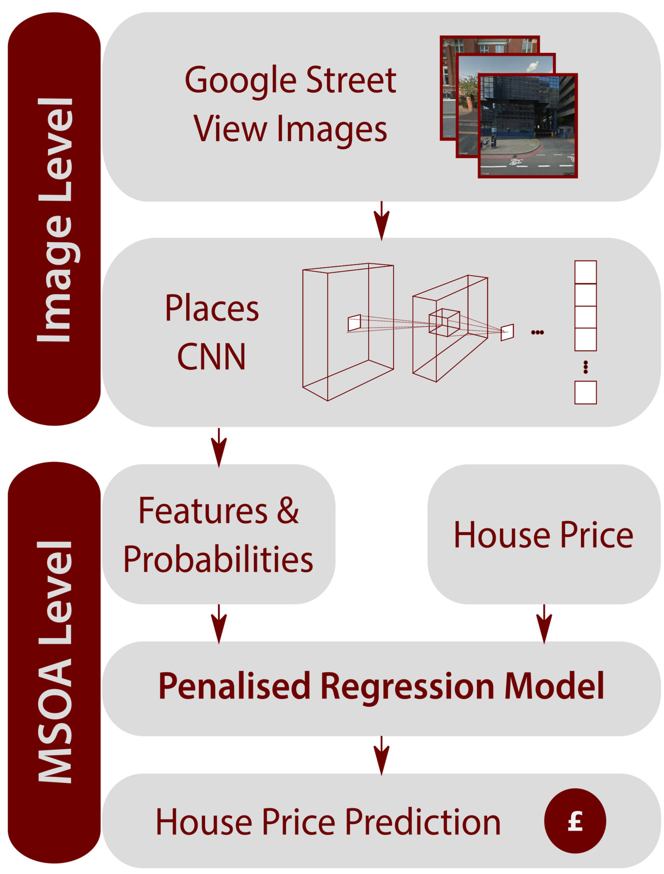

Figure 1: Methodology outline

2015-2016 Results

To infer current house prices, we use extracted feature values per MSOA and house price data from 2015-2016 in an elastic net model.

Figure 2 exhibits the top 20 features from the model that are positively and negatively associated with higher house prices. The coefficients make intuitive sense, as 'palace' is a positive coefficient and 'slum' is a negative coefficient. Figure 3 shows the observed and predicted median house prices from 2015-2016. The model identifies high median house prices in Central London but tends to underestimate them otherwise. Figure 4 displays the prediction errors. Notice that prices in the south east tend to be overestimated, and in the north west tend to be underestimated.

