RAISE THE ROOF: TOWARDS GENERATING LOD2 MODELS WITHOUT AERIAL SURVEYS USING MACHINE LEARNING

F. Biljecki [1*], Y. Dehbi [2]

[1] Urban Analytics Lab, National University of Singapore, Singapore – filip@nus.edu.sg
[2] Institute for Geodesy and Geoinformation, University of Bonn, Germany – dehbi@igg.uni-bonn.de

Commission VI, WG VI/10

KEY WORDS: LoD2, 3D city models, machine learning, roof, 3D GIS

ABSTRACT:

LoD2 models include roof shapes and thus provide added value over their LoD1 counterparts for some applications such as estimating the solar potential of rooftops. However, because of laborious acquisition workflows they are more difficult to obtain than LoD1 models and are thus less prevalent in practice. This paper explores whether the type of the roof of a building can be inferred from semantic LoD1 data, potentially leading to their free upgrade to LoD2, in a broader context of a workflow for their generation without aerial campaigns. Inferring rooftop information has also other uses: quality evaluation and verification of existing data, supporting roof reconstruction, and enriching LoD0/LoD1 data with the attribute of the roof type. We test a random forest classifier that analyses several attributes of buildings predicting the type of the roof. Experiments carried out on the 3D city model of Hamburg using 12 attributes achieve an accuracy of 85% in identifying the roof type from sparse data using a multiclass classification. The performance of binary classification hits the roof: 92% accuracy in predicting whether a roof is flat or not. It turns out that the two most useful variables are footprint area and building height (i.e. LoD1 models without any semantics, or LoD0 with such information), and using only them also yields relatively accurate results.

1. INTRODUCTION

Roof is a key part of a building, and its top surface is frequently modelled in 3D GIS, benefiting applications that require the knowledge of the roof type and geometry. However, such information is not always available, as in practice many CityGML/CityJSON (Gröger et al., 2012; Ledoux et al., 2019) datasets are modelled in LoD1 (without roof shapes) and without the roof type attribute, and acquiring them is mandated from LoD2. But even then the attribute about the roof type is often absent.

The research question we seek to answer in this paper is whether it is possible to predict the type of the roof of a building without roof measurements, i.e. from a certain set of attributes using machine learning. Inferring the roof type in that context may be important for four use cases.

First, generating LoD2 models without having to map them with traditional approaches. Bypassing aerial surveys is the main motivation for this work, and we focus on this aspect.

Second, for data quality and verification of existing LoD2 models (e.g. gauging the quality of the roof reconstruction), i.e. to detect whether a roof in existing data is modelled wrongly. It may happen that during an automatic reconstruction from lidar a roof type is misinterpreted resulting in a reconstruction of a rooftop with a wrong attribute and inaccurate geometry (Akca et al., 2010; Gooding et al., 2015). There is research on benchmarking the quality of the roof reconstruction in 3D city models, but it is mostly focused on geometric accuracy rather than thematic (Oude Elberink and Vosselman, 2011; Rottensteiner et al., 2014), so work such as this one might contribute in filling that void. Besides the imperfection of reconstruction workflows, the

results may be affected by noisy observations. An advantage of our approach is that we use a different set of data, which might be affected by a smaller magnitude of error, thus deriving better predictions of the roof type and therefore serving as a reference.

Third, roof reconstruction might be aided with information about the roof type. This work also goes hand in hand with the traditional approach of acquiring LoD2 models with data obtained from aerial surveys. For example, during the automatic reconstruction from lidar point clouds having the information about the type of the roof may be a useful lead for the reconstruction algorithms and might greatly contribute to the quality of the generation of LoD2 models. Again, our approach uses another set of data, so coupling our predictions with roof reconstruction may contribute to the overall increase of accuracy. This idea has been investigated by Henn et al. (2013).

Fourth, as 3D city models may have the information about the roof type in the form of attribute without their geometry, this work may also be useful in enriching LoD0 (e.g. cadastral datasets) and LoD1 models with this attribute.

The scope of our work is predicting the type of the roof, which is the first step for generating their geometry required for LoD2 models. Therefore, much of the paper is shaped around the context of generating LoD2 models from LoD0/LoD1 datasets without elevation measurements (the first use case).

2. BACKGROUND AND RELATED WORK

2.1 Use of building roof information

While LoD2 models containing roofs are in general more desirable than LoD1 models, a recent study suggests that the value of

---

*Corresponding author

details in 3D city models is driven by the intended use case (Biljecki et al., 2018). Another study (Wong and Ellul, 2018), involving a questionnaire, is more apposite suggesting that information on the rooftop (whether descriptive in the attribute form or as LoD2 geometry) may be useful for some use cases and stakeholders, but irrelevant for others. Therefore we have to consider the value of LoD2 over LoD1 from a use case context. At the same time, it should not be neglected that the increase of LoD also increases complexity (Ellul and Altenbuchner, 2013).

While in general LoD2 models may or may not improve the quality of spatial analyses, some analyses require LoD2 at minimum. For example, Weiler et al. (2019) argue that having an LoD2 is essential for rooftop photovoltaic estimations (and in a broader context for energy estimations, e.g. for more accurate calculations of the thermal boundaries of a building (Pasquinelli et al., 2019)). Similar conclusions are reached by Peronato et al. (2016).

Lindenthal (2017) analyses the urban form and proves that the shape homogeneity among neighbouring homes has implications on the valuation of real estate. The geometry of the roof is an important factor in the study.

Schröter et al. (2018) use the roof type to get more insight about buildings in characterising the vulnerability of residential buildings to flooding.

