

3D CITY MODELS FOR URBAN MINING: POINT CLOUD BASED SEMANTIC ENRICHMENT FOR SPECTRAL VARIATION IDENTIFICATION IN HYPERSPECTRAL IMAGERY

P.A. Ruben¹*, R. Sileryte², G. Agugiaro³

¹ Msc. Geomatics for the Built Environment, Delft University of Technology, The Netherlands - pablo.ruben.private@gmail.com

² Department of Urbanism, Delft University of Technology, The Netherlands - r.sileryte@tudelft.nl

³ 3D Geoinformation, Delft University of Technology, The Netherlands - g.agugiaro@tudelft.nl

Commission IV, WG 10

KEY WORDS: Semantic 3D city models, Enrichment, CityGML, Point cloud, Urban mining

ABSTRACT:

Urban mining aims at reusing building materials enclosed in our cities. Therefore, it requires accurate information on the availability of these materials for each separate building. While recent publications have demonstrated that such information can be obtained using machine learning and data fusion techniques applied to hyperspectral imagery, challenges still persist. One of these is the so-called 'salt-and-pepper noise', i.e. the oversensitivity to the presence of several materials within one pixel (e.g. chimneys, roof windows). For the specific case of identifying roof materials, this research demonstrates the potential of 3D city models to identify and filter out such unreliable pixels beforehand. As, from a geometrical point of view, most available 3D city models are too generalized for this purpose (e.g. in CityGML Level of Detail 2), semantic enrichment using a point cloud is proposed to compensate missing details. So-called deviations are mapped onto a 3D building model by comparing it with a point cloud. Seeded region growing approach based on distance and orientation features is used for the comparison. Further, the results of a validation carried out for parts of Rotterdam and resulting in KHAT values as high as 0.7 are discussed.

1. INTRODUCTION

Recent circular economy initiatives, such as urban mining (i.e. the reuse of building materials) create a strong need for information about the materials currently composing the built environment. In 2008 for instance, no less than 47% of the United Kingdom's CO₂ emissions resulted from the building sector. Among these, 83% are related to the usage and 15% to the construction (BIS, 2010). Therefore, not only the construction of new buildings but also the retrofitting of existing ones create a need for more sustainable approaches. One of these is the sourcing of materials already enclosed in the cities and thus released relatively close-by at the end of their life cycle (Baccini, Brunner, 2012). A critical question is how to find information on the presence and quantity of such materials. While new buildings increasingly start using material passports or Building Information Models (BIM) to keep a record of such information, huge stocks of materials are already available in older constructions without being (digitally) documented.

One building element of special attention is the roof as it requires regular maintenance and refurbishment, approximately every 10-25 years. Developments in the field of material science make reuse and recycling of used roof material possible at a local scale. Bituminous roofs can, for example, be mixed into new road pavements (Townsend et al., 2007). To allow efficient processes, thematic data at object (e.g. building) level is however required for two reasons. On the one hand, sourcing becomes much easier as a mere works notification is sufficient to know when a material is removed, avoiding the effort of on-site inspection. On the other hand, the (positive or negative) impact of circular approaches often depends on the exact geo-

graphic distribution (Patouillard et al., 2018). In fact, materials still have to be moved and the exact impact assessment requires a sufficient level of detail. Ideally, demolition and renovation activities should be coordinated to intentionally limit such movements.

While statistical disaggregation methods exist, they are not accurate enough at the object level. To accurately identify roofing materials, this paper builds upon research done in the field of (mainly hyperspectral) aerial imagery techniques. The available technology which allows for identification of materials at a pixel level, however, suffers from 'salt-and-pepper' noise (also referred to as 'salt-and-pepper' effect) caused by a mix of materials within a single pixel, e.g. caused by the presence of skylights, solar batteries or chimneys (Zhong et al., 2014, Priem, Canters, 2016). To overcome this effect, this paper proposes a priori identification of such pixels using a 3D city model which was previously enriched by means of a point cloud. This enrichment is required because in currently available 3D city models details such as chimneys are often left out.

The key idea described in this paper is to map geometrical 'deviations' by comparing a 3D city model with a point cloud. By identifying points, 2D geometries representing the 'deviations' are obtained and fused with the hyperspectral raster image geometry (see figure 1).

