EQUIRECTANGULAR 360°IMAGE DATASET FOR DETECTING REUSABLE CONSTRUCTION COMPONENTS

Ana Bendiek Laranjo[1,2], Jens J. Hunhevicz[1,2], Karsten Menzel[3], Catherine De Wolf[2]
[1]EMPA, Dübendorf, Switzerland
[2]ETH Zürich, Zurich, Switzerland
[3]Technische Universität Dresden, Dresden, Germany

Abstract

Insufficient as-built data hinders the transition of the architecture, engineering and construction (AEC) sector to a circular system. Combining reality capture and machine learning (ML) could help better detect reusable components. However, a comprehensive image dataset of on-site inventory for circular economy strategies has yet to be developed. This study introduces and describes the generation of a purpose-built, 360°dataset. Initial validation using the YOLOv8 object detection model demonstrates a 63.4% mean average precision (mAP50), making it viable for computer vision. Further exploration of automating building stock inventory using 360-degree images and ML for urban mining is needed.

Introduction

The architecture, engineering and construction (AEC) industry must change from a linear to a circular system to minimize its harmful impact and meet climate targets. Reusing resources is a key circularity strategy for reducing waste (European Council and European Parliament, 2008). In the built environment, this strategy implies the non-destructive recovery and reuse of building components according to their original purpose without a loss in their value (Hillebrandt et al., 2018). The extensive reuse of building components is currently challenged by the lack of sufficient data about material and substance composition and the as-built condition of the existing building stock (Çetin et al., 2021; Iacovidou et al., 2018; Uotila et al., 2021).

Urban mining projects need detailed (digital) information on their composition and dimensions to reuse and recycle building resources efficiently (Çetin et al., 2021; Honic et al., 2019; Uotila et al., 2021). This as-built data needs to be generated in advance, typically using reality-capture techniques such as light detection and ranging (Lidar) (Gordon et al., 2023; Xiong et al., 2022), 360°cameras (Gordon et al., 2023), and public digital imaging mining using Google Streetview (Raghu et al., 2023). Digital data for information extraction is typically processed further using photogrammetry for extracting dimensions (Gordon et al., 2023; Xiong et al., 2022) and machine learning (ML) for determining the material composition of the building stock (Raghu et al., 2023).

Among reality-capturing technologies, 360°cameras proved to be the most viable technology for capturing reliable information for deconstruction purposes in a comparative case study, because of their high accuracy and low noise (Gordon et al., 2023). Panoramic images have advantages over planar photos because of their comprehensive view of spaces, device cost-effectiveness, and quick data capture (Barazzetti et al., 2018; Gordon et al., 2023; Chou et al., 2020). Furthermore, 360°images are increasingly used for quick image analyses in computer vision applications (Barazzetti et al., 2018; Gordon et al., 2023; Chou et al., 2020). However, computer vision and 360°images have yet to be combined for circularity strategies.

This research examines whether combining building inventory using 360 °cameras with ML can be used to detect reusable building components. Yet since ML applications are only as good as the training data on which they are trained (Géron, 2023), the quality of the dataset is paramount. The relevance of training data can be defined as the extent to which it aligns with the data that the model is likely to encounter in the production phase (Witt, 2023). This concerns what is depicted and how it is annotated. The quality of the dataset refers to aspects influencing the generalization capabilities of a model to unseen data. The generated dataset should be suitable for the application in terms of representativeness, image quality, dataset size, and variance (Géron, 2023). For the envisaged ML application, a real-world dataset that simulates a survey of end-of-life (EoL) scenarios is necessary.

However, in a preliminary field study we found that existing panoramic datasets of buildings are not suitable for the above application for the following reasons: they are either synthetically generated, they are a compilation of different scenarios without any repetition, they are limited to specific parts of buildings (e.g., the façade), they are too homogeneous in quality to be captured on-site, or they have only very broad labels (e.g., doors are not subdivided by function into entrance doors, emergency exits, or room doors)

Therefore, our study takes a first step towards using 360°panoramic images in reusability assessment by generating and validating a 360°image dataset in a case study on the Technical University of Dresden (TU Dresden) campus. The custom dataset ("TUDataset") is preliminarily validated using the state-of-the-art YOLOv8 object detection model to identify building components. The validation explores whether the data collected during an on-site audit is qualitatively sufficient for object detection applications. It also explores the differentiation of component types within a component category. In so doing, this