The final example of a use case is general visualisation. From the visual point of view it is more attractive to have LoD2 models rather than LoD1.

2.2 Mapping building roofs

LoD2 are usually acquired using data obtained with airborne laser scanning or photogrammetry (or their combination), either manually or automatically using model-driven, data-driven or hybrid methods. There has been a great deal of research on that topic (Henn et al., 2013; Jung and Sohn, 2019). However, in 2019 lidar coverage (especially open data) still seems to be a luxury (Mathews et al., 2019), inhibiting the availability of LoD2 models.

Another interesting aspect to note is that the literature review suggests that much research dealing with LoD2 generation and roof reconstruction focuses on generating the geometry, while in many cases ignoring the thematic aspect (attribute) of the type of the roof without ever storing it as an explicit attribute.

Other techniques for procuring 3D models with roof shapes involve generating them from architectural/BIM datasets (Stouffs et al., 2018), and procedural modelling (Müller Arisona et al., 2013), but those account for just a fraction of datasets in practice, where the latter not always representing an accurate representation of the reality.

2.3 Related work

Henn et al. (2012) classifies building types from LoD1 models and other data using Support Vector Machines (SVM) for the semantic enrichment of low resolution building data. The obtained accuracy greatly depends on the building type class, ranging from 55% to above 90%. The same classifier (SVM) has been used by Römer and Plümer (2010) to identify architectural styles in 3D city models. Related work is done also by Neidhart and Sester (2004), Meinel et al. (2009), and Hecht et al. (2015).

In a subsequent work, Henn et al. (2013) find the most probable roof type with SVM, using this information to aid the roof
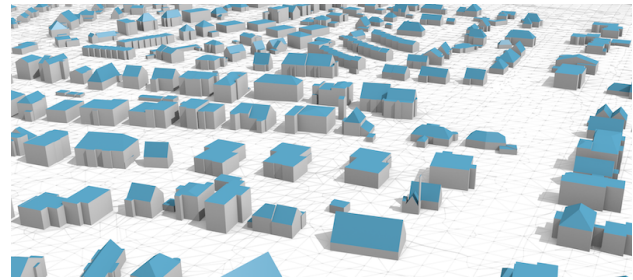

Figure 1. Semantic CityGML LoD2 model of Hamburg.

reconstruction. The achieved accuracy ranges from 70 to 97%, depending on the roof class. This work is quite relevant to ours, and we use some of the predictors and observations in our work. However, our work differs from theirs for not making use of some predictors that can be derived only from aerial surveys such as inclination of roof planes (our goal is to investigate whether we can obtain the same outcome without such information).

Awrangjeb et al. (2018) proposed a data-driven method for the 3D building roof reconstruction which has been subsequently used for building change detection. Jung et al. (2017) introduced an automatic 3D building reconstruction approach which in particular enforces geometric regularities on reconstructed roof models using MDL. Wu et al. (2017) applied a graph-based method based on hierarchical structure analysis of building contours acquired from bipartite graph matching for the reconstruction of building models from airborne LiDAR data.

Biljecki et al. (2017) reveal that it is possible to predict the height of buildings from 2D data with a reasonable accuracy, leading to an automatic generation of LoD1 models without elevation data. The attribute of the year of construction of a building can also be derived from other standard building attributes (Biljecki and Sindram, 2017), which we use in this work.

Allani-Bouhoula and Perrin (2008) reconstruct LoD2 models from LoD0 footprints (inferring the height of buildings and roof types) analysing local regulations. The achieved accuracy is between 88% and 91%. However, the short nature of the paper leaves many aspects unclear, e.g. it does not provide sufficient information for replication and the method seems to be very localised.

On the indoor side of 3D GIS, Loch-Dehbi et al. (2017) and Dehbi et al. (2018) predicted indoor models and developed a classification of room types and room shapes from sparse indoor data.

This literature review indicates that 2D and 3D data may reveal additional properties about the built environment, and we use related work as a basis to predict the roof type from similar data.
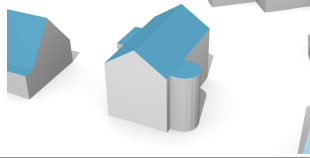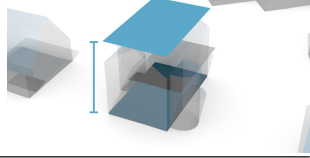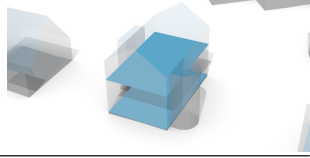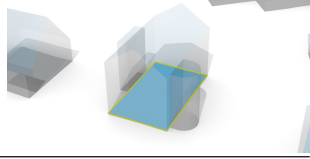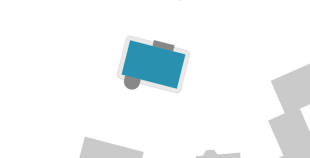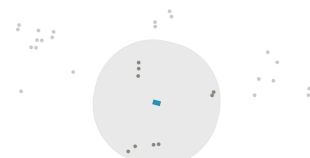
3. METHODOLOGY

3.1 Dataset and predictors

We have used the open datasets of Hamburg[1] for training and testing the classification. For the ground truth we used the CityGML LoD2 model (Figure 1), which contains the attribute describing the type of the roof for each building.

We extract a dozen of attributes that serve as predictors, potentially hinting at the type of the roof. They are listed in Table 1 describing them in details with examples.