After introducing the method itself, the results of a case study based on the 3D city model of Rotterdam (The Netherlands) are presented. The city model was provided by the municipality (Gemeente Rotterdam, 2016) and the point cloud used is the third version of the Dutch national AHN point cloud dataset (PDOK, 2016).

*Corresponding author

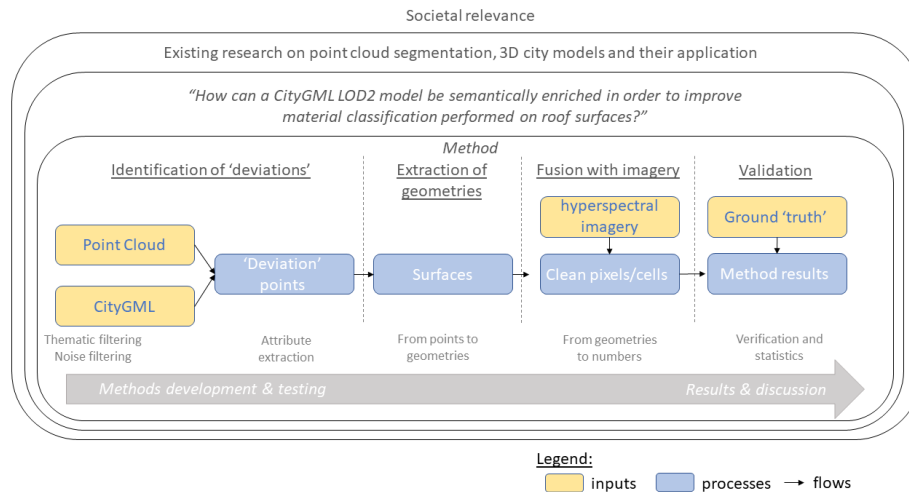


Figure 1. Framework of the research method in this paper ('clean' pixels are pixels containing a single material)

1.1 Usage of hyperspectral imagery for surface material identification

A methodology that has already proven effective to gather roof cover material data (and, more generally, surface cover materials) is classification applied to hyperspectral imagery. Several studies have demonstrated that reliable classification of different surface cover materials is possible (Heiden et al., 2007, Demarchi et al., 2014). While most studies were conducted at pixel level, successful application of this classification at building (polygon) level was demonstrated as well. By applying support vector classification on APEX (airborne prism experiment) imagery of Brussels, overall kappa values as high as 0.87 were obtained for the identification of selected surface cover materials (Priem, Canters, 2016).

As mentioned before, a challenge for the correct material classification are spectral deviations, also known as 'salt-and-pepper noise': classifiers are oversensitive to spectral variations that occur within a pixel, even when these are in the minority. Previous work at polygon level addresses this by applying a majority filter (Priem, Canters, 2016). However, such a majority filter can only be applied *after* the classification and requires pixels with spectral variation (and thus not the material covering the majority of the pixel) as classification result to be in the minority - which is not always the case as seen in the figure 2).

1.2 3D city models as an alternative to majority filter correction

As the majority filter might fail when a majority of pixels contain spectral variations (e.g. hangars with large roof windows – see figure 2), a *pre*-classification filter to identify these might be more suited. As this type of noise is not truly randomly distributed, 3D city models that are nowadays available for many cities can be used as input for such a filter, provided that they contain sufficient details. Within the CityGML standard, the required level of detail (LoD) is equivalent to LoD3. At the time of writing, however, LoD2 is the most widespread as the production of LoD3 city objects still involves a lot of manual (and thus expensive) work.




Figure 2. While the majority filter might work for the building on the top (among the pixels contained in the red polygon, only two are containing spectral variations - in this case chimneys), it is rather unlikely to work in situations with a high share of pixels containing spectral variations such as on the bottom (green: 'clean' pixels containing a single material, light green: 'clean' pixels located in shadow and thus only usable separately). Pixel size: 4*4 m; Background imagery: ©google earth/digitalglobe 2018

1.3 Related work

No previous work concerning the enrichment of 3D city models using point clouds was identified when redacting this paper. For inspiration, it was therefore decided to broaden the literature review and include some related topics.