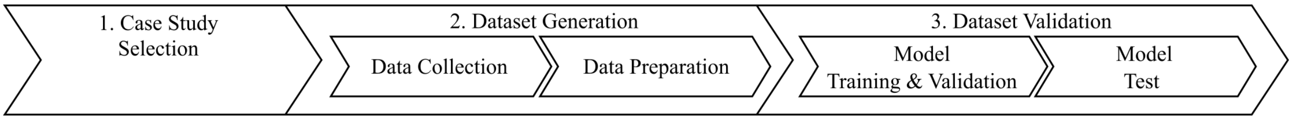

Figure 1: Research design of this study.

study gives valuable insights into the process of generating suitable 360°image datasets for building component reuse by discussing the chosen approach, potential for improvement, and next steps toward a full data-capturing ML pipeline.

Structure of the paper

The paper is structured according to the Figure 1. First, the following section describes the representative case study for EoL situations. It defines the data collection locations and the objects (i.e., the reusable building components) to be detected by the ML application. Second, the section “Dataset Generation” presents the approach and the image capture process to showcase how a 360°dataset can be created, and describes image processing for the specific use case of component reuse detection. Third, the “Dataset Validation” section describes how the approach is validated for computer vision. Experiments with the YOLOv8 object detection model verify the suitability of the dataset generation approach for the desired purpose of detecting reuse components. The methodology steps are presented in greater detail in the respective sections. Finally, the discussion reviews the contributions and further research steps toward a full pipeline of 360°image dataset creation for ML-based reuse detection.

Case Study Selection

As described in the introduction, the case study determines the relevance of the dataset for future input data from EoL buildings. The EoL refers to the final phase of a building’s existence (after the production and use phase) in its current form (Hillebrandt et al., 2018). For this study, the case of ”Technical University of Dresden (TU Dresden) campus” was chosen. It was designed to mirror urban mining projects by encompassing the entire building premises and emphasizing the frequent repetition of diverse building component designs within structures. It consists of five buildings on the TU Dresden campus that were selected according to the following considerations: around 32% of the residential and nonresidential buildings demolished in 2022 were built before 1948, 53% between 1949 and 1986, and only 14% between 1987 and 2010 (Statistisches Bundesamt, 2023). Accordingly, the case study was limited to the buildings on campus built between 1906 and 1990: two modern buildings, as well as three buildings of reform architecture from the early 20th century (see Table 1). Different EoL situations are presented: while the Beyer Building is partially gutted and currently unused, the Fritz-Förster building is only partially in use due to renovation. The Nürnberger Ei and the Schumann building are contain operating offices and equipment (and, therefore, more data clutter).

Finally, only publicly accessible buildings were considered to comply with privacy compliance for publishing. Accordingly, the building was inventoried outside of business hours. People were asked to leave the field of view, and images with recognizable faces were removed during data cleaning.

Dataset Generation

While the relevance of the data was ensured in the design of the case study, the dataset quality (see Introduction) is addressed in the dataset generation process.

Data collection

A total of 1112 relevant images were captured for the case study (see 1). The data was generated from scratch using the specialized OpenExperience 360°camera helmet (Figure 2). The helmet has two installed 180°cameras, whose individual images are seamlessly merged into a spherical or panoramic view by a stitching algorithm. Capturing images in an onsite audit using a conventional 360°camera ensures that the data represents the expected input data. The final ML application is expected to generalize to input data collected under similar technical conditions.

The equirectangular projection (ERP) on a two-dimensional pane was used for the dataset. The ERP of the images has a resolution of 7000 × 3500 pixels, a horizontal and vertical resolution of 96 dpi, and a bit depth of 24 (see Figure 3). To address image variance deficiency observed in existing datasets, the data was captured at different times of the day and under varying weather conditions, resulting in different lighting and shading conditions. Furthermore, the images were taken without a fixed object distance, often capturing the same room from different positions and a person’s viewpoint, generating a variance in the object’s

Figure 2: DIGIBAU 360°helmet camera used in this study. Source: (OpenExperience, 2023).