Table 1. Variables in the classification with their description and examples. Source of the open datasets: Hamburg[1] and OpenStreetMap.

	Predictors	The exemplary building comprises three building parts with different roofs. We focus on the main part.	
1	Function	In our dataset we encountered 168 building functions (for more details see AdV (2019)), but a single one (residential building) accounts for 58% of instances in the dataset, followed by garages (10%). The function is assigned to the entire building, so we make an assumption that it is equal for all building parts.	
2	Building height	We have used the building part height available as attribute in the dataset. It is usually also available as an attribute in cadastral records, and if not – it can be extracted from the geometry of LoD1 models. The specification of the data (AdV, 2019) states that the building height in the dataset is from the base of the building to the highest point of the roof, and its accuracy is ± 1 m.	
3	No. of storeys	Similarly as the attribute above, the information on the number of storeys above ground is available in the dataset and it is one of the standard pieces of information included in building datasets such as those from cadastre. While the number of storeys is also a measure of the vertical extent of the building as the height, it will later become obvious that it does not duplicate information.	
4-9	Footprint geometry	We calculate various geometric characteristics of the footprint: (4) area, (5) orientation, (6) complexity, (7), perpendicularity, (8) normalised perimeter index (Angel et al., 2010), and (9) side ratio. The advantage of these measures are that they are always available and they can be calculated also from LoD0 models. The geometry of the footprint follows the outline of the roof, so it may be important.	
10	No. of adjacent objects	We calculate the number of direct neighbours (i.e. number of features in the buffer of 1 m). In this case the main building part of the building has two adjacent objects.	
11	No. of neighbours	We set up a buffer of 30 m around each building part to calculate the number of neighbours. The composition of the surrounding context of a building may be a useful predictor of the form of the roof.	
12	OSM objects	The number of amenities within 100 m may be relevant as well. The source of the data is OpenStreetMap. The consideration of amenities is inspired by the work of Henn et al. (2012) for the classification of building type. Amenities are used in predictions in GIS, e.g. Bakillah et al. (2014) used OSM amenities to predict the population count.	
	Roof type (label)	The dataset has 6 types of roofs (classes), see Figure 2 for all examples. While the modelling guidelines (AdV, 2019) describe 15 roof types in Germany, in our dataset we have found less classes than that. For additional experiments (Section 4.2) we also added a new label with binary classes: whether a roof is flat or not.	

(a) Flat roof.　(b) Gabled roof.　(c) Hipped roof.　(d) Mixed form roof.　(e) Pitched/shed roof.　(f) Pyramid roof.
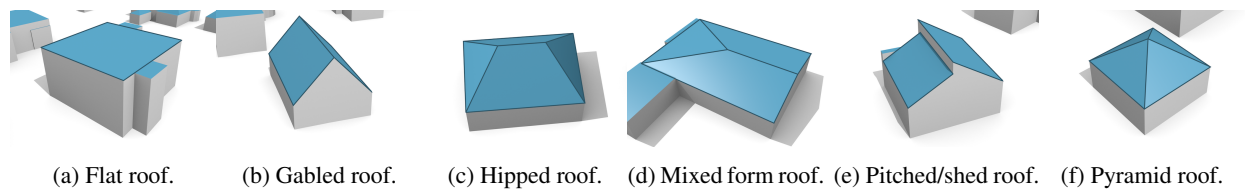
Figure 2. The six types of roofs (i.e. classes) in our study dataset.

Except some leads from related work (Section 2.3) we do not have any indication that these are by any means useful in the classification, but most of them do not cost anything to obtain and are always available (also in LoD0 datasets). It is relevant to note that while CityGML provides code lists for attributes (Gröger et al., 2012), in this dataset the roof type and building function are informed separately of the standard with local code lists.

About half of the predictors are related to the geometry of the footprint (essentially LoD0). Analysing the geometrical properties of footprints has been a key indicator in other 3D GIS work using machine learning (Dehbi et al., 2016; Hu et al., 2018). We deem that the footprint is key information, primarily because the outline of the rooftop follows the shape of the footprint, and might be instrumental in the classification. For example, from common sense we can *a priori* speculate that it is unlikely that a building with a circular-like footprint will have a gabled or pitched roof. This observation has also been highlighted by Henn et al. (2013). Furthermore, we have used OpenStreetMap (OSM) for one predictor, extracting the surrounding context.

The majority of buildings in the study area comprise multiple building parts with each part having a single rooftop. This is an important aspect because a building may have multiple rooftops and therefore it may comprise multiple roof types (see the example of the building in Table 1). Therefore in this work we focus on building parts and treat them separately. While most attributes are assigned to building parts, another relevant characteristic is that building function is attributed to the building level rather than building part. Therefore we assume that all building parts within the same building have the same function.

We use random forest classification (Breiman, 2001), and implement it in Python with scikit-learn (Pedregosa et al., 2011). For the data management we use 3DCityDB (Yao et al., 2018). In total we have about 707k buildings and building parts, and we use 70% of them for training the classification, and the rest for testing it and assessing the performance of the work. We tested other models such as SVM obtaining similar results, but we decided not to clutter the paper with additional information that improves the results only marginally.