A related approach is the integration of a Multi-View Stereo Mesh Model (MVSM) with an existing 3D city model which has been developed to improve solar potential analysis (Willenborg et al., 2018). For each roof surface, using both Euclidean




Figure 3. Overview of the method used for semantic enrichment of a 3D city model by using a point cloud

distance and angle orientation calculations, a set of candidate mesh faces is established, segmented and expanded. This work builds upon an earlier master thesis which investigated more use cases (Tryfona, 2017). These include semantic segmentation, attributes transfer and mesh straightening/simplification.

Several authors established taxonomies for feature recognition in point clouds (Nguyen, Le, 2013, Vosselman et al., 2004, Wang, Shan, 2009). By providing a structured overview, these publications allow the comparison of different approaches and selection of the current one (attribute- and region-based). Edge-based approach was considered too, but several limitations were identified (see part 2.3). Moreover, discussions of challenges such as the presence of crease, jump and smooth edges (Nurunabi et al., 2012) also contributed to exploratory research. Eventually, the study of the mentioned references led to the approach shown in figures 1 and 3.

2. SEMANTIC ENRICHMENT OF A 3D CITY MODEL USING A POINT CLOUD

The deviation identification is based on a seeded region growing (performed on the relevant portion of the point cloud representing the building footprint) and occurs in several steps (figures 1 and 3). First, a selection based on vertical distance between the 3D city model and the point is performed and an orientation-based region growing is applied. As the latter calculation is more computation-intensive than a simple distance calculation, performing it for a limited set of points allows to considerably reduce computation time. Then, identified ‘deviations’ are converted into a 2D geometry and merged with the hyperspectral imagery data. By doing so, pixels containing spectral variations can be filtered out *before* classification.

2.1 Seed selection by vertical distance measuring

The first step consists in creating point cloud subsets. First, a bounding box is created for each building and a subset is formed with all points containing it (only the horizontal, not the vertical dimension is considered). Second, a region of interest is created for each roof surface (spanning from -1 to +20 m vertically from the roof surface, see figure 4). This one defines a sub-subset which relates to the respective roof surfaces. For each point within this sub-subset, the vertical distance to the roof is calculated. In all cases whenever a user-predefined *threshold distance* is exceeded, the points are labelled as distance deviations.



Figure 4. Illustration of the region of interest for a given roof surface

2.2 PCA-based region growing

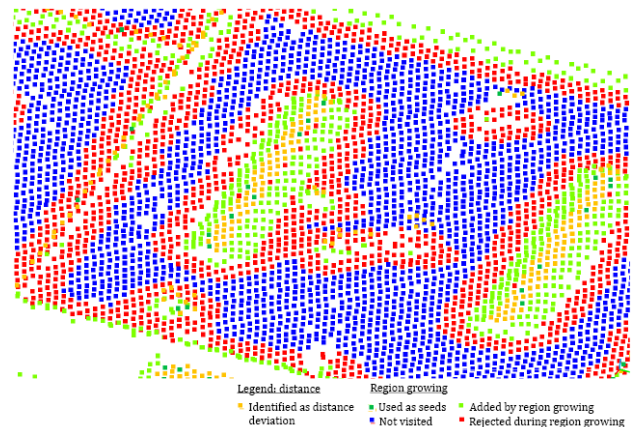


Figure 5. Illustration of the seed selection and region growing algorithm on a building with half-cylindrical roof windows that were absent in the 3D model. Settings used: *distance threshold* 20 cm and *angular threshold* 5.73°.

The identified seeds are then used as an input for the region growing. For each k-nearest (first level) neighbour of the seeds, the k-nearest (second level) neighbours are selected and an orientation test is performed. This one consists in a principal components analysis (PCA) to calculate the respective normal vector. Using the dot product, the deviation regarding the roof orientation is calculated. If the deviation angle is bigger than a user-predefined *angular threshold*, the point is added to the region, otherwise, the point is rejected. As illustrated in figure 5, the growing operation stops once the region is surrounded by rejected points and can thus not be expanded anymore.

