Automated Equipment Recognition and Classification from Scattered Point Clouds for Construction Management Applications

J. Chen[a], Y. Fang[b], and Y. K. Cho[b]

[a] School of Electrical and Computer Engineering, Robotics Institute, Georgia Institute of Technology
[b] School of Civil and Environmental Engineering, Georgia Institute of Technology
E-mail: jchen490@gatech.edu, yihaifang@gatech.edu, yong.cho@ce.gatech.edu

Abstract –

Recognizing construction assets from as-is point cloud data of construction environment provides essential information for engineering and management applications including progress monitoring, safety management, supply-chain management, and quality control. This study proposes a fast and automated processing pipeline for construction target assets recognition from scattered as-is point clouds. The recognition tasks can be subdivided into object detection, which involves computing the bounding box around each construction equipment, and object classification, which involves labelling point cloud clusters from discrete equipment categories, such as backhoe loader, bulldozer, dump truck, excavator and front loader. The object detection step consists of point cloud down-sampling, segmentation and clustering. For the object classification step, machine learning methods were employed to determine class membership probability using features derived from the ESF (Ensemble of Shape Functions) descriptor. The classifiers were trained on synthetic point clouds generated from CAD (Computer-aided Design) models. The method was validated using laser scanned point clouds from an equipment yard. The test results demonstrate promising advancements towards semantic labelling and scene understanding of point cloud data.

Keywords –

Object recognition; Object classification; Scattered point clouds; Laser scanning; Machine learning.

1 Introduction

In the past decades, construction practices and research have been actively embracing emerging technologies to improve project performance in productivity, quality and safety. The technical categories that have been extensively studied and tested in the construction industry include data sensing, simulation, information modeling, and visualization. From these categories, as-built data sensing and visualization are considered by many industry practitioners and academia experts as one of the most promising technologies that will greatly expand the utilization of site information [1][2]. For instance, 3D point clouds produced by laser scanners and other data acquisition technologies have been widely used for acquiring and generating as-built site information to support applications such as construction quality assessment and control [3][4], construction progress tracking [5][6], building energy analysis [7], construction hazard recognition [8], structural health monitoring [9] [10] and highway asset management [11] [12]. Each application has a different focus of target objects (e.g., building components, materials, equipment, workers, and traffic signs) on the as-built point clouds.

Regardless of the applications, recognizing objects from point clouds is often the very first step for post data processing and analyses. Although object recognition is a well-studied problem in the field of computer vision, recognizing objects from point clouds of an uncontrolled environment such as construction sites remains a challenging and unsolved task. Many previous algorithms for object recognition were validated based on complete scans of small objects (e.g. fruit, household objects) obtained in controlled lab settings. In uncontrolled settings with the presence of occlusion and other objects that share similar features with the target object, many state-of-the-art recognition algorithms fail to maintain their good recognition rate and speed [13]. These problems become even more serious when acquiring as-built 3D data from construction sites where the presence of huge building structure and a large amount of equipment and materials introduce potential occlusions to the point cloud acquired regardless of what technology is used. As construction projects become larger and more complex, obtaining complete 3D data becomes even more difficult [14]. In addition, similarities among different types of construction equipment (e.g., bulldozer vs. loader) present a formidable challenge to object recognition algorithms that aim to distinguish between them.

This study proposes an effective and efficient object recognition pipeline that is robust to incomplete point clouds and shape variances of construction assets including structures, materials, and equipment. This paper starts with a thorough literature review on as-built data acquisition and modeling, followed by a detailed description of methodology applied in the proposed approach. The validation of the proposed approach on two testbeds is introduced, and the limitations and future work on the proposed method are discussed.

2 Literature Review

2.1 As-built Data Acquisition in Construction

As-built data of building components and construction assets are important for various applications in construction engineering and management. Depending on different representations of as-built scene, approaches for acquiring as-built data can be categorized as 2D image or video based methods, and 3D range data or point cloud based methods. This section will review multiple technologies belonging to these two categories.