3.2 Discovery

To develop the classification approach and to give a better understanding about the used dataset we have carried out a discovery, resulting in useful information and a few insightful plots. Figure 3 suggests that low-rise buildings dominate the study area, likely being highly influenced by the large share of garages. Flat roofs dominate the landscape (Figure 4), again likely garages play a significant role. The share of roof types is unequal (e.g. pyramid roofs account for just 0.4% of the data) resulting in an imbalanced training dataset. It also appears that the building part height plays an important role (e.g. 78% of building parts with flat roofs are low-rises, in contrast with 59% for mixed-form roofs).
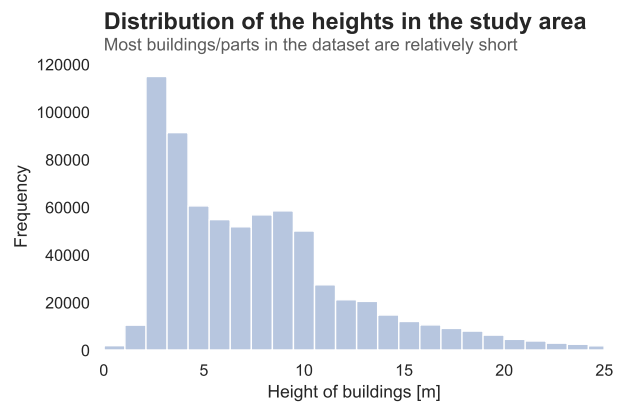

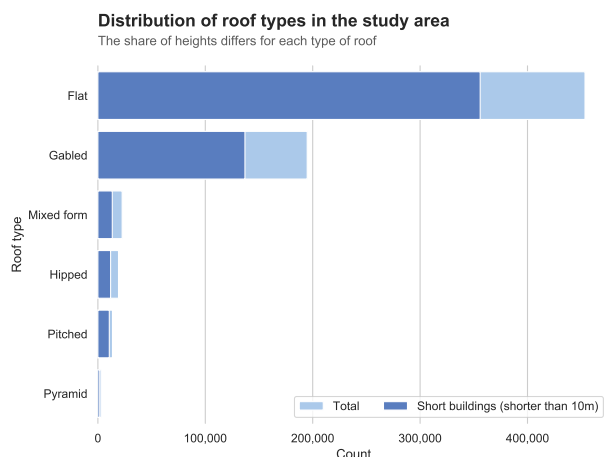Figure 3. Distribution of heights of buildings/parts.

Figure 4. Share of roofs by type in the study area.

Figure 5 further indicates the relevance of including the height, but this time also coupled with the information on the number of storeys.

4. RESULTS

4.1 Multiclass classification

The classification gives an accuracy of 85% in predicting the roof type using the 12 variables listed in Table 1. The accuracy is somewhat lower than the one obtained by Henn et al. (2013), but it is difficult to compare the methods and results side by side because their method uses information obtained using lidar data, while exactly the opposite is the goal of our paper (to develop a method that does not require elevation data). We believe that the slight difference in their favour is owing to the inclusion of some variables that can be obtained only from point clouds.
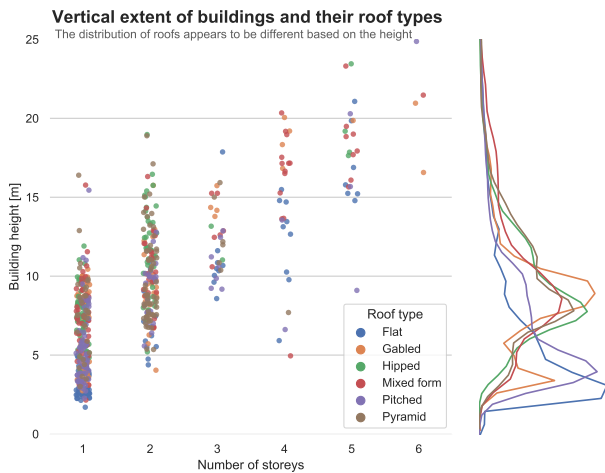
Figure 5. The vertical extent of the building part gives an indication of the roof type. The plot shows a 0.1% random subset.

The performance is given for each roof type in Figure 6. While the results are as in related work depending on the roof type, it is interesting to note that the performance between roof types differs between our results and those derived by Henn et al. (2013). Their most reliable classification is of gabled roofs (97%), while ours is for flat roofs (also 97%). The detection of hipped and pyramid roofs has poor performance in both methods, likely due to the imbalanced dataset. Many roofs that are not flat and gabled are misclassified as gabled roofs.
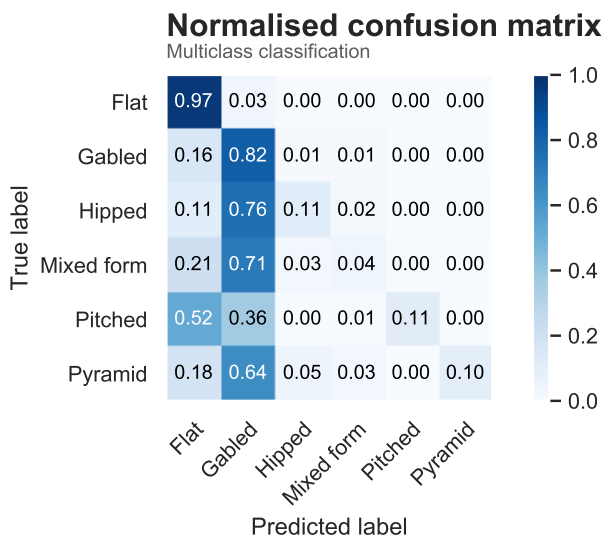

Figure 6. Performance of the multiclass classification.

The next important insight in the classification is the relevance of each predictor. By far the two most important variables are footprint area and building height, with footprint complexity being the least useful indicator of the roof type.