2.3 Point to plane transformation using Voronoi diagrams

In the previous step, deviation regions are defined as a group of points, not as surfaces. To account for variations in point density, a conversion to a 2D shape is required. After considering concave hull, convex hull, boundary extraction (see below for identified limitations) and Voronoi diagram approaches, the latter was selected for its ability to adapt to the local context. It also offers the advantages to work without user settings and to support regions of any size (starting from a single point), variations in point density and shapes containing holes. A drawback is that the Voronoi diagram of all points (including ones not labelled as ‘deviations’) needs to be generated which, in terms of computation, makes it heavier than other approaches. Additionally, the cases of cells being infinite or going outside of the original roof geometry needs to be taken into account (see figure 6). In the end, the cells containing a deviation point are then extracted and adjacent cells are merged into a single geometry (per deviation region).




Figure 6. By building the Voronoi diagram, the translation of points marked as ‘deviations’ (red) into a 2D shape becomes possible. As some Voronoi cells might go beyond the roof surface or even be infinite, cropping with the roof surface geometry (in blue) is necessary

It should be noted that the identification of pixels containing a mapped ‘deviation’ might also be done by directly checking whether the pixel contains points labelled as deviations or not. However, the approximation of the area occupied by deviations would be rougher. Moreover, the usage of the Voronoi diagrams allows points close to the border of a pixel to appear as deviations in the neighbouring pixel(s), too.

Another approach that was explored is the boundary extraction by the usage of a minimum spanning tree (Wang, Shan, 2009). Using a convex-hull test on local neighbourhoods, border points are first identified and then used to build a graph based on k-nearest neighbourhood. Ultimately, the minimum spanning tree of this graph is generated and used as a border representation. As can be seen in figure 7, this approach does, however, have difficulties in identifying some critical border points. Furthermore, it appears to be challenging to use this approach for shapes with a width of a limited number of points only.

2.4 Identification of mixed pixels in aerial imagery

Once the ‘deviations’ are available as 2D geometries, the fusion with the pixels of hyperspectral aerial imagery is finally performed. To estimate the degree of deviations of the pixels contained in a roof surface, several steps involving the orthorectified pixel centre points must be followed. To avoid the appearance of interpolation artefacts, geometry correction is preferred

¹unless stated otherwise, all background imagery: ©PDOK landsdekkende luchtfoto’s 2016-2018




Figure 7. The identified border points (in green on the top image) of the deviation regions (in red) do not always include all critical corner points. Also, shapes with a width of only a few points (such as roof edges) result in border points positioned in a way that leads to the minimum spanning tree confusing opposed edges (bottom image). Please note: connecting the start and end of the minimum spanning tree was not performed within the feasibility study.¹




Figure 8. Using the orthorectified center points of a pixel and of its neighbours, a hexagon is defined. The points defining this one are subsequently orthorectified

to raster resampling. In practice, this means that each pixel is transformed into a hexagon, each vertex being defined with regard to the centre points of the neighbouring pixels (see figure 8). The vertices of the hexagon are then orthorectified and the intersection with the 2D ‘deviation’ geometry is computed. Using the total area of the hexagon, the area percentage of ‘deviations’ within a pixel is calculated. Based on this percentage and depending on the sensitivity of the classification algorithm subsequently used, the usability of a pixel to determine its surface material can be determined.

3. CASE STUDY: 3D CITY MODEL OF ROTTERDAM

An area located in the south of Rotterdam was selected as a case study due to its diversity in terms of built environments, i.e. in terms of neighbourhood functions, building shapes and roof materials. Within this study area, 41 buildings with roofs containing 831 pixels were selected as a representative set in terms of surface, function, height and visible quantity of deviations. While diversity strengthens the validity of results, a limitation must also be noted: big buildings (from about 6000 m² on) had

to be excluded due to the high number of pixels contained in the roof surfaces which would have represented more than 10% of the total number of pixels. This was necessary to ensure reliability as results might otherwise become too dependent on the characteristics of a specific big building.




Figure 9. Illustration showing the buildings south of the Maas river in Rotterdam selected for the validation set.