Table 1: Data generation protocol on TU Dresden campus

Generation Date	Building	Location	Construction/ Renovation Year	#Images
30.03.2023	Fritz-Foerster-Bau	Mommsenstraße 6, 1069 Dresden	1926/2022	280
30.03.2023	Nürnberger Ei	Nürnberger Straße 31a, 01187 Dresden	1996	37
31.03.2023	Georg-Schumann-Bau	Münchner Platz 3, 01187 Dresden	1906	369
31.03.2023	Haus 116	August-Bebel-Straße 30, 01219 Dresden	1970/2013	159
14.04.2023	Beyer-Bau	George-Bähr-Straße 1, 01069 Dresden	1913	277
				1122

Figure 3: Unprocessed ERP image included in the TUDataset.

appearance, scale, and occlusion. Finally, the strongly varied resolution of the images resulted in different representations of the same object.

Dataset Preparation

The generated raw data was prepared using the Roboflow[1] online tool through data cleaning, annotation, set partition, and pre-processing and augmentation further described below.

Data Cleaning

In object detection, images or videos are used as inputs, and the features are extracted from the information in the pixels. The selected YOLO (You Only Look Once) model employs a feature extraction method that does not require the prior definition and cleaning of features. Therefore, the data cleaning was limited to eliminating duplicates and excluding low-quality and blurred images.

Data annotation

Object detection is a supervised machine learning application that requires a fully labeled dataset for the training. What is recognized and the granularity of differentiation is determined in the annotation process.

The EU policy fails to provide regulations and prerequisites for the reuse of components without prior (destructive) testing. Thus, to determine what could be reused within the case study, a field search including online reused material marketplaces (Concular, Restado. Bauteilnetz, Ebay Kleinanzeigen, SALZA, and Bauteilclick), practice reports, and guidelines was conducted to identify building components categories with assumed general or ”default” reuse potential.

[1]https://roboflow.com/

Table 2: Selection criteria for component identification.

Criteria	Description
Relevance in practice	The building components should be found in large quantities in every building.
Post-processing	The component should be reusable per default, with little refurbishment and no testing needed.
Typology	The components should have standard features, but different configurations. Shapes and colors should differ
Indoor	The availability of data demands a restriction to indoor components.

Because this study is a proof-of-concept, the scope of the object detection model was set to recognize only a selection of reusable components. To fit the ML requirements on representativeness, quantity, and variance, the selection criteria in Table 2 were developed. Out of the pool of default reuse components, five categories adhere to the established selection criteria: doors, windows, radiators, sanitary objects, and lights. These component categories are considered reusable without further destructive testing. They present a large intra-class variation (e.g., many different types of windows) and were captured in large numbers in the case study.

The next step determines the granularity of differentiation within the selected component categories. Existing indoor datasets, such as the PanoContext Zhang et al. (2014), Indoor360 Chou et al. (2020) etc., differentiate only in rough component categories, such as doors and windows. However, function, material, and design are decisive for component reuse. The selected reusable component categories, ”windows, doors, sanitary objects, and electrical

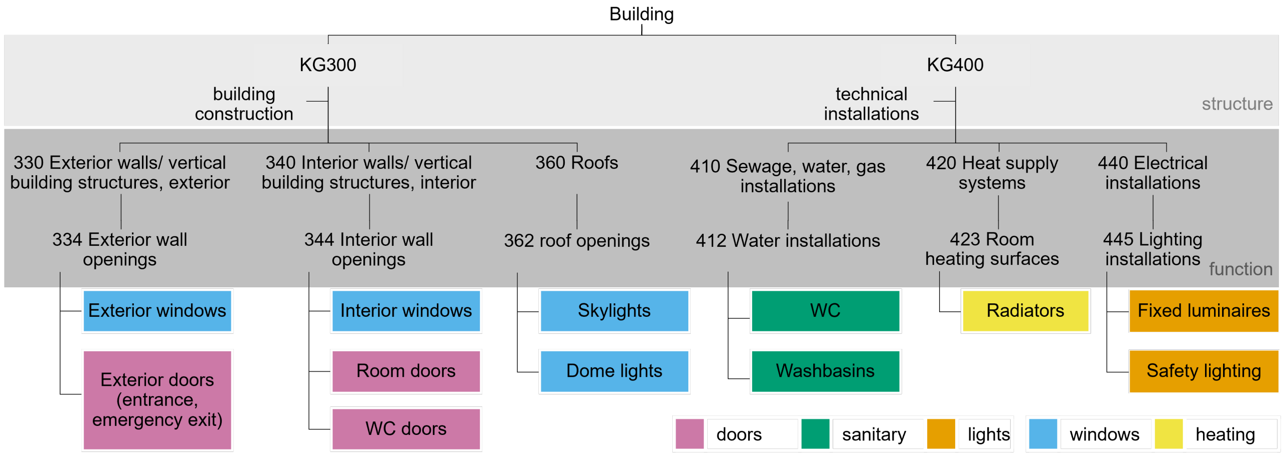

Figure 4: Categorization of reusable component selection according to DIN276.