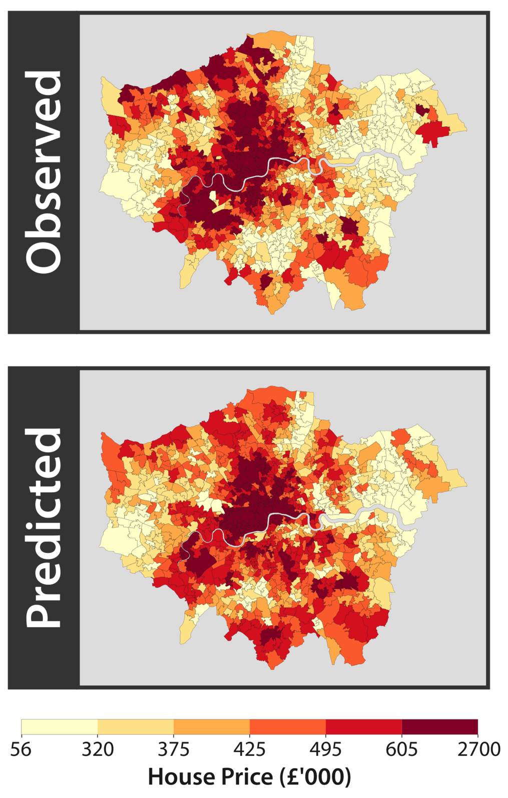

Figure 3: Observed and predicted median house prices

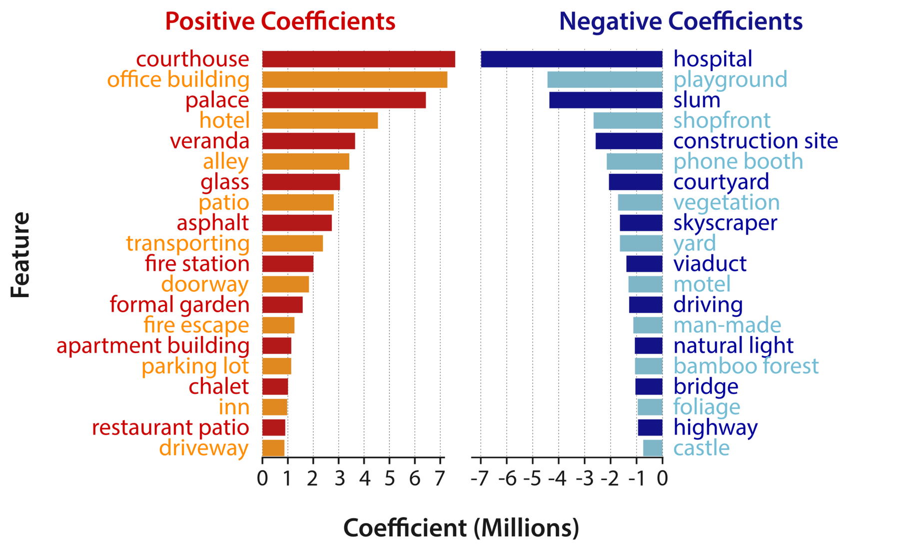

Figure 2: Top 20 model coefficients from current median house prices

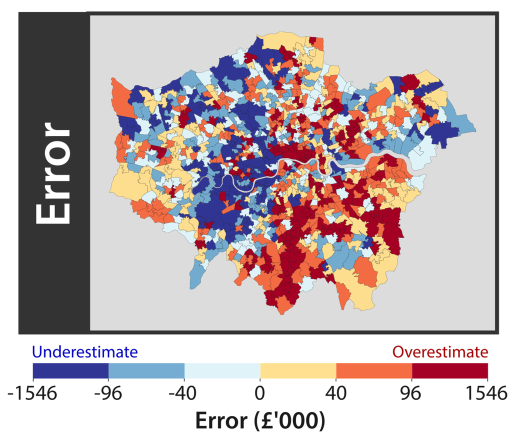

Figure 4: Prediction errors from current median house prices

Places CNN

Figure 5: Interesting features

Although many features associated with London in our model make intuitive sense, some are unexpected. Features outside of the top 20 that are associated with London include 'windmill' and 'rice paddy'. To investigate why this is the case, we explore the top 3 images labelled 'rainforest', 'windmill', 'rice paddy', and 'golf course', as shown in Fig. 5. According to the CNN, lush groups of trees are rainforests; large white spaces are windmills; green spaces behind fences are rice paddies; green spaces with trees are golf courses. The CNN identifies patterns and shapes in images and gives them labels, but the labels themselves are not always accurate.

2005-2016 Results

To infer the relative change in house prices over the past decade, we use extracted feature values per MSOA and calculated changes in relative rank from 2005-2006 and 2015-2016 in an elastic net model.

Figure 6 shows the top 20 features from the model that are positively and negatively associated with higher relative changes. Some coefficients have swapped. 'Hospital' was previously a negative coefficient and 'office building' was a positive coefficient. Figure 7 shows the observed and predicted relative change. There appears to be a center versus suburbs effect, which is overgeneralised by the model and shown in the prediction errors in Fig. 8.