As input, the 3D city model produced by the municipality of Rotterdam (Gemeente Rotterdam, 2016) was used together with the AHN3, the third national Dutch point cloud dataset (PDOK, 2016), which are both produced based on aerial LiDAR acquisitions and available as open data. The AHN3 dataset is available as a set of tiles containing points already classified in five classes. For this work, the point cloud used was limited to the class ‘buildings’. It is worth mentioning that the point cloud used for the production of the 3D city model was obtained during a flight campaign separate from the AHN3, as a much higher density was required, i.e. approximately 30 points/m² vs. 8 points/m² for the AHN3 in the study area.

Additionally, a hyperspectral imagery dataset was used for the pixel geometries. It results from a flight carried out within the APEX on 17 September 2014. The GSD (ground sampling distance) is 3.5-3.8 m and data were obtained already orthorectified. Orthorectification as mentioned in part 2.4 was performed using a trigonometric equation and potential errors were estimated.

The ‘ground truth’ was obtained by visual evaluation of very-high-resolution (≤ 25 cm) aerial images (see figure 10). The yearly national acquisitions for the years 2016-2018 (PDOK, 2018) were used together with Google Earth imagery offering higher contrast and acquired in 2015. Where possible, Google Street View imagery from October 2016 was used, too.

In total, two performance indicators were selected for the validation of the proposed method. First, the KHAT statistic (Congalton, 1991) was chosen as it allows to have a good idea of the overall performance. Second, the commission error rate for clean pixels (i.e. without spectral variations) was selected to estimate the number of ‘failures’ which are likely to lead to material identifications affected by the ‘salt and pepper effect’.

3.1 Variable selection

For both the *distance threshold* and *angular threshold* used in the algorithm (see sections 2.1 and 2.2), settings were selected manually by visual analysis of results computed for a pilot



Figure 10. Example of a visualization obtained during the validation. Mapped ‘deviations’ are shown in yellow and can be overlaid with the (modified) pixels resulting from hyperspectral imagery pixels and contained in the building roofprint (geometry in blue). Settings used: *distance threshold* 20 cm and *angular threshold* 2°.

building. For each, a ‘strict’ and a ‘loose’ value were identified (i.e. 20 cm and 40 cm for the *distance threshold* - see figure 11; 2° and 5° for the *angular threshold* - see figure 12). The idea is to identify a representative performance spectrum despite the limited time which allowed only a limited number of combinations. The results presented in the next section are therefore highly specific to the values chosen.




Figure 11. Regarding the *distance threshold*, at 20 cm, confusion between ‘deviations’ and slightly sloped roofs that were reconstructed as flat ones occurred (top, red spot on the left). On the other hand, at 40 cm (bottom), such confusion did not appear anymore but critical seeds of the most important ‘deviation’ objects were still identified (e.g. all missing windows have a number of red points).

Other variables that had to be set manually are the k-nearest neighbour settings used for the region growing and the plane fitting. In both cases, a constant value of 10 was chosen. For the region growing, this chosen parameter value is based on a pilot test case shown in figure 5. It shows that the identified ‘deviation’ region is indeed surrounded by a layer with a thickness of 3 or more rejected points. For the plane-fitting, the choice is based on the pilot shown in figure 13.

3.2 Results

The sensibility analysis of the different setting combinations is presented in table 1 and figure 14. In total, four setting combinations were used, resulting from the two pairs obtained in the pilots of figure 11 (20 cm and 40 cm) and 12 (2° and 5°). It was performed using nominal data (presence of spectral variations vs. identification of ‘deviations’) and a pixel as a unit



Figure 12. Regarding the *angular threshold*, values of 1.13° (top) lead to a clear overestimation of 'deviations', while at 2.56 degrees (middle), the estimation seems to be rather correct. Therefore, 2° was selected as a 'strict' setting. At 5.73° (bottom), clear holes can be recognized inside geometries. Therefore, 5° were selected as a 'loose' setting.