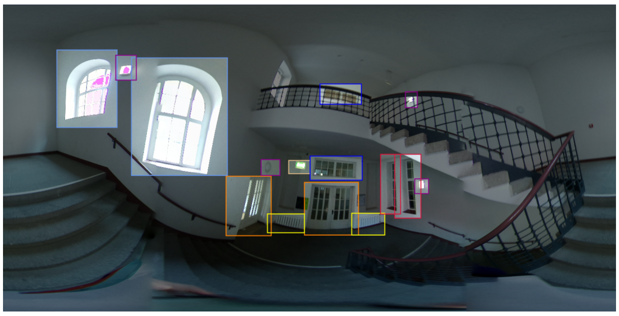

Figure 5: Each image label has a different bounding box indicating different component types.

installation", needed further differentiation to meet these requirements.

This study used the German AEC sector's building cost structure of the national standard DIN 276:2018-12 (Deutsches Institut für Normung, 2018) as a reference for the annotation. The component selection (windows, doors, plumbing, lights, and heating) was further differentiated into cost groups according to their associated building structure and function within the building or system, as seen in Figure 4. Categorization by cost group will not only standardize the work in different projects by reference to a public standard, but it could also allow for rapid cost determination for selective reconstruction.

In the annotation, the selected reusable component categories were set as super-categories, and the different cost groups were set as component classes. Furthermore, the component types were differentiated within these classes according to the material and design. For example, if the frame materials differed, two external glass doors with the same design would belong to different types. Hence, the labeling was structured according to component category, class (e.g., exterior door), and type (e.g., type 21). This resulted in 136 labels, ergo, different component types.

Dataset Partition

This study employed hold-out validation. It is an evaluation technique in which the dataset is split into three subsets: training, validation and test set. The training set is used to fit the models, the validation set is used to estimate the prediction error for model selection, and the test set is used to assess the generalization error of the final model configuration. The test data set is "held-out" and only used once, as reusing it can result in a substantial underestimation of the true test error. The TUDataset images were therefore divided into 70% training (792), 20% validation (226), and 10% testing (116) images.

Data pre-processing and augmentation

Applying pre-processing techniques to the training, validation, and testing sets ensures that the machine learning model learns and infers based on consistent image properties. Inference refers to the process of generating predictions.

First, the images were auto-orientated, removing the EXIF (Exchangeable Image File Format) data from images to ensure that they are displayed in the same manner as they are stored on the disk. Then, the image size was stretched to 640 × 640 pixels. Finally, the last step consists of data augmentation or training set expansion. This technique artificially increases the training set size and is a regularization method, reducing overfitting. To improve the model's tolerance to position, orientation, and size changes, the augmented instances should be as realistic as possible and ideally be indistinguishable from non-augmented instances by the human eye. Following Zhao et al. (2021), several augmentation techniques considering the particularity of the equirectangular projections were applied to the training and validation sets: the images were flipped horizontally and sheared $\pm$ 15°horizontally and $\pm$ 15°vertically both on the image level and bounding box level (see Figure 7). The outputs per training sample were set to 3, creating three altered images for every instance and resulting in a final dataset size of 2718 (3 × 792) images.

Dataset Validation

After generating the dataset, the suitability of the on-site 360°images for computer vision methods, and con-

(a) Scaling

(b) Horizontal flipping

(c) Image level shearing (+15°horizontally, +15°vertically)

(d) Image level shearing (15°horizontally, -15°vertically)

(e) Image level shearing (-15°horizontally, +15°vertically)

(f) Image level shearing (-15°horizontally, -15°vertically)

Figure 7: Different Augmentation Techniques

sequently for ML-aided reusability assessment, was validated using an object detection model.