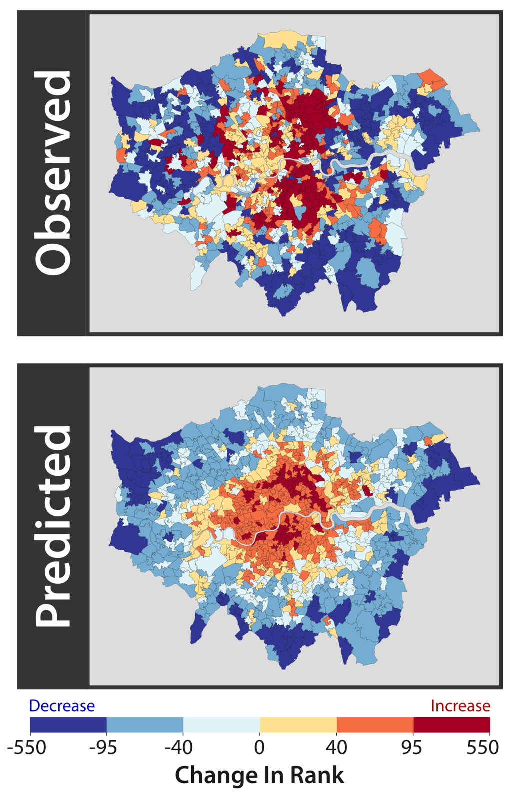

Figure 7: Observed and predicted relative change in house prices

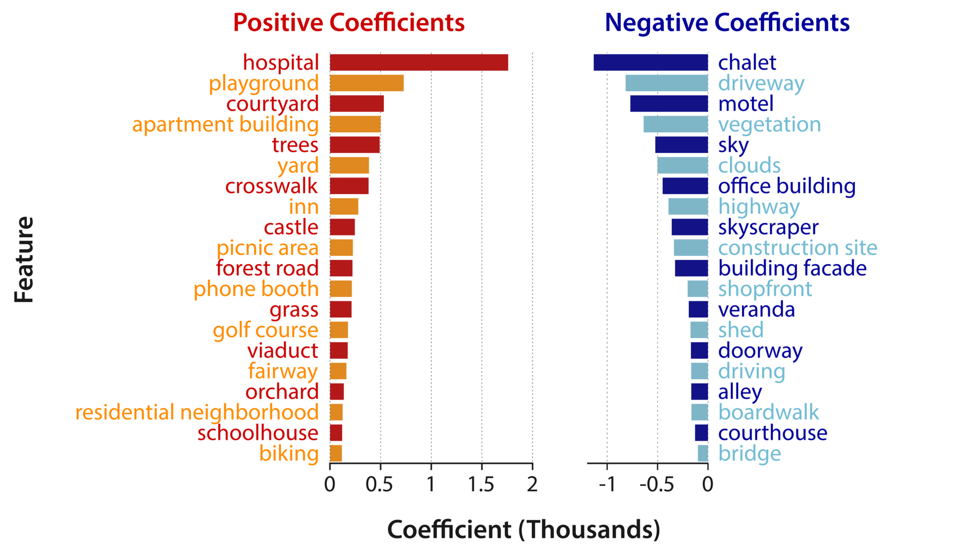

Figure 6: Top 20 model coefficients from relative change in house prices

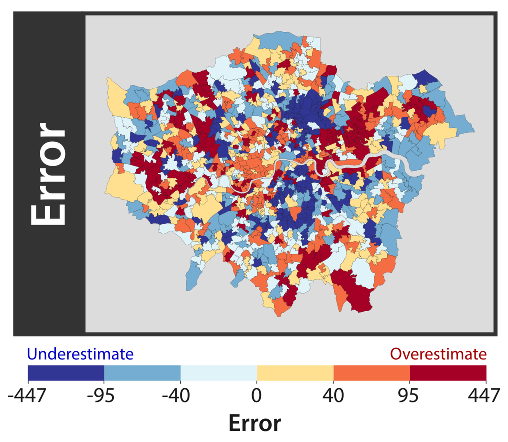

Figure 8: Prediction errors from relative change in house prices

Conclusions

We aimed to answer two questions at the beginning and the answers are clear – we can infer current house prices and identify which areas have recently seen relatively high increases in house prices from Google Street View images. Since new information can be captured for them, we could incorporate Google Street View images into existing models to complement them.

For further information, contact Bhavan Chahal (b.chahal@warwick.ac.uk).

References

[1] Zhou B, Lapedriza A, Xiao J, Torralba A and Oliva A (2014) Learning deep features for scene recognition using places database. In: Proceedings of the 27th International Conference on Neural Information Processing Systems. Montreal, Canada, pp. 487–495

[2] Patterson G, Xu C, Su H and Hays J (2014) The sun attribute database: Beyond categories for deeper scene understanding. International Journal of Computer Vision108(1-2): 59–81.

[3] https://www.gov.uk/government/statistical-data-sets/price-paid-data-downloads

EPSRC
Engineering and Physical Sciences Research Council

WARWICK
THE UNIVERSITY OF WARWICK