

Road Design Layer Detection in Point Cloud Data for Construction Progress Monitoring

Steven VICK¹, Ioannis BRILAKIS²

¹ PhD Candidate, Department of Engineering, University of Cambridge, Trumpington Street, Cambridge, UK, CB2 1PZ, Email: sv364@cam.ac.uk

² Laing O'Rourke Reader, Department of Engineering, University of Cambridge, Trumpington Street, Cambridge, UK, CB2 1PZ, Email: ib340@cam.ac.uk

Abstract

Poor performance in transportation construction is well-documented, with an estimated \$114.3 billion in global annual cost overrun. Studies aimed at identifying the causes highlighted traditional project management functions like progress monitoring as the most important contributing factors. Current methods for monitoring progress on road construction sites are not accurate, consistent, reliable, or timely enough to enable effective project control decisions. Automating this process can address these inefficiencies. The detection of layered design surfaces in digital as-built data is an essential step in this automation. A number of recent studies, mostly focused on structural building elements, aimed to accomplish similar detection but the methods proposed are either ill-suited for transportation projects or require labelled as-built data that can be costly and time consuming to produce. This paper proposes and experimentally validates a model-guided hierarchical space partitioning data structure for accomplishing this detection in discrete regions of 3D as-built data. The proposed solution achieved an F1 Score of 95.2% on real-world data confirming the suitability of this approach.

Keywords: Transportation Construction; Progress Monitoring; Drones; Automation

INTRODUCTION

A recent study estimated that \$3.2 trillion per year in infrastructure investment is required to keep pace with global GDP growth through 2030, with approximately 16% of that investment (\$512 billion) required for road construction (Dobbs et al., 2013). Poor performance in this field is well-documented, with more than forty independent studies and government audits over the past three decades identifying as many as 83 separate causes for transportation construction cost overrun (e.g. Siemiatycki, 2009; Cantarelli et al., 2010; Memon et al., 2012; Salling & Leleur, 2015). One study

26 analysed projects in Great Britain, Denmark, Sweden, and Norway, concluding that 77% experienced cost overruns
27 with the average amount being 29% of original contract value (Salling & Leleur, 2015). Assuming these levels of
28 performance continue, there is a potential for \$114.3 billion in annual cost overrun on road construction projects
29 worldwide ($\$512B \times 77\% \times 29\%$). While the causes of these overruns are varied, traditional project management
30 functions have been highlighted as the most important contributors (González et al., 2014). Project control is one such
31 function that directly impacts project performance (Del Pico, 2013). It's a cyclical process consisting of progress
32 monitoring and implementation of corrective actions. Automating the first part of the project control cycle, progress
33 monitoring, can enable more timely and effective corrective actions aimed at keeping a project on time and within
34 budget.

35 Progress monitoring is the cyclical collection, analysis, and reporting of work complete at a given point in
36 time, compared to the planned worked complete. The main processes of progress monitoring are: (1) collecting as-
37 built data, (2) processing it into the form required for analysis, and (3) comparing it with the desired (or "as-planned")
38 state of the project. Current transportation progress monitoring practice is largely "non-spatial" involving no or very
39 sparse measurement of the physical progress on site. Project engineers, surveyors, crew foremen, and other on-site
40 personnel collect as-built data using paper or electronic checklists and daily reports, verbal updates, site photographs,
41 material delivery receipts, inventory reports, and subcontractor invoices. These manual data collection methods rely
42 on the experience, training, and subjective assessment of the individuals performing the inspection. Additionally,
43 collecting and parsing data from these disparate sources into meaningful activity-level inputs and compiling them into
44 project-level performance metrics is a time consuming and arduous task that consumes as much as 30-50% of a project
45 manager's time (Navon & Sacks, 2007). Progress determinations are made in increments of distance along the
46 centreline of the road. For example, a project utilizing a 20m increment might report completion of a layer to station
47 01+20 (i.e. 120m from the beginning of the road). One drawback of this approach is that progress (or lack of progress)
48 between the stations may be left unreported. Ultimately, non-spatial progress monitoring methods are too error-prone,
49 inconsistent, and burdensome to enable timely and effective implementation of corrective project control actions.