The use of 360°images poses specific challenges for object detection. The projection of the spherical image to a plane distorts the objects depending on their distance and angle to the camera viewpoint. Integrating optimized distortion-aware convolution layers could handle these geometric deformations (Li et al., 2023). However, for the purpose of this proof-of-concept study, the conventional state-of-the-art object detection model YOLO (Redmon et al., 2016) was deemed appropriate. Different versions of this one-stage-detector have already been applied in the field of object detection in 360°images (Chou et al., 2020; Yang et al., 2018), as well as being frequently used as a benchmark model (e.g. in Zhao et al. (2019, 2020)). Due to constraints on Graphics Processing Unit (GPU) capacity, the Yolov8s model pre-trained on the COCO dataset (Lin et al., 2014) was selected.

Model Training

This study adopted an iterative training and validation process proposed by Géron (2023): the training configurations were tweaked based on the performance on an independent validation set to avoid overfitting, wherein a model becomes excessively optimized for the test set and fails to generalize to novel data. The model was evaluated on the test set only after selecting a final configuration. This iterative strategy is simplified in Figure 6. The training runs, and their validation and final testing were locally implemented in PyTorch using an NVIDIA RTX A4000 GPU.

The training process for the YOLOv8 model aims to minimize the training and generalization errors. In each run, the cost or loss function compares the predicted bounding box outputs with the actual outputs on the training set (training loss) and the validation set (validation loss). The generalization error is calculated as the difference between validation and training losses and is a key metric indicating the model's ability to perform well on new, unseen data.

The model's training performance, speed, and accuracy depend on using various hyperparameters and configurations that need manual configuration. This study adopted the validation approach proposed by Smith (2018) that examines the learning rate, batch size, momentum, weight decay, cyclical learning rates (cosine learning rate scheduler), and cyclical momentum. Accordingly, training and validation loss are analyzed during training to detect indications of underfitting and overfitting and determine the optimal combination of hyperparameters.

An initial configuration was used as a benchmark to adjust the training arguments, i.e., learning rate, batch size, momentum, weight decay, cyclical learning rates (cosine learning rate scheduler), and cyclical momentum. Nine configurations were trained and validated. The models were assessed by their mean average precision (mAP), precision, recall, and F1 score. A final configuration with the following hyperparameters was selected: a learning rate of 0.0007, weight decay of 0.001, the use of a cosine annealing learning rate scheduler, dropout regularization, and the Adam optimization algorithm.

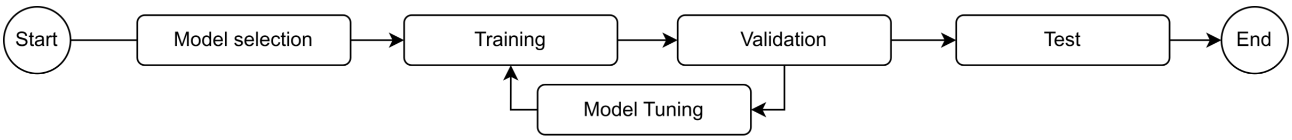

Figure 6: The applied iterative model training approach.

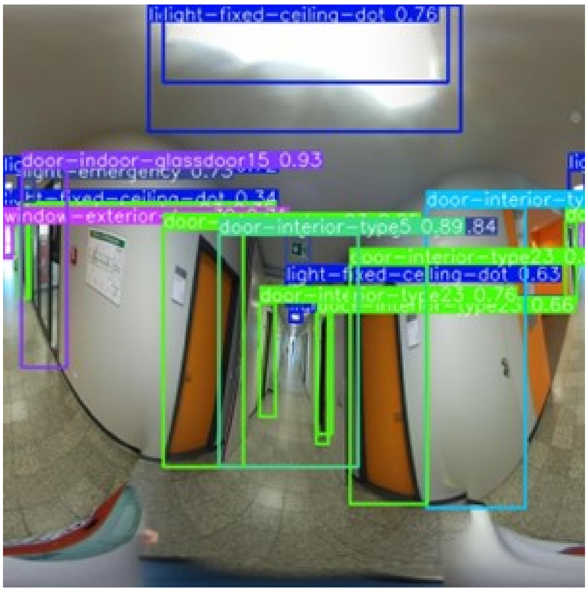

Figure 8: The model output with colored bounding boxes according to their label and superscribed with the confidence of the selected label.