4.2 Binary classification

The classification presented so far was a multiclass classification, considering 6 types of roofs. In this section, we consider the binary classification: determining whether the roof is flat or not. There are multiple reasons for working in this direction. First,
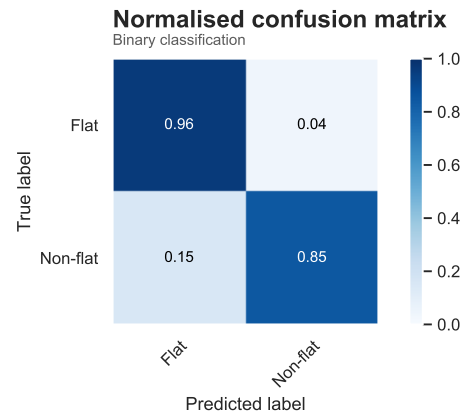

Figure 7. Performance of the binary classification.
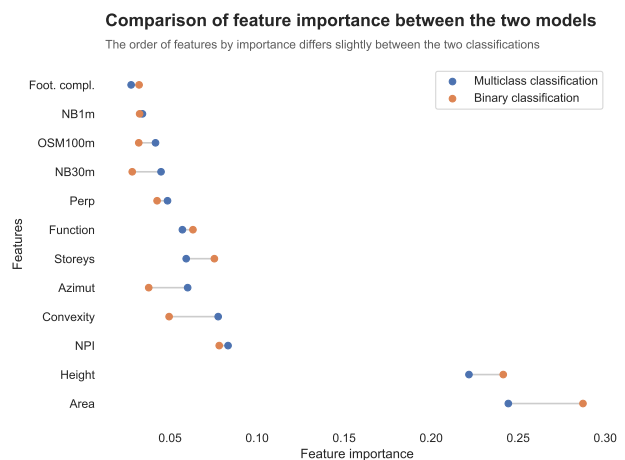

Figure 8. Feature importance of the two approaches. Some predictors are more important than others.

reducing the number of classes may result in a higher accuracy. Second, the detailed taxonomy of roofs may not always be important, and despite the reduced amount of information the binary decision whether the roof is pitched or flat may still be useful, and having such predictions with a higher accuracy may be more relevant than the classification presented in Section 4.1.

The accuracy we achieve for this classification is higher: 92%. The confusion matrix is given in Figure 7, again asserting that the accuracy varies across classes.

We expect that there are use cases in which this binary classification will be useful (also meaning that no roof geometry is actually needed for the use case, just the attribute on the roof type). For example, estimating the suitability for greening a roof (flat roofs are more suitable for vegetating them than any other type of pitched roof), establishing drone landing sites, and planning evacuation of buildings from air.

It seems that the feature importance between the two classifications differs resulting in their different ranking (Figure 8).

In both classifications, the most important predictor is the one (footprint area) not related to 3D, which is available in LoD0, giving confidence that in the future we might achieve LoD2 models from LoD0 data, and not only from LoD1 (elaborated more
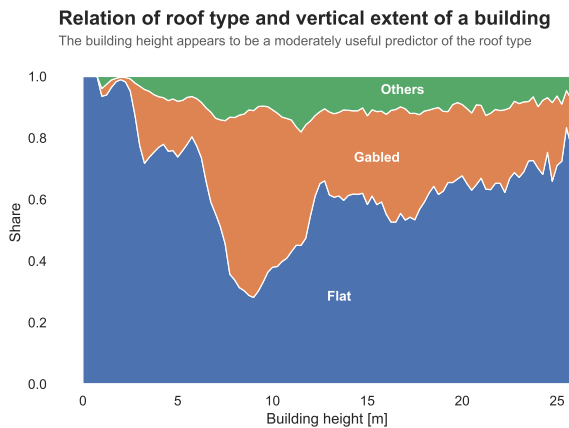
Figure 9. Share of roof types by height.

in Section 5.2). The importance of the building height however should not be discounted, as it comes second in importance. As it is clear from previous examples this is also common sense: the taller the building is, it is less likely that it has a non-flat roof. However, the relationship is not trivial (see Fig. 9). Some of the other predictors are also relevant, but to a much smaller extent.

4.3 Dealing with very sparse data

Figure 8 gives a valuable insight in the relative importance of the predictors used (Table 1). We have investigated the performance of the classification using only the two most important predictors (area and height), developing two additional classifications (multiclass for the 6 roof classes, and binary). The multiclass classification performance is 75%, while the binary comes at 81%.

5. DISCUSSION AND FUTURE WORK

5.1 Accuracy

The achieved accuracy levels vary depending on the number of involved classes and the specific roof class, with the best scenario being identifying whether a roof is flat or not using 12 predictors (we achieve an accuracy of 92%). The usefulness of these results ultimately depends on the intended use and applications (see the discussions in Sections 1, 2.1, and 4.2). For some of these, a degree of error could be more than acceptable, while not for others.

On the other hand, we do not have access to the performance measures of automatic roof reconstruction methods using lidar data in terms of detecting the categorical information about the roof type, but we can take the liberty to speculate that they are not perfect either and our results actually might not deviate that much from them. For future work it will be beneficial to compare the performance using additional predictors, and to compare the method to the identification of the roof type from methods using point clouds for the reconstruction (or even combine the two, as suggested in the introduction of the paper as the third use case).