50 The civil engineering community recently examined using a variety of commercially-available spatio-
51 temporal sensors to support progress monitoring. Ultra-wide band (UWB) (e.g. Shahi et al., 2012), radio frequency
52 identification (RFID) (e.g. Razavi & Moselhi, 2012), barcode (Cheng & Chen, 2002), global positioning system

53 (GPS), and various combinations of these sensors (e.g. Razavi & Haas, 2010) have been used to track materials,
54 equipment, and other resources as indicators of progress. While these sensors can provide useful information, they
55 don't directly measure the physical progress on the site. Additionally, their installation and maintenance can be
56 burdensome for project management as tags and readers must be applied, monitored, and repositioned on a wide range
57 of resources as the project evolves.

58 An alternative approach that directly measures the progress on site uses spatial data collection technologies.
59 Spatial progress monitoring involves surveying the work performed, calculating the volumes, areas and/or distances
60 that form the basis of the analysis, and comparing the results with planned progress to determine project status. The
61 most common data collection technologies used are Robotic Total Stations, Global Navigation Satellite System
62 (GNSS) receivers, Light Detection and Ranging (LiDAR) scanners, and Aerial Photogrammetry (NCHRP, 2013).
63 However, the scale of most transportation construction projects dictates the use of methods capable of quickly
64 surveying large areas to an acceptable level of accuracy and resolution, defined as 1-5 cm three-dimensional (3D)
65 point accuracy and more than 150 points/m² (Olsen et al., 2013). Table 1 summarizes current technologies capable of
66 achieving at or near this level of performance. Note that performance and cost can vary due to a variety of factors for
67 each technology, so general ranges are provided based on a review of academic studies, technical reports, and
68 manufacturer specifications. Unmanned aerial photogrammetry stands out as a cost-effective data collection solution
69 capable of meeting the demands of spatial progress measurement on large transportation sites.

70 Although these data collection solutions exist, automatically recognizing progress in the as-built scene is a
71 difficult task. The authors decompose automatic progress monitoring for road construction into the following steps:
72 (1) aligning and registering the as-built data in a common coordinate system with the design model, (2) detecting 'if'
73 and 'where' each road design layer occurs in the as-built data, (3) comparing this with the as-planned status, and (4)
74 generating the progress report. Commercially-available solutions exist for the first step, typically using GPS and/or
75 surveyed control targets to place the as-built data in the design model's geographic coordinate system (e.g. Trimble,
76 2017; Pix4d, 2017). The second step, however, remains an unsolved problem and is the focus of this study. Note that
77 the authors define road layer detection as a process that identifies which, and exactly where, each design layer is
78 present in the as-built point cloud data (PCD). While no methods exist for automatically detecting road design surfaces
79 in as-built data, similar functions are performed in automated manufacturing quality control processes. Such systems

80 (e.g. GOM Inspect (2017)) compare as-built data from recently-manufactured components with 3D models to detect
81 quality deviations beyond acceptable levels. However, these applications benefit from a variety of controlling factors
82 that cannot feasibly be replicated on transportation construction sites. For example, they are designed to operate
83 indoors with known and stable lighting conditions. Additionally, the components analysed are relatively small and
84 can be carefully measured from all required perspectives.

85 Detecting road design layers in unlabelled spatial data requires a-priori information regarding the expected
86 elevation of each layer throughout the corridor. Civil Information Modelling (CIM) is a mature technology capable of
87 providing such information, and is gaining wider support from government agencies and contractors. The U.S. Army
88 Corps of Engineers, for example, requires use of CIM tools in design and construction of all horizontal civil works
89 projects (US Army, 2013). CIM models represent horizontal construction as combinations of triangulated surfaces
90 defined by dynamic design objects such as alignments, profiles, and design cross sections. This enables extraction of
91 individual model surfaces defining the finished grade of activity layers typically found in road construction. In current
92 practice, CIMs are used mostly for design visualization and automated guidance of earthwork equipment (Schneider,
93 2013). Since CIM model surfaces can be very large and complex, an interesting problem to consider is how to best
94 divide them into smaller sections to allow incremental progress detection.