Model Test

Finally, the chosen model configuration is tested on the hold-out test set, and its metrics are evaluated. The test run was performed on the test subset of 116 images and 975 instances. An example of the model's output in the test run can be seen in Figure 8. The model achieved an mAP at an intersection over union (IoU) threshold of 0.5 (mAP50) of 0.634, indicating a satisfactory overall performance in correctly detecting and localizing objects (see Table 3).
Furthermore, precision and recall have higher values than in the validation (see Figure 9). The PR curve reveals that a significant number of class labels exhibit a comparatively high PR score, indicated by the curves above the blue average PR score line, and a subset of classes with considerably inferior PR curves. This combination leads to an overall mAP@50 of 63.4% for all classes.
It indicates a model demonstrating relatively high precision at the beginning and overall good performance in precision and recall trade-off. In the class-wise evaluation of the model, in the worst-performing classes, zero instances were detected, resulting in low precision, recall, and mAP scores. This indicates that the model does not effectively detect and localize instances of these classes. These classes may have visual attributes that are difficult for the model to distinguish, leading to recognition errors. Furthermore, insufficient training data for these classes could also contribute to poor performance, as the model is not sufficiently exposed to different examples for effective learning.

Table 3: Test set performance metrics.

Class	Images	Instances	Box(P	R	mAP50	mAP50-95)
all	116	975	0.721	0.596	0.634	0.371

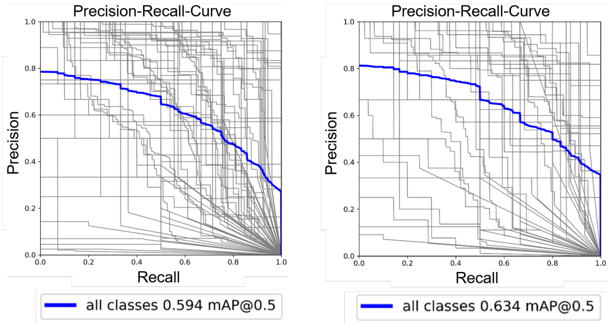

(a) Precision-Recall-Curve in training

(b) Precision-Recall-Curve in test-run

Figure 9: Performance metrics of the test run.

The diagonal in the confusion matrix (Figure 10) suggests that most classes are confidently predicted. However, the outliers indicate that some classes are incorrectly assigned to another specific class with high confidence. This can be recognized by a single point in a row and the coloring of the point in the graph. For these classes, more images or higher quality images are needed so that the model extracts features that lead to differentiation from the other component types more efficiently. Certain classes, such as class 135 (window-interior-type8) and class 136 (window-interior-type9), lack fundamental distinguishing features, which is visible in the erroneous assignment to very different component types (doors, lights, windows).

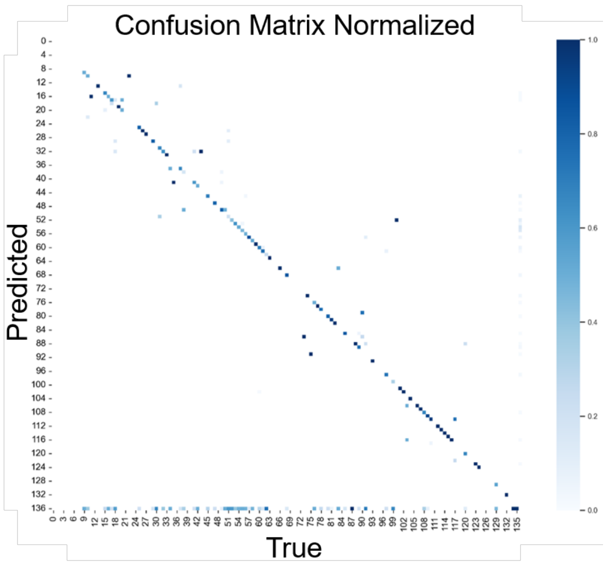

Figure 10: Normalized confusion matrix in test-run. Class IDs 1–52 = door types, 53–59 = lights, 60–64 = sanitary objects, 65–136 = window types.

Discussion

The implemented YOLOv8 object detection model shows that satisfactory accuracy can be achieved using conventional object detection models on equirectangular projections for differentiating component types within a class. The test metrics indicate effective assignment to building component categories (see Fig. 4), with most errors occurring in differentiating component types within the same component class. For instance, misclassified windows are mostly labeled as another window type rather than a different class, like doors or lamps. The model's overall performance indicates that the differentiation into different component types needed for the reusability assessment is possible. However, refining misclassifications, minimizing false positives and negatives, and augmenting training data for complex classes would improve its accuracy in recognizing component types.