2.1.1 2D Image and Video-based Methods

Recognizing objects in 2D images or videos has been extensively studied in the field of computer vision. Various techniques are available for object recognition, pose identification, and location tracking from 2D images or videos captured from of-the-shelf cameras. Given its advantages of low cost and minimal infrastructure, 2D image or video based techniques have been applied in construction for various operation-level applications including construction progress monitoring [5], construction entity detection and tracking [15], and productivity and safety analyses [16][17]. Various techniques have been explored to recognize objects (e.g., equipment, workers) on construction sites from 2D image or video, such as hue, saturation, and value (HSV) color space [18], background subtraction [19], histogram of color (HOC) [20], and template matching [21].

2.1.2 3D Range Data and Point Cloud based Methods

Compared to 2D presentation of construction site, 3D as-built data provide additional information and flexibility for visualization, as-built modeling, and other applications that heavily depend on spatial data. The most popular technologies for obtaining as-built 3D data on construction sites are 3D range imaging camera, photogrammetry, and laser scanning. This section will briefly review the characteristics of these three technologies in their applications in construction environments.

Based on different techniques such as stereo triangulation, structured light, and time-of-flight, a 3D range imaging camera produces a range image with pixel values representing the intensity and range. Although the 3D range imaging camera has been widely applied in the fields of virtual reality and gaming, very few investigations of its benefits in construction have been undertaken. Such research focused on range data processing for real-time site spatial modeling [22]. Compared to laser scanners, 3D range imaging cameras are more portable and less expensive but they operate in a much closer range, and do not provide reliable range images when direct sunlight is present, which makes it difficult to apply on large construction sites [23][24]. While understanding the positive and negative aspects of each data acquisition technology in a construction environment is important for real-time applications, it is also vital to consider the difficulties of data processing and extracting useful semantic information out of as-built data, regardless of the data acquisition methods.

2.2 Object Recognition from a Point Cloud

The typical pipeline of point cloud processing for object recognition consists of the following steps. First, shape descriptors are computed for each object instance or class to be recognized and stored in an offline database. Second, point cloud data is examined by an object detection algorithm to identify the contained object instances or object classes and query descriptors are computed for each detected object. Finally, objects from the model database with high descriptor similarity are aligned with the objects recognized in the point cloud to produce a match [1]. A major disadvantage of instance-based object recognition is that it is not as applicable and effective when significant shape variance exists among objects [1]. The most common approach in class-based object recognition is to use global shape descriptors which are less discriminative than semi-local descriptors but better adapted to shape variance. However, a major challenge is that this approach cannot handle partial or incomplete point cloud caused by occlusion or clutter [25]. To resolve this problem, one solution is to use descriptors that are more robust to shape changes [25]. To recognize mobile construction assets such as workers and equipment, [7] proposed a projection-recognition-projection (PRP) method to automatically recognize construction equipment from the point cloud. The 3D point cloud is projected to a 2D space where the geometric features represented by a local SURF descriptor are compared to a prepared template database for recognition. This method is very effective and efficient for recognizing target objects that are known to be present on the construction site. For unknown objects with high shape variance, however, the performance of this method is limited.

2.3 Shape Distribution and Global Shape Descriptor

The concept of shape distribution was introduced by [26] to describe the geometric signature of an object. Shape distributions are sampled from multiple shape functions that measure global geometric properties of an object. Distance between two randomly chosen points (D2), for instance, is a robust shape function for distinguishing different object classes. It should be noted that the performance of this method can be compromised when the number of classes to be recognized increases or only a partial view of the object is available in the point cloud. [27] proposed a robust global shape descriptor named ESF (Ensemble of Shape Function) for recognizing a variety of object classes (e.g., mugs, fruits, cars) from point clouds. This descriptor specifically addresses the challenge of using partial views in incomplete point clouds for object recognition through a construction of three distinct shape functions describing distance, angle, and area distributions. Each shape distribution is represented by an ESF histogram that can be used to efficiently retrieve the k-closest matches from a pre-defined class database using the k-nearest neighbor algorithm (k-NN). Although the ESF descriptor shows fairly good recognition rates for small-scale objects such as apples (98.45%) and mugs (99.46%), it shows limited capability in recognizing large objects such as cars (43.64%).