95 The computer graphics, medical imaging, and robotics fields regularly deal with spatial subdivision of
96 complex 3D spaces. Rasterization and voxelization are two popular methods for accomplishing this task. Rasterization
97 uses a perspective-projection of 3D objects onto a 2D plane and then divides this projection into a regular grid. In
98 computer graphics a grid of picture elements (pixels) is used. Voxelization divides 3D space into a network of
99 regularly-spaced cuboid volume elements (voxels). A common voxel organizational structure known as the octree
100 recursively divides a cube-bounded 3D space into eight sub-cubes down to a specified subdivision threshold (criteria
101 are typically dimension-based or occupancy-based). Voxels are the smallest subdivision, or the 'leaves', of this
102 hierarchical tree structure. The octree structure allows rapid searching and indexing of 3D space, and implicitly
103 describes how various objects are distributed within complex 3D scenes (Watt, 2000). These benefits make octrees an
104 attractive option for use in detecting 3D design surfaces in as-built point cloud data, but the trade-off between accuracy
105 and efficiency must be kept in mind. For example, an octree of finer resolution (i.e. smaller leaf voxel size) more

106 accurately represents a 3D object but at the expense of increased memory usage and slower voxel traversal operations
107 (Watt, 2000).

108 In summary, the basic building blocks for implementing an improved automated progress monitoring
109 approach exist in current practice. LiDAR and photogrammetry are viable means for obtaining survey-quality as-built
110 data on large construction sites, while CIM provides the basis for a digital comparison of the as-built data with the as-
111 planned state. Detecting the CIM design layers in as-built data is the next unsolved problem in realizing this
112 automation. This paper proposes and evaluates a novel model-guided and sparse hierarchical space partitioning
113 approach for accomplishing road design layer detection in discrete regions of as-built point cloud data. The following
114 section reviews recent research relevant to this task. The proposed solution is then described in detail and evaluated
115 using both synthetic and real-world data from the construction of a residential road in Cambridge, UK. Finally, the
116 researchers summarize the experimental results, draw conclusions, and discuss future research goals.

117 **BACKGROUND**

118 The specific technical problem addressed in this paper is detection of road design surfaces in as-built data to support
119 automated progress monitoring. The following review evaluates related methods proposed in recent studies based on
120 their level of automation and applicability for layered road design surfaces.

121 A group of related studies focused on detecting building construction design objects using as-built PCD
122 collected by laser scanning (Bosché, 2010; Kim et al., 2013; Turkan et al., 2014; Kalasapudi et al., 2017),
123 photogrammetry (Golparvar-Fard et al., 2015; Tuttas et al., 2014; Braun et al., 2015), or both (Han et al., 2018). Design
124 object detection in the 3D data was accomplished in different ways. This study divides the automated design surface
125 detection approaches into two main categories: (1) point-correspondence and (2) space partitioning. The following
126 sub-sections review the literature in these two categories before examining the sparse research specifically related to
127 transportation design surfaces.

128 *Point Correspondence Approaches*

129 Point-correspondence approaches (Bosché, 2010; Kim et al., 2013; Turkan et al., 2014) generate as-planned point
130 clouds from 3D model elements and search for the nearest as-built correspondences within a threshold distance to
131 determine if an as-planned point is present in the as-built scene. Design surface triangulation vertices are sparse, so

132 as-planned point clouds must be either subsampled on the design surface to closely match the density of the as-built
133 cloud (e.g. Kim et al., 2013), or the as-built points must be projected onto the nearby surfaces before selecting the
134 nearest projection as the correct as-planned point location (e.g. Bosché, 2010). The researchers mostly used thresholds
135 on the number of confirmed point correspondences to determine if a model element was detected. These methods
136 showed promise for automatically detecting planned structural building elements in 3D PCD, particularly when
137 factoring in a-priori connectivity and construction sequencing information like Kim et al. (2013). However, they
138 focused on all-or-nothing detection of the individual components. Such an approach works well for detecting relatively
139 small components constructed within a progress monitoring cycle (e.g. a column), but fails to capture the kind of
140 incremental progress needed to support effective project control on road construction sites with very large design
141 surfaces.