5.2 Road map to LoD2 models without aerial surveys

In this paper we have obtained the roof type of buildings, but not the geometry, i.e. we identify the roof type as a categorical attribute. Generating the geometry of the roof and fitting on a building resulting in an LoD2 model will be part of future work, composing a broader project in automatically deriving LoD2 models from sparser datasets (LoD0/LoD1) using machine learning

(Figure 10). In that context, with this paper we have made a substantial advancement — predicting the roof type. However, while the roof type is the most important characteristic of a roof, it is not the only one. In order to generate the geometry, obtaining real LoD2 models, we would need to infer additional information, e.g. height of the eaves and the ridge, orientation of the geometry of the roof, and other measures (depending on the type of the roof, e.g. flat roofs would have none of these, but hipped roofs would require also additional measures such as ridge length, see Figure 2c). Achieving that goal will require predicting additional parametres, and the required information and complexity of that task will depend on the roof type.

5.3 Limitations

5.3.1 Lack of reliable data and uniformity This paper addresses the coverage of LoD2 models, and attempts to improve it by establishing a new method for their generation with machine learning. At the same time, we need LoD2 models with the ground truth of roof types so we can train our classification system. A limitation of the work is that there might not be many suitable datasets to do so, and the performance of the classification should improve with the increase in the diversity of the data (i.e. combining data from multiple countries). While the number of LoD2 datasets is growing, they are not always rich in semantics. For example, they may have the roof geometry mapped in great detail, but unfortunately at the same time omitting the attribute on the roof type, required for our method.

In the particular case of Hamburg we do not have many roof types. The dataset is dominated by two roof classes (Figure 4), inhibiting the diversity of training data and usefulness in other geographic areas. Therefore, as we have been dealing with a dataset that has a disproportionate number of classes, the classifier is biased towards specific roof types.

Finally, there is no guarantee that the training data is free of errors. For example, Figure 3 suggests that there are errors in the used dataset, since some building heights are unrealistically short. The same probably goes for other predictors we used, such as the number of storeys, for which we have no guarantee that they are entirely accurate. For future work, a comprehensive inspection of the data and manual verification of the attributes is advised.

5.3.2 Multiplicity of roofs LoD1 models may come in multiple forms based on the complexity and partitioning of the footprint (e.g. see the discussion in (Biljecki et al., 2016)). In our experiments we have used building parts that are obtained from LoD2 models, and those include refined building parts with each having a separate rooftop. However, in LoD0 and LoD1 datasets of the study area, buildings are modelled as single objects even if they are composed of distinct parts that have variable roof heights and roof types (see Figure 11 for the LoD1 representation of our exemplary building from Table 1, which in this case is covered by only one footprint while in reality it is composed of three parts with three rooftops). Therefore in such cases a reliable classification might not be possible. However, that is a problem also in the traditional approach of the reconstruction of LoD2 models from lidar data (Alexander et al., 2009).

Future work will require automatic segmentation of footprints to detect building parts with distinct rooftops. Most of the work partitioning a building involves 3D data such as point clouds (Dehbi and Plümer, 2011; Commandeur, 2012; Oude Elberink et al., 2013), thus doing that with footprints alone will be complex.
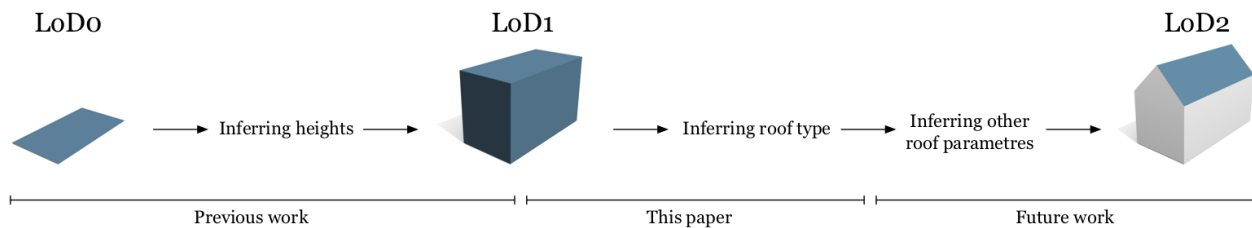
Figure 10. Proposed pipeline from LoD0 footprints to LoD2 models without aerial survey measurements. For previous work (inferring heights of buildings from footprints, generating LoD1 models) see (Biljecki et al., 2017).
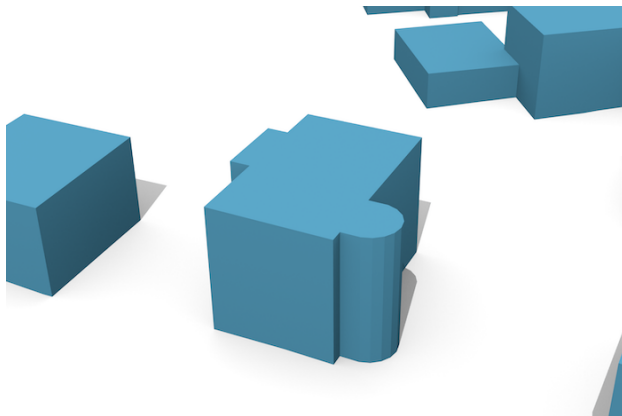

Figure 11. LoD1 counterpart of the LoD2 building exemplified in Table 1. The multiplicity of roofs within the same footprint might hamper future work.

6. CONCLUSIONS

Our work suggests that generating LoD2 models without roof measurements might be a viable possibility in future because we show that it is possible to predict the roof type from sparse data i.e. lower LoD datasets (LoD0 and LoD1). The implementation is not complex and it was done entirely using open-source software and open data, facilitating replication. The work is also useful for other purposes such as supporting roof reconstruction and verifying existing data. The impediments of the work are imbalanced data and multiplicity of roof types in a building.