2.4 Training in Object Recognition

In addition to a robust description of objects, another critical component in object recognition is training the object recognition algorithm. Typically, the training process consists of the following four steps: 1) gather a training set which consists of a number of point clouds of the object classes that need to be recognized and label the point clouds with desired classifications; 2) determine a descriptor to represent the objects in the training set; 3) select the structure of the learned function and corresponding learning algorithm (e.g., support vector machine, k-nearest neighbor, decision tree); and 4) determine training parameters and run the learning algorithm on the gathered training set. Thus, a new input point cloud can be recognized by mapping its descriptor to the function inferred by the previously labeled training data. Similar to other applications adopting supervised learning, the availability and collection of training set is a bottleneck due to the lack of a centralized repository for point clouds. This problem can be addressed by using synthetic point clouds generated from 3D CAD models [27].

3 Methodology

To address the limitations in existing object recognition techniques, this study proposes a fast and automated processing pipeline for construction target assets recognition from scattered as-is point clouds. Figure 1 shows the processing pipeline for our object recognition approach. The pipeline is divided into an object detection stage and an object classification stage. The two stages are further discussed in detail in the following sections.

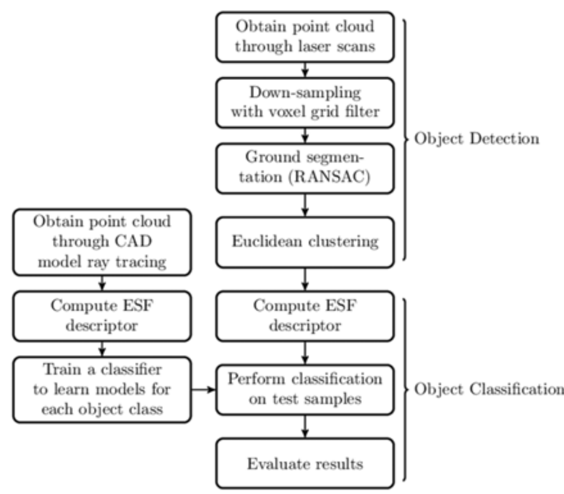

Figure 1: Flowchart for our object recognition approach

3.1 Object Detection in Laser Scanned Point Cloud

Laser scanned point clouds from construction sites often consist of millions of points, which renders the post-processing step computationally demanding. Thus, a voxel grid representation is used to store point cloud data in condensed form in order to achieve an efficient object detection pipeline. Next, the RANSAC (Random Sample Consensus) technique [28] is employed to filter out points belonging to the ground plane. In cases where the ground is a uniform flat plane or approximately flat, a single RANSAC pass is sufficient to obtain plane parameters while rejecting outliers. For cases where the ground consists of multiple planes with different gradients, multiple RANSAC passes are required to detect each plane. Finally, the Euclidean distance metric [29] is used to combine neighboring points into point cloud clusters for each object.

3.2 Object Classification

The object classification stage involves generating a training database of point cloud samples from objects with known categorization and matching query point cloud samples to the appropriate category based on computed shape descriptors. Further details concerning the classification process are discussed in the following sections.

3.2.1 Synthetic Point Cloud Generation

Sample 3D CAD models of our desired classification categories were downloaded online from 3D Warehouse. In this study, construction equipment samples were collected from *backhoe_loader*, *bulldozer*, *dump_truck*, *excavator (backhoe)*, and *front_loader* categories (Figure 2). A ray casting technique was used to sample points along the model surface with respect to virtual laser scanners placed at multiple view locations around the model. For each equipment model, point clouds were generated from multiple view angles and the points were perturbed with Gaussian white noise to simulate the data collection process in real world scenarios. This can also be viewed as a form of regularization which allows the classifier to be robust to data variability.