142 *Space Partitioning Approaches*

143 Other researchers partitioned registered as-built/as-planned scenes and identified points falling into the partitioned
144 regions to facilitate object detection. Some used the boundaries of each structural element model (e.g. a square column)
145 to conduct this partitioning (e.g. Zhang & Arditi, 2013). More commonly used approaches partition the 3D space
146 containing the as-built/as-planned scene using voxels. Some (Tuttas et al., 2014; Braun et al., 2015) built an octree
147 over the as-built PCD and searched for points within orthogonal distance thresholds of rasterized planar segments (e.g.
148 rectangles or triangles) extracted from each structural element's constituent surfaces. Although this approach is not
149 "all-or-nothing," it does limit the detection decision to each planar surface segment that makes up the design object.
150 Consider the effect of this on a road design surface that has very large planar regions (think of a large stretch of flat
151 highway) as well as areas where the surface orientation changes in much smaller intervals (think of a banked highway
152 around a curve). In this case, the surface detection would happen in uneven and inconsistent intervals along the length
153 of the road, resulting in progress reports that are difficult to interpret and act upon.

154 An approach that more consistently partitioned larger and/or more complex scenes (Golparvar-Fard et al.,
155 2015) voxelized the entire 3D space containing the registered as-is/as-built scene before traversing the voxels (twice:
156 once each for as-planned and as-built data) from various unordered camera positions and labelling them based on the
157 presence of design elements and as-built points. This method made use of ubiquitous daily job site images for as-built
158 data collection, and even enforced visibility and visual consistency constraints to account for static and dynamic

159 occlusions in the scene. However, the uniform nature of the implemented voxelization is ill-suited for large and
160 closely-layered road design surfaces. Consider the notional case illustrated in Figure 1, showing a longitudinal (y-z)
161 2D section of a small asphalt road design region. Realistic design layer thicknesses are represented to scale with the
162 voxel size shown, approximating the design specifications for the road construction site evaluated later in this paper.
163 With a notional 10 cm leaf voxel size, the thin upper design surfaces occupy the same voxel in multiple instances. As-
164 built point and as-planned surface occupancy in this case would not be sufficient to identify which design surface is
165 found in the as-built scene, presenting a problem for the detection approach used by Golparvar-Fard et al. (2015).
166 Voxelizing the scene at 1/5 the previous resolution (2 cm shown in Figure 1) can address this concern, but results in
167 a 125x increase in the total number of voxels that need to be traversed and analysed. Using a smarter space partitioning
168 approach, like an octree that only subdivides branch cells containing as-planned data, can reduce this complexity but
169 other issues still remain. First, use of regular cuboid voxels small enough to distinguish between thin asphalt layers
170 requires extremely high-density as-built data. For example, a 2 cm leaf voxel size in Figure 1 would alleviate the
171 problem of voxels occupied by multiple design layers, but would require as-built data with a spatial resolution of at
172 least 2,500 pts/m² to ensure at least one point could be found in each voxel along the road corridor, not accounting for
173 noise and measurement error. This limits the data collection options to the most expensive and time-consuming
174 methods (Table 1) that are more susceptible to occlusions due to the ground-level perspective. Second, the cubic-grid
175 structure of the voxelization produces voxels that do not uniformly fit the design surfaces. In some instances, a surface
176 intersects only the corner of a voxel while in others it cuts directly through the middle. As a result, as-built points
177 corresponding to a specific design surface are only sought above the surface in some cases, while in others only the
178 region below is considered. This could results in further inconsistencies with the object detection decision, particularly
179 when voxelizing at finer resolutions to ensure unique as-planned voxel labels.

180 *Transportation-Specific Research*

181 Just one study sought to compare road design surface layers to as-built point cloud data (Kivimaki & Heikkila, 2015).
182 This study, however, stopped short of true design surface detection as the matching of surveyed points to the relevant
183 design surface was accomplished by assigning a specific code linking each point to a surface. Such point codes are
184 only possible on terrestrial survey equipment where the exact nature of each point can be coded into the machine prior
185 to taking the measurement. The authors concede that manual matching of point clouds to surfaces is required for data
186 that lack the point code information; a drawback on the efficiency and applicability of this approach.