Figure 13. Results obtained using PCA on different k-nearest neighbour settings. The points of which the normal vector diverged more than 5° from the roof are shown in red colour. While between settings 5° (left) and 10° (middle) differences exist (see zones inside black circles), they are minimal between 10° (middle) and 15° (right).

of measurement. Here the results for the KHAT statistic reach from 0.57 (for the settings 20 cm & 5°) to 0.70 (for 40 cm & 2°). Overall, there is a difference of about 0.11 between the KHAT statistics obtained for the 40 cm and the 20 cm settings. Between 2° and 5° , this difference is only about 0.02. For the commission error rate of the pixels without spectral variations,

Settings: distance = 20 cm; angle = 2°			
		truth (spectral variations)	
		no	yes
prediction (identified 'deviations')	no	186	23
	yes	128	494
KHAT statistic		0.59	
commission error rate 'no'		0.11	
Settings: distance = 20 cm; angle = 5°			
		truth (spectral variations)	
		no	yes
prediction (identified 'deviations')	no	212	63
	yes	102	454
KHAT statistic		0.57	
commission error rate 'no'		0.23	
Settings: distance = 40 cm; angle = 2°			
		truth (spectral variations)	
		no	yes
prediction (identified 'deviations')	no	247	48
	yes	67	469
KHAT statistic		0.70	
commission error rate 'no'		0.16	
Settings: distance = 40 cm; angle = 5°			
		truth (spectral variations)	
		no	yes
prediction (identified 'deviations')	no	282	100
	yes	32	417
KHAT statistic		0.68	
commission error rate 'no'		0.26	

Table 1. Error matrices (for the four setting combinations) showing the results of the sensibility analysis using nominal data and the cell as a unit of measurement




Figure 14. Comparison of selected performance indicators for the four setting combinations

the opposite situation can be observed: the difference between the 40 cm and 20 cm setting is only 0.03-0.05 while the one between the 2° and 5° settings reaches 0.10-0.12. Results for this performance indicator vary between 0.26 (for 40 cm & 2°) and 0.11 in the best case (for 20 cm & 2°).

4. FINDINGS & LIMITATIONS

Overall, the results suggest two optimisation approaches. In fact, the results show that, depending on the final goal, the priority might either be given to the KHAT statistic or the commission error rate of the pixels without spectral variations. The lowest possible KHAT values are achieved by the looser settings with 40 cm. This suggests that for accurate quantification of materials rather loose thresholds should be used. On the other hand, if the aim is merely to find the presence of materials, false positives should be avoided. Here, the lowest commission error rate of the pixels without spectral variations is achieved by stricter settings of 20 cm & 2°.

Another finding of this case study is that the GSD of the hyperspectral imagery used in the case study (3.5-3.8 m) can be deemed acceptable for the purpose of urban mining. Very few situations in which all pixels contained spectral variations occurred.

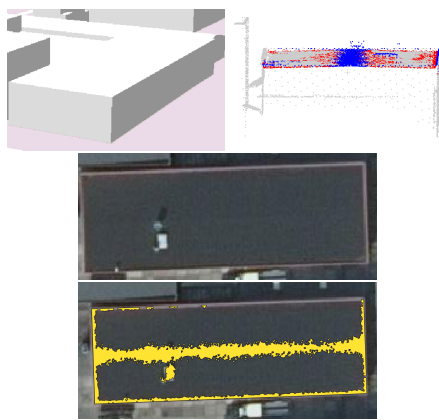


Figure 15. Example of a building that was wrongly reconstructed with a flat roof (top left), leading to the identification of ‘deviations’ (middle and bottom). Settings used: *distance threshold* 20 cm and *angular threshold* 5°. Top right: point cloud - blue: seed points selected based on distance and red: points added during PCA-based region growing.

For the 3D city model, quality requirements are rather high as wrongly reconstructed roofs lead to ‘deviations’ being detected where modelling quality is poor. While the vast majority of buildings fulfilled this requirement, slightly sloped roofs reconstructed as flat ones were encountered a few times in the case study (see figure 15).

This work aimed to prove the validity of the approach in a diversified built environment. If applying the method to a different context, it is thus recommended to treat the results with care as differences might exist. Furthermore, the results are specific to the settings chosen for the tolerances. As this was done with the aim of identifying a performance spectrum, more careful optimization might eventually increase performance.

Additional limitations due to geometric reasons exist. First, some chimneys and other small deviations might be missed due to their size in relation to the point cloud density. In the case of Rotterdam, assuming a point cloud of circa 8 evenly distributed points/m², and according to the well-known Nyquist-Shannon theorem, this corresponds to a GSD of circa 35 cm. In other words, the smallest detectable object has a planar length of approximately 70 cm (equivalent to an area of circa 0.5 m²). Further limitations arise in case of objects which were occluded

during LiDAR acquisition. In addition, spectral variations of which all points have a vertical distance below the threshold cannot be identified. This can be the case for elements such as flat-roof windows, solar panels or gutters.