A limitation of this study is the component-level annotation, which is time-consuming and, thus far, only includes component types from the case study. Using the object detection model trained on the TUDataset for other buildings will face challenges in detecting new building component types. Therefore, for subsequent research, the dataset annotation will be edited to be broader by eliminating further typification according to design and material. Only the categories of the DIN276 (see Figure 4 should be considered: e.g., exterior windows, skylights, dome lights, exterior doors (entrance, emergency), interior doors (room doors, WC doors), toilets, sinks, radiators (vertical, horizontal), lights (fixed luminaires), and safety lights. Combining existing labels into broader ones will introduce further intra-class variation, promising a better generalizing model and ensuring applicability to new buildings and EoL situations. While a simplified annotation process is considered to facilitate the up-scaling of the proposed approach to further building inventories, future studies to assess and validate scalability to other building contexts and larger datasets are needed.

This study aimed to detect the component category and type. However, more as-built information, such as the components' dimensions, material, and condition, should be extracted for reuse assessment. Therefore, photogrammetry and further feature extraction will be explored in upcoming research.

Furthermore, this study only used the object detection model to validate the approach. Further research will focus on the functionality and performance of the model. Different established object detection models should be compared in terms of robustness in dealing with large datasets and detection performance for the final model selection. Furthermore, various approaches integrating spherical layers into convolutional neural networks have been proposed in research (Cohen et al., 2018; Li et al., 2023; Zhao et al., 2019). Considering these additional models is promising for effectively addressing the distortions in ERP images and potentially improving the detection performance.

Finally, some limitations stem from the case study selection. First, resource constraints required a limitation on the included architectural styles, potentially compromising the generalizability of the results to buildings from different eras. Therefore, the selection criteria should be reviewed to integrate diversity of architectural styles and building functionalities, and the TUDataset should be continuously adapted to represent the variety of existing building stock. Second, the selection of component categories examined in this research represents only a small subset, and a broader range of non-destructive, non-toxic component categories could potentially be reintegrated into a circular economy. Third, it is based on the assumption that all areas and spaces within a building are accessible. Consequently, only the components captured by the camera can be identified. Therefore, the data introduces uncertainties in the number of components similar to those in traditional on-site inspections and component identification processes.

Conclusion

The main contribution of this research is the introduction of an approach to generate well-suited 360°image datasets for circular component reuse. The selected case study, centered around five buildings on the TU Dresden campus, served as a representative scenario for end-of-life sites. The dataset generation process involved using a specialized OpenExperience 360°camera helmet, capturing 1112 relevant images at different times of the day and under varying weather conditions. Data preparation, including cleaning, annotation, set partition, and pre-processing, was conducted to ensure the technical and relevance requirements of datasets.

The resulting fully-labeled TUDataset consists of approximately 2,400 panoramic images in equirectangular projection and 136 object classes, facilitating research and practical applications in object detection. Data quality and relevance requirements were addressed in the case study's design and during the dataset generation process. The YOLOv8 object detection model validated the real-world dataset, achieving a mAP50 of 0.634 on the test set.

The workflow developed in this research can be extended to other reusable building components, thereby contributing to developing new cost-effective and scalable approaches for building component recovery. Exploiting machine learning and 360°images for the inventory and reuse assessment could save time and costs, making circular strategies more competitive.

Overall, this study demonstrates that 360°images generated on site are suitable for effectively detecting different building component classes (windows, doors, sanitary, lamps, radiators), and provides the foundation for the next steps toward automated reusable component detection.

Acknowledgements

We thank the Technical University of Dresden for facilitating this study's data generation and implementation.

References

Barazzetti, L., Previtali, M., and Roncoroni, F. (2018). Can we use low-cost 360 degree cameras to create accurate 3D models? pages 69–75. Copernicus GmbH.

Chou, S.-H., Sun, C., Chang, W.-Y., Hsu, W.-T., Sun, M., and Fu, J. (2020). 360-Indoor: Towards learning real-world objects in 360° indoor equirectangular images. Proceedings - 2020 IEEE Winter Conference on Applications of Computer Vision, WACV 2020. Publisher: Institute of Electrical and Electronics Engineers Inc.

Cohen, T. S., Geiger, M., Koehler, J., and Welling, M. (2018). Spherical CNNs. Proceedings of the International Conference on Learning Representations.

Deutsches Institut für Normung (2018). DIN 276 Kosten im Bauwesen.