187 *Gaps in Knowledge, Objective, and Research Questions*

188 The existing body of research is heavily-skewed towards structural building components in as-built PCD. Surfaces for
189 such components are typically characterized by relatively-small orthogonally intersecting planar sections (e.g. the
190 faces of a concrete foundation wall). Road design surfaces are larger and more complex, typically involving changes
191 in slope along both the longitudinal and transverse axes of the road corridor. The closely-layered construction,
192 particularly in the thin upper asphalt layers, also differentiates road design surfaces from building structural
193 components. These differences limit the applicability of existing building-focused detection methods on as-built data
194 for road projects. Considering the state of practice and body of research reviewed in the preceding sections, we identify
195 the following gaps in knowledge: (1) no method exists for automatically measuring progress on road construction
196 projects, (2) existing methods for automatically detecting design surfaces in as-built data are not well-suited for
197 identifying and distinguishing individual layered road design surfaces, (3) none of the proposed methods tackle
198 incremental progress detection, and (4) there is no formal understanding of the challenges and limitations associated
199 with adopting an automated progress monitoring approach on linear transportation projects.

200 The objective of this study is to address gaps 2-4 by developing and testing a novel data structure for
201 automatically and incrementally detecting layered road design surfaces in unlabelled as-built PCD. The authors
202 examine the following research questions to accomplish this objective: (1) how can the 3D as-planned space be
203 partitioned to allow for consistent and incremental road layer detection?, (2) how can as-built PCD be accurately
204 classified within the partitioned regions, accounting for noise and uncertainty in the as-built scene?, and (3) what
205 combination of design and/or input parameters will be required to produce the most desirable results?

206 **PROPOSED SOLUTION**

207 The authors propose a novel solution for detecting layered road design surfaces in discrete regions of as-built point
208 cloud data. This solution implements a new model-guided and sparse hierarchical space partitioning data structure
209 named *BrickTree*, a combination of the authors' names and a nod to the rectangular cuboid shape of the leaf voxels in
210 this structure. Figure 2 illustrates the key processes in implementing this approach. The inputs to this solution are (1)
211 an as-built point cloud, scaled in real-world units and registered in the CIM model's coordinate system, and (2)
212 triangulated road layer surfaces extracted from the CIM model. The rest of this section describes in further detail the
213 four processes of the proposed solution.

214 *The BrickTree Data Structure*

215 The first step in generating the *BrickTree* applies a rigid-body transformation to the layered design surfaces,
216 positioning them just above the global origin with the longitudinal direction of the road corridor aligned with the *y*-
217 axis and the transverse direction aligned with the *x*-axis. A 2D grid, named the *projectGrid*, is then created on the *x-y*-
218 plane below the registered and aligned design surface layers. The *projectGrid* serves as the basis of further partitioning
219 as each square grid cell defines the *x-y* boundaries of a *rasterBranch* in the hierarchical tree structure (Figure 3). As
220 such, each *rasterBranch* is a vertical projection of a *projectGrid* cell, and is defined within the hierarchy as a collection
221 of *leafVoxels*. *LeafVoxels* are defined as 3D rectangular space partitioning elements that subdivide each *rasterBranch*
222 region. The *gridCellSize* parameter determines the granularity of the incremental surface detection decision, which
223 occurs at the *rasterBranch* level where only one design layer should be identified for a given dataset. Note that this
224 could result in small classification errors in regions of transition between as-built layers, and selection of the
225 *gridCellSize* parameter should consider this possible error in determining the acceptable level of granularity for the
226 final progress report. The authors note that such errors (e.g. detecting a surface within a branch's boundaries when
227 only half of it is actually present) are common for all spatial partitioning methods, and generally deemed acceptable
228 if properly controlled considering the inaccuracy and subjectivity of current progress monitoring practice. The
229 *gridCellSize* is set to 0.5m and controlled for in the following experiments based on conversations with local project
230 management personnel that indicated this would be an acceptable level of granularity for the final progress
231 determination, allowing more detailed reporting of progress than the centreline increment methods used in current
232 practice.

233 Each *rasterBranch* is partitioned into *leafVoxels* at intersections with the layered design surfaces. The goal
234 of this model-guided voxelization is to ensure each *leafVoxel* enables consistent searching for as-built points in the
235 regions above and below their corresponding surface. This study constructs axes-aligned voxels, relying on the
236 assumption that road surfaces can be reasonably approximated as flat within the regions defined by the *gridCellSize*.
237 This simplifying assumption reduces the complexity of point-to-voxel assignment operations by limiting them to a
238 series of simple indexing calculations in the *x*, *y*, and *z*-directions. The degree to which this assumption affects the
239 performance of the proposed method and how well it generalizes to the possible combinations of road camber and
240 grade is the focus of further research by the authors, and thus is not reported in this paper.

