



UNIVERSITY OF  
CAMBRIDGE

Centre for Digital Built Britain

2017/18 Mini-Project

Machine Learning and AI in the Built  
Environment

Final Report

Thies Lindenthal, University of Cambridge [[htl24@cam.ac.uk](mailto:htl24@cam.ac.uk)]

Erik B. Johnson, University of Alabama

## Abstract

This project improved the foundations for applying tried-and-tested machine learning (ML) approaches to the built environment. If anything, the cost of developing ML systems has dramatically fallen in the last years – and is expected to further decrease in the future. Still, collecting data for training and evaluation of machine learning system remains complex and costly. Deploying a trained model for inference at a large scale, soliciting feedback on model predictions and managing the flow of expert feedback into new iterations of a model is a also challenge.

This mini project reduced the cost of creating and deploying ML systems by creating versatile and extendable API's, data management infrastructure and mobile apps. This progress already facilitated new academic research projects in urban economics and real estate. A future version of the API's might be commercialised in areas like mortgage origination, insurance claim processing or property tax (non-UK, though) estimation.

## Research Question

How can we cost-efficiently scale ML research in the built environment from a one-city scope to the national level, manage training and evaluation data, deploy production systems for inferences and collect feedback from human experts efficiently?

## Methodology

This project can be structured into three main segments, each addressing a core challenge for researchers utilising ML in the built environment. The deliverable of the project is a set of interrelated software components.

1. **Collect expert input efficiently:** This project developed an app to collect and manage training data from a diverse group of human experts. An intuitive user interface enables non-technical users to collect pictures of buildings via import, device camera, randomised Google Street View images), to classify pictures efficiently and to evaluate results after automatic classifications. The app is written using the open source Ionic “Progressive Web App” framework<sup>1</sup>, which makes it truly cross-platform (iOS, Android, web browsers) and easy to amend.

For now, the app supports one model: an automatic vintage estimator for UK residential real estate (Lindenthal & Johnson, 2018)<sup>2</sup>. The app can be adapted to other models within minutes. A trial version of the app can be accessed (preferably with a mobile device) at this web address: <https://www.cremll.com/app>

Classifications can be defined and managed via centralised configuration files (in JSON format) which turned out to be a fast and convenient approach.

2. **Integrate with other software/apps:** An API can be called from external applications to estimate building level attributes based on a dynamic set of models. The current version of the API is hosted on Amazon AWS and supports the direct upload of pictures, which are then classified using one or more trained ML models. Industry-strength open source APIs for

---

1 <https://ionicframework.com/pwa>

machine learning models already exist (e.g. Tensorflow Serving) which we then combine with a thin layer of customised code to allow for parallel inference from multiple models. The API returns classification scores and additional information in JSON format which can be parsed by e.g. apps or software used in research or the industry.

For API addresses and credentials, please contact the author.

3. **Store and display.** We continuously collect imagery of individual properties across England, Scotland and Wales from Google Street View, refining a methodology from Lindenthal & Johnson (2018)□. A set of “workers” hosted on Amazon AWS EC2 instances identifies individual buildings on Google Street View and sends the images to a central API, which classifies the pictures and stores the results centrally (spatial databases). Previously, image collection and classification has been done on local computers in our offices, which obviously does not scale well.

The current status of our data collection and classification can be followed in real time at this interactive map: <https://thies.carto.com/builder/359c34d0-9d3b-481a-b3b9-abb75bced8e6/embed>

## Discussionwebserver

The UK is in an exceptional position when analysing the built environment from an economic perspective: Property transactions for residential real estate are readily available public data (Land Registry, 2017)□. However, the level of detail for each transaction is very low. Only basic building attributes are recorded: Property type, freehold status, newly built vs. re-sale. Over the last year, I have utilised Big Data and advances in machine learning to augment these sales data using remotely sensed information, e.g. deriving building size and volume from LIDAR (Lindenthal, 2017b)□ or architectural homogeneity from 3D city models (Lindenthal, 2017a)□

A new working paper (Lindenthal & Johnson, 2018)□ established a method to extract pictures of individual buildings from Google Street View (previous research has been constrained to the street or block level). Using deep convolutional neural networks we built a machine learning model for detecting the buildings’ vintage (Georgian/Victorian/...) from these pictures. We were able to classify all of Cambridge’s buildings and to estimate price premia for certain styles.

To scale this work to the national level and also to increase the level of detail, a core layer of IT infrastructure is needed. This project allowed us to develop exactly this: An API to access centrally hosted ML models, an app to make the API conveniently accessible, feedback channels, and data storage and management systems.

Are we re-inventing the wheel by writing software for ML use cases? Why not use existing services like Google’s ML platform? We expect that in a few years from now, our current system will indeed be outdated and obsolete. Currently, commercial platform providers are not versatile enough to support our requirements. Ironically, one cannot run the latest versions of Google’s object detection API on Google ML.

## Conclusion

All in all, the cost of “prediction machines” (Agrawal, Gans, & Goldfarb, 2018) will continue to fall, making ML ubiquitous. We hope to contribute to this trend in the field of real estate and urban economics research through new software tools and digital infrastructure. Two new research projects are already utilising the new systems:

- Together with Mike Langen (Maastricht University), we propose a new method to collect structural property characteristics, using image recognition and machine learning. Based on a training dataset of 10,000 Google Street View images, our algorithm is able to detect and extract property characteristics, such as number of floors, building style, windows, garden, etc. By explicitly focusing on extracting property-level characteristics without land registry information as a requirement, our algorithm is easier to implement and more precise than previous algorithms when it comes to property applications, making it an attractive choice for urban studies. Given the ubiquity of Google Street View, our approach is transferable into many markets, allowing researchers to enrich existing datasets with building-level information or collect new data in a cheap and efficient manner.
- A graduate student of ours is evaluating the deterioration of high streets in economically struggling cities by automatically identifying shops, classifying these shops into categories, estimating vacancies and linking these data to asset prices and measures of urban well-being.

## Related and Further Work

The app and API are currently targeted to classification tasks (images/sound/maps/...). The natural next step would be to also support object detection use cases. Mike Langen, my co-author from Maastricht University is currently exploring how to detect individual buildings from street level imagery. Once we have established a feasible workflow, we will integrate object detection into the API and the app. Further work also includes scaling of the image collection process to cover more of the UK's buildings at a greater speed.

## Acknowledgements

*This project was supported by a mini-projects award from the Centre for Digital Built Britain, under InnovateUK grant number 90066*

## References

- Agrawal, A., Gans, J., & Goldfarb, A. (2018). *Prediction Machines: The Simple Economics of Artificial Intelligence*. Harvard Business Review Press. Retrieved from <https://books.google.co.uk/books?id=Y9LFtAEACAAJ>
- Land Registry. (2017). Price Paid Data. Retrieved from <http://landregistry.data.gov.uk/app/ppd>
- Lindenthal, T. (2017a). Beauty in the Eye of the Home-Owner: Aesthetic Zoning and Residential Property Values. *Real Estate Economics*, 1–26. <http://doi.org/10.1111/1540-6229.12204>
- Lindenthal, T. (2017b). *Estimating Supply Elasticities for Residential Real Estate in the United Kingdom*. Reference Module in Earth Systems and Environmental Sciences (Comprehens). Elsevier Inc. <http://doi.org/https://doi.org/10.1016/B978-0-12-409548-9.09682-2>
- Lindenthal, T., & Johnson, E. B. (2018). *Machine Learning, Building Vintage and Property Values*. Department of Land Economy Working Paper. University of Cambridge.