5. DISCUSSIONS

An important factor of influence is the GSD of the hyperspectral aerial imagery. As it increases, so does the number of (unusable) pixels containing spectral variations and the accuracy of material quantification thus decreases. On the other hand, detection of material presence only requires a limited number of pixels without deviations – given that they can be reliably identified. With the approach developed in this thesis, the acceptable GSD for material presence detection might thus increase, lowering the number of acquisition flights and the data processing required for a given area. When doubling the flight altitude and keeping the same speed (but doubling the sensor’s sampling time), the GSD and the mapped area (per unit of flight time) are doubled, while the number of the pixels (for a given area) is divided by four.

The geometric limitation to objects with a planar length of more than 70 cm highlights the need for higher point cloud densities. However, modern LiDAR sensors can already provide much denser points clouds than the 8 points/m² used here. This is also the case with equivalent products from dense stereo matching. A factor limiting the usability of the latter is the vertical accuracy which is generally lower than for LiDAR. In practice, this would mean a higher vertical distance threshold, reducing the deviations that can be reliably identified.

The final observation relates to the CityGML standard as it does not currently support the storage of ‘deviations’. Workarounds were identified (e.g. using *BuildingInstallation* or *Generic CityObject* objects) but this is a rather “unorthodox” approach. A need for clearer guidelines concerning semantic enrichment could be considered in the future versions of the standard. From a broader perspective, a similar observation can be made for the storage of building materials which is not supported either. While some initial work has been done, as within the Energy Application Domain Extension (Agugiaro et al., 2018) which covers material aspects, it is not fully suited for applications dealing with circular material use.

6. CONCLUSIONS AND FUTURE RESEARCH

This research has demonstrated the potential of a semantically enriched 3D city model to replace majority filters in situations with limited pixels containing a single material. The case study and subsequent validation have shown the potential of semantically enriching LoD2 city models using a point cloud (method shown in figure 3). For the variable settings selected for the validation, KHAT values as high as 0.70 were obtained for distance-based seed selection followed by orientation-based region growing. Results indicate that, depending on the final goal, different settings should be used. For accurate quantification of materials, the lowest possible KHAT values - and thus rather loose thresholds - should be used. If, however, the aim is just to find the presence of materials, rather strict thresholds should be used to avoid false positives as much as possible.

Future research should focus on the third dimension of ‘deviations’ (e.g. median vertical distance). This is required for discrimination between shadowed and sunlit pixels, a step applied

before hyperspectral imagery classification (Priem, Canters, 2016) and mentioned in figure 2. Moreover, means to identify 'deviations' which are not geometric should also be researched. An option would be to use the intensity of return - a parameter typically included during LiDAR acquisitions.

Finally, beyond urban mining, other use cases exist for the enrichment of 3D city models with roof "deviations". The first example is automated quality checks, to identify parts of the 3D model that require improvements. As shown in 15, the quality of 3D modelling is not always satisfying and the method might be used to identify these cases. The second example is the improvement of solar potential calculation, similarly to related research that was introduced earlier (Willenborg et al., 2018). By adding missing elements such as chimneys or roof windows, shapes and areas of usable surfaces can be determined more precisely. As shadows are critical here too, additional motivation exists to include the third dimension of 'deviations'.

ACKNOWLEDGEMENTS

The authors would like to thank the city of Rotterdam (especially Mr. Christian Wisse and Mr. Christian Veldhuis), the APEX project (especially Mr. Andreas Hueni and Prof. Dominik Brunner) and the reviewers for their contributions, as well as Mr. Kaixuan Zhou for co-reading the master thesis behind this paper.