European Council and European Parliament (2008). Directive 2008/98/EC of the European Parliament and of the Council of 19 November 2008. Official Journal of the European Union, (L312).

Gordon, M., Batallé, A., Wolf, C., Sollazzo, A., Dubor, A., and Wang, T. (2023). Automating building element detection for deconstruction planning and material reuse: A case study. Automation in Construction, 146:104697.

Géron, A. (2023). Hands-on machine learning with Scikit-Learn, Keras, and TensorFlow. O'Reilly, Beijing; Boston; Farnham, third edition, oktober 2022 edition.

Hillebrandt, A., Riegler-Floors, P., Rosen, A., and Seggewies, J.-K. (2018). Atlas Recycling. Edition Detail. Detail Business Information GmbH, München, er ste auflage edition.

Honic, M., Kovacic, I., and Rechberger, H. (2019). Improving the recycling potential of buildings through Material Passports (MP): An Austrian case study. Journal of Cleaner Production, 217:787–797. Publisher: Elsevier BV.

Iacovidou, E., Purnell, P., and Lim, M. K. (2018). The use of smart technologies in enabling construction components reuse: A viable method or a problem creating solution? Journal of Environmental Management, 216:214–223.

Li, M., Meng, M., and Zhou, Z. (2023). RepF-Net: Distortion-Aware Re-projection Fusion Network for Object Detection in Panorama Image. In Wang, L., Gall, J., Chin, T.-J., Sato, I., and Chellappa, R., editors, Computer Vision – ACCV 2022, pages 508–523, Cham. Springer Nature Switzerland.

Lin, T.-Y., Maire, M., Belongie, S., Bourdev, L., Girshick, R., Hays, J., Perona, P., Ramanan, D., Zitnick, C. L., and Dollár, P. (2014). Microsoft COCO: Common Objects in Context.

OpenExperience (2023). DIGIBAU 360° Helmkamera.

Raghu, D., Bucher, M. J. J., and Wolf, C. (2023). Towards a 'resource cadastre' for a circular economy – Urban-scale building material detection using street view imagery and computer vision. Resources, Conservation and Recycling, 198:107140.

Redmon, J., Divvala, S., Girshick, R., and Farhadi, A. (2016). You only look once: Unified, real-time object detection. Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2016-December. Publisher: IEEE Computer Society.

Smith, L. N. (2018). A disciplined approach to neural network hyper-parameters: Part 1 – learning rate, batch size, momentum, and weight decay. US Naval Research Laboratory Technical Report.

Statistisches Bundesamt (2023). Weniger Abriss: 2022 fielen so wenige Wohnungen aus dem Bestand wie noch nie seit 1992. Pressemitteilung N050, Statistisches Bundesamt, Deutschland.

Uotila, U., Saari, A., and Junnonen, J.-M. (2021). Investigating the barriers to laser scanning implementation in building refurbishment. Journal of Information Technology in Construction, 26:249–262.

Witt, J. (2023). How to Improve Your Computer Vision Model.

Xiong, Z., Gordon, M., Byers, B., and Wolf, C. (2022). Reality capture and site-scanning techniques for material reuse planning. page 88.

Yang, W., Qian, Y., Cricri, F., Fan, L., and Kamarainen, J.-K. (2018). Object Detection in Equirectangular Panorama.

Zhang, Y., Song, S., Tan, P., and Xiao, J. (2014). PanoContext: A Whole-Room 3D Context Model for Panoramic Scene Understanding. pages 668–686. Springer, Cham.

Zhao, P., You, A., Zhang, Y., Liu, J., Bian, K., and Tong, Y. (2020). Spherical Criteria for Fast and Accurate 360° Object Detection. Proceedings of the AAAI Conference on Artificial Intelligence, 34(07):12959–12966.

Zhao, Q., Chen, B., Xu, H., Ma, Y., Li, X., Feng, B., Yan, C., and Dai, F. (2021). Unbiased IoU for Spherical Image Object Detection.

Zhao, Z. Q., Zheng, P., Xu, S., and Wu, X. (2019). Object Detection with Deep Learning: A Review. IEEE Transactions on Neural Networks and Learning Systems, 30(11):3212–3232. Publisher: Institute of Electrical and Electronics Engineers Inc.

Çetin, S., Wolf, C., and Bocken, N. (2021). Circular Digital Built Environment: An Emerging Framework. Sustainability, 13(11):6348.