Several ideas for future work are presented in Section 5, and are mostly focused on achieving the LoD0 → LoD2 pipeline (Figure 10). Additional directions for future work include conflating other sources of data increasing the number of predictors. Furthermore, we did not apply the inferred patterns to test the classification in another distant geographic area such as another continent. While we do not have a proof at the moment, we believe that the classifier should work with a comparable performance at least in nearby (German) cities, owing to the similar architecture.

ACKNOWLEDGEMENTS

References

AdV, 2019. Prüfplan für Gebäudemodelle LoD1/LoD2 Ergebnis der Projektgruppe '3D-Geobasisdaten' der Arbeitsgemeinschaft der Vermessungsverwaltungen der Länder der Bundesrepublik Deutschland (AdV). Technical report.

Akca, D., Freeman, M., Sargent, I. and Gruen, A., 2010. Quality assessment of 3D building data. *The Photogrammetric Record* 25(132), pp. 339–355.

Alexander, C., Smith-Voysey, S., Jarvis, C. and Tansey, K., 2009. Integrating building footprints and LiDAR elevation data to classify roof structures and visualise buildings. *Computers, Environment and Urban Systems* 33(4), pp. 285–292.

Allani-Bouhoula, N. and Perrin, J.-P., 2008. Architectural rules for three-dimensional reconstruction. In: *CGIM '08 Proceedings of the Tenth IASTED International Conference on Computer Graphics and Imaging*, Innsbruck, Austria, pp. 237–242.

Angel, S., Parent, J. and Civco, D. L., 2010. Ten compactness properties of circles: measuring shape in geography. *The Canadian Geographer* 54(4), pp. 441–461.

Awrangjeb, M., Gilani, S. and Siddiqui, F., 2018. An effective data-driven method for 3-d building roof reconstruction and robust change detection. *Remote Sensing* 10(10), pp. 1512.

Bakillah, M., Liang, S., Mobasheri, A., Jokar Arsanjani, J. and Zipf, A., 2014. Fine-resolution population mapping using OpenStreetMap points-of-interest. *International Journal of Geographical Information Science* 28(9), pp. 1940–1963.

Biljecki, F. and Sindram, M., 2017. Estimating building age with 3D GIS. *ISPRS Ann. Photogramm. Remote Sens. Spatial Inf. Sci.* IV-4-W5, pp. 17–24.

Biljecki, F., Heuvelink, G. B. M., Ledoux, H. and Stoter, J., 2018. The effect of acquisition error and level of detail on the accuracy of spatial analyses. *Cartography and Geographic Information Science* 45(2), pp. 156–176.

Biljecki, F., Ledoux, H. and Stoter, J., 2016. An improved LOD specification for 3D building models. *Computers, Environment and Urban Systems* 59, pp. 25–37.

Biljecki, F., Ledoux, H. and Stoter, J., 2017. Generating 3D city models without elevation data. *Computers, Environment and Urban Systems* 64, pp. 1–18.

Breiman, L., 2001. Random Forests. *Machine Learning* 45(1), pp. 5–32.

Commandeur, T., 2012. Footprint decomposition combined with point cloud segmentation for producing valid 3D models. Master's thesis, Delft University of Technology.

Dehbi, Y. and Plümer, L., 2011. Learning grammar rules of building parts from precise models and noisy observations. *ISPRS Journal of Photogrammetry and Remote Sensing* 66(2), pp. 166–176.

Dehbi, Y., Gojayeva, N., Pickert, A., Haunert, J. H. and Plümer, L., 2018. Room shapes and functional uses predicted from sparse data. *ISPRS Ann. Photogramm. Remote Sens. Spatial Inf. Sci.* IV-4, pp. 33–40.

Dehbi, Y., Gröger, G. and Plümer, L., 2016. Identification and Modelling of Translational and Axial Symmetries and their Hierarchical Structures in Building Footprints by Formal Grammars. *Transactions in GIS* 20(5), pp. 645–663.

Ellul, C. and Altenbuchner, J., 2013. LOD 1 VS. LOD 2 – Preliminary investigations into differences in mobile rendering performance. *ISPRS Ann. Photogramm. Remote Sens. Spatial Inf. Sci.* II-2/W1, pp. 129–138.

Gooding, J., Crook, R. and Tomlin, A. S., 2015. Modelling of roof geometries from low-resolution LiDAR data for city-scale solar energy applications using a neighbouring buildings method. *Applied Energy* 148, pp. 93–104.

Gröger, G., Kolbe, T. H., Nagel, C. and Häfele, K.-H., 2012. OpenGIS City Geography Markup Language (CityGML) Encoding Standard. Version 2.0.0, Open Geospatial Consortium, OGC Doc. No. 12-019.

Hecht, R., Meinel, G. and Buchroithner, M., 2015. Automatic identification of building types based on topographic databases – a comparison of different data sources. *International Journal of Cartography* 1(1), pp. 18–31.

Henn, A., Gröger, G., Stroh, V. and Plümer, L., 2013. Model driven reconstruction of roofs from sparse LIDAR point clouds. *ISPRS Journal of Photogrammetry and Remote Sensing* 76, pp. 17–29.

Henn, A., Römer, C., Gröger, G. and Plümer, L., 2012. Automatic classification of building types in 3D city models - Using SVMs for semantic enrichment of low resolution building data. *Geoinformatica* 16(2), pp. 281–306.

Hu, X., Fan, H. and Noskov, A., 2018. Roof model recommendation for complex buildings based on combination rules and symmetry features in footprints. *International Journal of Digital Earth* 11(10), pp. 1039–1063.