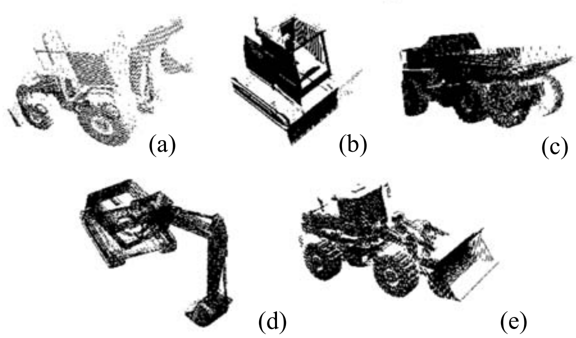

Figure 2: Synthetic point clouds of construction equipment: (a) backhoe loader, (b) bulldozer, (c) dump truck, (d) excavator, (e) front loader

3.2.2 ESF Descriptor

The ESF descriptor describes a 3D point cluster with a 640-dimensional set of features encompassing angle, area and distance shape functions. ESF is an example of a global shape descriptor since it encodes information about an object by considering points in the entire point cloud without making any prior assumptions about the object geometry. In our study, ESF was selected over competing descriptors such as VFH (Viewpoint Feature Histogram) or SHOT (Signature of Histograms of Orientations) due to its better performances in speed and robustness with respect to noise and partial views. Figure 3 shows the distribution of descriptor values of five construction equipment categories considered in this study. The x-axis describes each dimension of the descriptor whereas the y-axis represents the descriptor value at that dimension. ESF descriptor calculation is performed using the C++ implementation provided in the PCL (Point Cloud Library) software suite [30].

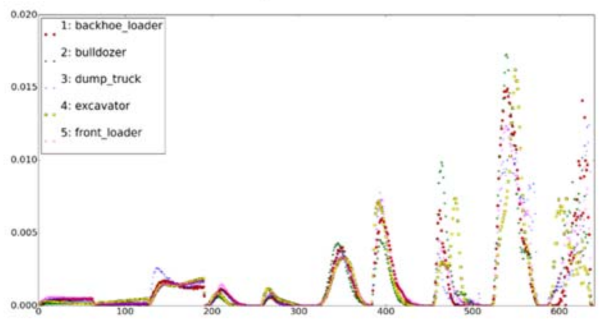

Figure 3: ESF descriptor distribution with separate plots representing the mean descriptor values for each category

3.2.3 Machine Learning

In the final stage of object classification, class labels are assigned to each input point cloud cluster by training and applying machine learning classifiers. In addition to the k-nearest neighbor (k-NN) algorithm used in the original ESF descriptor study [27], we also considered the approach of using discriminative classifiers such as logistic regression [31] and support vector machine [32]. The k-NN method involves picking *k* training samples closest to the test sample in terms of Euclidean distance in descriptor space. The output class label is then determined through uniform voting over class labels of the *k* nearest training samples. On the other hand, the logistic regression approach estimates the class label probability using logistic functions and learns a set of weights for descriptor elements from the training samples. Whereas, the support vector machine (SVM) method involves constructing a set of separating hyperplanes in high-dimensional space to perform classification. In our implementation of SVM, the linear kernel was used instead of the RBF (Radial Basis Function) kernel since the number of features in the descriptor is large. To achieve multi-class classification, the one-vs-all and one-vs-one scheme was used for logistic regression and SVM respectively.

After training and testing samples are obtained, the corresponding ESF descriptor is calculated for each point cloud cluster. For both training and testing, the descriptor values are scaled uniformly to a range of 0 to 1 and provided as input to each machine learning classifier. The output class labels and class probabilities are computed and stored for evaluation purposes. In this study, the training and testing processes were implemented using the Python toolkit *scikit-learn* [33].

4 Results and Validation

The validation dataset is composed of laser-scanned point clouds collected from a construction equipment yard. The data is collected from two sites with seven scans each and fused into a single point cloud of around 8 million points using registration techniques. The recognition results are shown in Figure 4. Point cloud clusters associated with recognized objects are identified with a bounding box labelled with the object class and computed class probability. Quantitative evaluation of the classification performance is discussed in the next section.