REFERENCES

- Agugiaro, G., Benner, J., Cipriano, P., Nouvel, R., 2018. The Energy Application Domain Extension for CityGML: enhancing interoperability for urban energy simulations. *Open Geospatial Data, Software and Standards*, 3(1), 2.
- Baccini, P., Brunner, P. H., 2012. *Metabolism of the anthroposphere: analysis, evaluation, design*. MIT Press.
- BIS, 2010. Estimating the amount of CO₂ emissions that the construction industry can influence.
- Congalton, R. G., 1991. A review of assessing the accuracy of classifications of remotely sensed data. *Remote sensing of environment*, 37(1), 35–46.
- Demarchi, L., Canters, F., Cariou, C., Licciardi, G., Chan, J. C.-W., 2014. Assessing the performance of two unsupervised dimensionality reduction techniques on hyperspectral APEX data for high resolution urban land-cover mapping. *ISPRS Journal of Photogrammetry and Remote Sensing*, 87, 166–179.
- Gemeente Rotterdam, 2016. 3d city model of rotterdam. www.3droterdam.nl.
- Heiden, U., Segl, K., Roessner, S., Kaufmann, H., 2007. Determination of robust spectral features for identification of urban surface materials in hyperspectral remote sensing data. *Remote Sensing of Environment*, 111(4), 537–552.
- Nguyen, A., Le, B., 2013. 3d point cloud segmentation: A survey. *RAM*, 225–230.
- Nurunnabi, A., Belton, D., West, G., 2012. Robust segmentation in laser scanning 3d point cloud data. *Digital Image Computing Techniques and Applications (DICTA), 2012 International Conference on*, IEEE, 1–8.

Patouillard, L., Bulle, C., Querleu, C., Maxime, D., Osset, P., Margni, M., 2018. Critical review and practical recommendations to integrate the spatial dimension into life cycle assessment. *Journal of cleaner production*, 177, 398–412.

PDOK, 2016. Actueel hoogtebestand nederland 3. <https://downloads.pdok.nl/ahn3-downloadpage/>.

PDOK, 2018. Landsdekkende luchtofo's 2016-2018. <https://www.pdok.nl/introductie/-/article/luchtfoto-pdok>.

Priem, F., Canters, F., 2016. Synergistic use of LiDAR and APEX hyperspectral data for high-resolution urban land cover mapping. *Remote sensing*, 8(10), 787.

Townsend, T., Powell, J., Xu, C., 2007. Environmental issues associated with asphalt shingle recycling. *Construction Materials Recycling Association, US EPA Innovations Workgroup*.

Tryfona, M.-S., 2017. Bidirectional enrichment of citygml and multi-view stereo mesh models. Master's thesis, Delft University of Technology, Faculty of Architecture.

Vosselman, G., Gorte, B. G., Sithole, G., Rabbani, T., 2004. Recognising structure in laser scanner point clouds. *International archives of photogrammetry, remote sensing and spatial information sciences*, 46(8), 33–38.

Wang, J., Shan, J., 2009. Segmentation of lidar point clouds for building extraction. *American Society for Photogrammetry. Remote Sens. Annual Conference, Baltimore, MD*, 9–13.

Willenborg, B., Pültz, M., Kolbe, T. H., 2018. Integration of semantic 3D city models and 3D mesh models for accuracy improvements of solar potential analyses. *International Archives of the Photogrammetry, Remote Sensing & Spatial Information Sciences*.

Yao, Z., Nagel, C., Kunde, F., Hudra, G., Willkomm, P., Donaubaer, A., Adolphi, T., 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), 1–26.

Zhong, Z., Fan, B., Duan, J., Wang, L., Ding, K., Xiang, S., Pan, C., 2014. Discriminant tensor spectral-spatial feature extraction for hyperspectral image classification. *IEEE Geoscience and Remote Sensing Letters*, 12(5), 1028–1032.

APPENDIX

The work presented in this paper mainly results from a Master thesis defended in July 2019. The detailed report can be found at <http://resolver.tudelft.nl/uuid:0fd2ffec-f25a-4391-be68-e816cc2611dd>. The implementation was performed in a python environment using several libraries (scipy, numpy, shapely, laspy, psicopg2). For the storage and querying of the city model itself, the 3D City Database (Yao et al., 2018) for PostgreSQL/PostGIS was used. Moreover, a repository with the implementation code is available at <https://github.com/Flyalbatros/3D-Models-in-Urban-Mining>.