Jung, J. and Sohn, G., 2019. A line-based progressive refinement of 3D rooftop models using airborne LiDAR data with single view imagery. *ISPRS Journal of Photogrammetry and Remote Sensing* 149, pp. 157–175.

Jung, J., Jwa, Y. and Sohn, G., 2017. Implicit regularization for reconstructing 3d building rooftop models using airborne lidar data. *Sensors* 17(3), pp. 621.

Ledoux, H., Arroyo Ohori, K., Kumar, K., Dukai, B., Labetski, A. and Vitalis, S., 2019. CityJSON: A compact and easy-to-use encoding of the CityGML data model. *Open Geospatial Data, Software and Standards* 4(4), pp. 1–12.

Lindenthal, T., 2017. Beauty in the Eye of the Home-Owner: Aesthetic Zoning and Residential Property Values. *Real Estate Economics*.

Loch-Dehbi, S., Dehbi, Y. and Plümer, L., 2017. Estimation of 3D Indoor Models with Constraint Propagation and Stochastic Reasoning in the Absence of Indoor Measurements. *ISPRS International Journal of Geo-Information* 6(3), pp. 90.

Mathews, A. J., Frazier, A. E., Nghiem, S. V., Neumann, G. and Zhao, Y., 2019. Satellite scatterometer estimation of urban built-up volume: Validation with airborne lidar data. *International Journal of Applied Earth Observation and Geoinformation* 77, pp. 100–107.

Meinel, G., Hecht, R. and Herold, H., 2009. Analyzing building stock using topographic maps and GIS. *Building Research & Information* 37(5-6), pp. 468–482.

Müller Arisona, S., Zhong, C., Huang, X. and Qin, H., 2013. Increasing detail of 3D models through combined photogrammetric and procedural modelling. *Geo-spatial Information Science* 16(1), pp. 45–53.

Neidhart, H. and Sester, M., 2004. Identifying building types and building clusters using 3-D laser scanning and GIS-data. *Int. Arch. Photogramm. Remote Sens. Spatial Inf. Sci.* XXXV/B4, pp. 715–720.

Oude Elberink, S. and Vosselman, G., 2011. Quality analysis on 3D building models reconstructed from airborne laser scanning data. *ISPRS Journal of Photogrammetry and Remote Sensing* 66(2), pp. 157–165.

Oude Elberink, S., Stoter, J., Ledoux, H. and Commandeur, T., 2013. Generation and Dissemination of a National Virtual 3D City and Landscape Model for the Netherlands. *Photogrammetric Engineering and Remote Sensing* 79(2), pp. 147–158.

Pasquinelli, A., Agugiaro, G., Tagliabue, L. C., Scaioni, M. and Guzzetti, F., 2019. Exploiting the Potential of Integrated Public Building Data: Energy Performance Assessment of the Building Stock in a Case Study in Northern Italy. *ISPRS International Journal of Geo-Information* 8(1), pp. 27.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M. and Duchesnay, É., 2011. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research* 12(Oct), pp. 2825–2830.

Peronato, G., Bonjour, S., Stoeckli, J., Rey, E. and Andersen, M., 2016. Sensitivity of calculated solar irradiation to the level of detail: insights from the simulation of four sample buildings in urban areas. In: *Proceedings of the 32nd International Conference on Passive and Low Energy Architecture*, Los Angeles, CA, United States, pp. 685–690.

Römer, C. and Plümer, L., 2010. Identifying architectural style in 3d city models with support vector machines. *Photogrammetrie-Fernerkundung-Geoinformation* 2010(5), pp. 371–384.

Rottensteiner, F., Sohn, G., Gerke, M., Wegner, J. D., Breitkopf, U. and Jung, J., 2014. Results of the ISPRS benchmark on urban object detection and 3D building reconstruction. *ISPRS Journal of Photogrammetry and Remote Sensing* 93, pp. 256–271.

Schröter, K., Lüdtke, S., Redweik, R., Meier, J., Bochow, M., Ross, L., Nagel, C. and Kreibich, H., 2018. Flood loss estimation using 3D city models and remote sensing data. *Environmental Modelling & Software* 105, pp. 118–131.

Stouffs, R., Tauscher, H. and Biljecki, F., 2018. Achieving Complete and Near-Lossless Conversion from IFC to CityGML. *ISPRS International Journal of Geo-Information* 7(9), pp. 355.

Weiler, V., Stave, J. and Eicker, U., 2019. Renewable Energy Generation Scenarios Using 3D Urban Modeling Tools—Methodology for Heat Pump and Co-Generation Systems with Case Study Application. *Energies* 12(3), pp. 403.

Wong, K. and Ellul, C., 2018. User requirements gathering for a national 3D mapping product in the United Kingdom. *ISPRS Ann. Photogramm. Remote Sens. Spatial Inf. Sci.* IV-4/W6, pp. 89–96.

Wu, B., Yu, B., Wu, Q., Yao, S., Zhao, F., Mao, W. and Wu, J., 2017. A graph-based approach for 3d building model reconstruction from airborne lidar point clouds. *Remote Sensing* 9(1), pp. 92.

Yao, Z., Nagel, C., Kunde, F., Hudra, G., Willkomm, P., Donaubauer, A., Adolphi, T. and Kolbe, T. H., 2018. 3DCityDB - a 3D geodatabase solution for the management, analysis, and visualization of semantic 3D city models based on CityGML. *Open Geospatial Data, Software and Standards* 3(1), pp. 208.