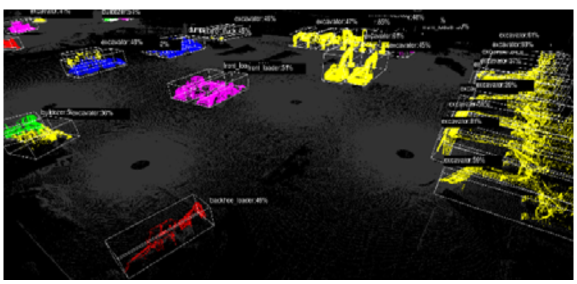

Figure 4: Testbench sample recognition results

In the object classification stage, recognition performance was quantified using the recall rate metric, which is the number of correctly identified samples divided by the total number of samples for each object class. Table 1 demonstrates that Logistic Regression and SVM classifiers [32] perform better compared to the original technique of k-nearest neighbors proposed by Wohlkinger and Vincze (2011) with respect to the number of correctly classified samples for each object category. The SVM classifier shows the highest overall recall rate for both and thus is used for all subsequent experiments.

Table 1: Number of correct classifications based on classification technique

Category	Sample number	K nearest neighbor	Logistic Regression	SVM
Backhoe Loader	3	1 (33%)	2 (67%)	2 (67%)
Bulldozer	3	0 (0%)	3 (100%)	1 (33%)
Dump truck	2	0 (0%)	1 (50%)	2 (100%)
Excavator	23	23 (100%)	19 (83%)	22 (96%)
Front Loader	9	3 (33%)	4 (44%)	4 (44%)
Total	40	27 (68%)	29 (73%)	31 (78%)

The effect of variables such as size of training data and point cloud resolution on the recall rate can also be investigated. Table 2 shows that the recall rate shows a general increase as the number of samples of training data increases. Table 3 shows that the recall rate increases with the resolution of both the training data and the testing data. This suggests that the object recognition outcome could theoretically be improved given a larger amount of training data and higher resolution laser scans.

Table 2: Recall rate for different sizes of training data

Training Size (Number of samples)	Recall Rate
48	7%
96	62%
144	60%
192	57%
240	68%
288	75%
336	78%

Table 3: Recall rate for different combinations of training data resolution and testing data resolution

Testing res. / Training res.	Downsampled by 20%	Downsampled by 5%	Original
40 scan lines	55%	53%	60%
200 scan lines	68%	70%	75%
400 scan lines	70%	73%	78%

5 Discussion

In terms of recognizing construction equipment, we demonstrated an overall classification recall rate of around 78% for our equipment yard testbench. Recall rates was improved using discriminative classifiers such as logistic regression and SVM compared to the original technique of k-nearest neighbors [27]. We further identified the number of training data and point cloud resolution as variables that could affect the recall rate. Classification performance for construction equipment was analysed using the confusion matrix shown in Table 4. Overall, the recall rate is high since there is less than one instance of wrong classification for all categories except for *bulldozer* and *front_loader*. The *excavator* category showed high classification performance due to its distinctive boom and bucket shape. On the other hand, the *front_loader* category showed less effective classification due to high shape variance and thus is often confused with other categories such as *bulldozer* and *excavator*. Overall, the classification method is shown to be theoretically sound and robust. The processed outcomes would be more viable given a larger amount of training data and more complete laser scans.

Table 4: Confusion matrix for construction equipment classes

Predicted / Actual	Loader	Bulldozer	Dump truck	Excavator	Front Loader
Backhoe loader	2	1			
Bulldozer	1	1		1	
Dump truck			2		
Excavator				22	1
Front loader		2	1	2	4

6 Conclusion

Our main contributions in this study are to integrate pre-processing, detection and classification steps of object recognition and classification from scattered point clouds into a single efficient recognition pipeline and to develop improvements over existing methods. We found that scan resolution, scan completeness and the availability of real world training data are key variables positively correlated with the classification recall rate. We were able to achieve promising recognition rates even without a large and comprehensive data set, as common machine learning applications often operate in the big data realm. This was achieved through a more rigorous processing of training data and better classification techniques. For future work, we envision expanding the training dataset, fine tuning the shape descriptor, and extending the method to attach multiple semantic labels such as coarse and fine categorization to each object in a construction site to aid the current construction management practices.

7 Acknowledgement

This material is based upon work supported by the National Science Foundation (Award #: CMMI-1358176). Any opinions, findings, and conclusions or recommendations expressed on this material are those of the authors and do not necessarily reflect the views of the NSF. This research was also supported by the Chung-Ang University Research Grants in 2015.

References

[1] P. Tang, D. Huber, B. Akinci, R. Lipman, and A. Lytle, “Automatic reconstruction of as-built building information models from laser-scanned point clouds: A review of related techniques,” Autom. Constr., vol. 19, no. 7, pp. 829 – 843, Jul. 2010.

[2] F. Bosché, “Automated recognition of 3D CAD model objects in laser scans and calculation of as-built dimensions for dimensional compliance control in construction,” Adv. Eng. Informatics, vol. 24, no. 1, pp. 107–118, 2010.

[3] L. Bu and Z. Zhang, “Application of Point Clouds From Terrestrial 3D Laser Scanner for Deformation Measurements,” Int. Arch. Photogramm. Remote Sens. Spat. Inf. Sci., vol. XXXVII, pp. 545–548, 2008.

[4] M. Ahmed, C. T. Haas, and R. Haas, “Toward low-cost 3D automatic pavement distress surveying : the close range photogrammetry approach,” NCR Res. Press, vol. 1313, no. November 2011, pp. 1301–1313, 2012.

[5] M. Golparvar-Fard, F. Peña-Mora, C. a. Arboleda, and S. Lee, “Visualization of Construction Progress Monitoring with 4D Simulation Model Overlaid on Time-Lapsed Photographs,” J. Comput. Civ. Eng., vol. 23, no. December, pp. 391–404, 2009.

[6] Y. Turkan, F. Bosche, C. T. Haas, and R. Haas, “Automated progress tracking using 4D schedule and 3D sensing technologies,” Autom. Constr., vol. 22, pp. 414–421, 2012.

[7] Y. K. Cho and M. Gai, “Projection-Recognition-Projection Method for Automatic Object Recognition and Registration for Dynamic Heavy Equipment Operations,” J. Comput. Civ. Eng., vol. 28, no. 5, p. A4014002, Sep. 2014.

[8] S. Fekete, M. Diederichs, and M. Lato, “Geotechnical and operational applications for 3-dimensional laser scanning in drill and blast tunnels,” Tunn. Undergr. Sp. Technol., vol. 25, no. 5, pp. 614–628, 2010.

[9] H. S. Park, H. M. Lee, H. Adeli, and I. Lee, “A new approach for health monitoring of structures: Terrestrial laser scanning,” Comput. Civ. Infrastruct. Eng., vol. 22, pp. 19–30, 2007.

[10] D. Huber, “ARIA: the Aerial Robotic Infrastructure Analyst,” 2014. [Online]. Available: http://spie.org/newsroom/technical-articles/5472-aria-the-aerial-robotic-infrastructure-analyst. [Accessed: 07-Oct-2015].

[11] J. Gong, H. Zhou, C. Gordon, and M. Jalayer, “Mobile Terrestrial Laser Scanning for Highway Inventory Data Collection,” in Computing in Civil Engineering (2012), 2012, vol. Computing, pp. 545–553.

[12] S. Pu, M. Rutzinger, G. Vosselman, and S. Oude Elberink, “Recognizing basic structures from mobile laser scanning data for road inventory studies,” ISPRS J. Photogramm. Remote Sens., vol. 66, no. 6 SUPPL., pp. S28–S39, 2011.

[13] B. Matei, Y. Shan, H. S. Sawhney, Y. Tan, R. Kumar, D. Huber, and M. Hebert, “Rapid object indexing using locality sensitive hashing and joint 3D-signature space estimation.,” IEEE Trans. Pattern Anal. Mach. Intell., vol. 28, no. 7, pp. 1111–1126, 2006.

[14] C. Kim, C. Kim, and H. Son, “Automated construction progress measurement using a 4D building information model and 3D data,” Autom. Constr., vol. 31, pp. 75–82, 2013.

[15] M. W. Park, A. Makhmalbaf, and I. Brilakis, “Comparative study of vision tracking methods for tracking of construction site resources,” Autom. Constr., vol. 20, no. 7, pp. 905–915, 2011.

[16] J. Gong and C. H. Caldas, “An object recognition, tracking, and contextual reasoning-based video interpretation method for rapid productivity analysis of construction operations,” Autom. Constr., vol. 20, no. 8, pp. 1211–1226, 2011.

[17] S. S. Han, S. S. Lee, and F. Peña-Mora, “Vision-Based Detection of Unsafe Actions of a Construction Worker: Case Study of Ladder Climbing,” J. Comput. Civ. …, vol. 27, no. 6, pp. 635–644, 2013.

[18] J. Zou and H. Kim, “Using Hue , Saturation , and Value Color Space for Hydraulic Excavator Idle Time Analysis,” vol. 21, no. August, pp. 238–246, 2007.

[19] S. Chi and C. H. Caldas, “Automated Object Identification Using Optical Video Cameras on Construction Sites,” Comput. Civ. Infrastruct. Eng., vol. 26, no. 5, pp. 368–380, 2011.

[20] M. Memarzadeh, M. Golparvar-Fard, and J. C. Niebles, “Automated 2D detection of construction equipment and workers from site video streams using histograms of oriented gradients and colors,” Autom. Constr., vol. 32, pp. 24–37, 2013.

[21] L. Cole, D. Austin, and L. Cole, “Visual Object Recognition using Template Matching,” Australas. Conf. Robot. Autom., 2004.

[22] J. Gong and C. H. Caldas, “Data processing for real-time construction site spatial modeling,” Autom. Constr., vol. 17, no. 5, pp. 526–535, 2008.

[23] M.-W. Park, C. Koch, and I. Brilakis, “Three-Dimensional Tracking of Construction Resources Using an On-Site Camera System,” J. Comput. Civ. Eng., vol. 26, no. 4, pp. 541–549, 2012.

[24] F. Remondino and D. Stoppa, TOF Range-Imaging Cameras | Fabio Remondino | Springer. Springer, 2013.

[25] S. Ruiz-Correa, L. G. Shapiro, M. Meila, G. Berson, M. L. Cunningham, and R. W. Sze, “Symbolic signatures for deformable shapes,” IEEE Trans. Pattern Anal. Mach. Intell., vol. 28, no. 1, pp. 75–90, 2006.

[26] R. Osada, T. Funkhouser, B. Chazelle, and D. Dobkin, “Matching 3D models with shape distributions,” Proc. - Int. Conf. Shape Model. Appl. SMI 2001, pp. 154–166, 2001.

[27] W. Wohlkinger and M. Vincze, “Ensemble of shape functions for 3D object classification,” 2011 IEEE Int. Conf. Robot. Biomimetics, pp. 2987–2992, 2011.

[28] R. Schnabel, R. Wahl, and R. Klein, “Efficient RANSAC for point-cloud shape detection,” Comput. Graph. Forum, vol. 26, no. 2, pp. 214–226, 2007.

[29] R. B. Rusu, “Semantic 3D Object Maps for Everyday Manipulation in Human Living Environments,” 2009.

[30] R. B. Rusu and S. Cousins, “3D is here : Point Cloud Library (PCL),” in IEEE International Conference on Robotics and Automation (ICRA), 2011.

[31] C. Lin, H. Yu, and F. Huang, “Dual Coordinate Descent Methods for Logistic Regression and Maximum Entropy Models,” Mach. Learn., no. 85, pp. 41–75, 2011.

[32] C. Chang and C. Lin, “LIBSVM : A Library for Support Vector Machines,” in ACM Transactions on Intelligent Systems and Technology, 2011, pp. 2:27:1–27:27.

[33] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, “Scikit-learn: Machine Learning in {P}ython,” J. Mach. Learn. Res., vol. 12, pp. 2825–2830, 2011.