241 To construct the *leafVoxels*, the authors orthogonally project the *projectGrid* onto each surface, storing the
 242 centroid and normal direction of each projected grid cell for the following steps. The z -value of this centroid coupled
 243 with the *searchDistance* parameter defines the height of the constructed *leafVoxels*, while the horizontal limits of the
 244 grid cells define the voxel boundaries in the x and y directions. Selecting the optimal *searchDistance* value is a focus
 245 of experiments in the following section, but an important consideration is the thickness of the thinnest design layer as
 246 this sets an upper limit on the value in order to limit voxel overlap. The final step in constructing the *BrickTree* sets
 247 the upper z -limit of the outer 3D bounding box as the maximum point elevation of the voxels in the top surface layer.
 248 With the *BrickTree* constructed, the proposed solution moves on to the next step: populating the *leafVoxels* with as-
 249 built data.

250 *Populating the BrickTree with as-built points*

251 There are two general strategies for assigning as-built points to a hierarchical space partitioning tree: (1) traverse the
 252 point cloud to determine where in the tree structure each point lies, or (2) traverse the tree structure to identify and
 253 assign points in the vicinity of each voxel. The axes-aligned *leafVoxel* structure allows points to be rapidly indexed
 254 and assigned to discrete 3D regions. For this reason, the authors chose to traverse the as-built cloud, directly assigning
 255 points to *leafVoxels*. Because the voxels are not evenly-spaced in the z -direction, rapid indexing only occurs at the
 256 branch level. For a given as-built point, p_b , the process calculates the branch index by comparing it to the origin
 257 (lower-left point) of the *projectGrid*, p_o , using the following equations where n_{i_cells} is the number of *projectGrid*
 258 cells in the x -direction. Note that the bracket notation used indicates floor-rounding to the nearest integer value.

259
$$i_{index} = \left\lfloor \frac{p_{bx} - p_{ox}}{gridCellSize} \right\rfloor \quad (1)$$

260
 261
$$j_{index} = \left\lfloor \frac{p_{by} - p_{oy}}{gridCellSize} \right\rfloor \quad (2)$$

262
 263
$$branch_{index} = i_{index} + j_{index} \times n_{i_cells} \quad (3)$$

264 With the branch identified, the method then compares the z -component of the built point, p_{bz} , to the z -limits of the
 265 *leafVoxels* to determine which, if any, to assign the point to. If an as-built point is detected inside a *leafVoxel*, the
 266 voxel is marked as 'detected'. Note that this initial labelling is reviewed and updated in the subsequent branch
 267

268 classification step. After the algorithm processes all as-built points falling within the *BrickTree* bounding box and
269 applies the *pointThreshold* to label the *leafVoxels*, the branch classification processes begins.

270 *Branch classification*

271 Noise and other errors could result in multiple *leafVoxels* being marked 'detected' in the same *rasterBranch*. A
272 classification decision must be made to correct this conflict as it's only possible to detect one design surface per branch
273 in an as-built scene. Sophisticated decision boundaries can be learned using supervised or unsupervised machine
274 learning approaches, but these require a large number of training datasets from a variety of construction sites. While
275 the authors aim to collect further data and explore machine learning based classification in future work, the current
276 study implements a simple statistical threshold decision boundary as a proof of concept for the overall framework.
277 Three possible rules are considered: maximum points (i.e. select the voxel with the most points), minimum distance
278 (i.e. select the voxel with the minimum average orthogonal distance between its as-built points and the design surface),
279 and minimum vertical variance (i.e. select the voxel whose points have the lowest variance in the z-direction).
280 Regardless of the decision rule used, the algorithm labels each *rasterBranch* with the surface identified and updates
281 the *leafVoxel* labels in the branch accordingly. The optimal classification rule out of the three considered is
282 experimentally evaluated later in this paper.

283 *Outlier correction*

284 The final step in the surface detection process, outlier correction, aims to improve the final result's accuracy using a
285 priori knowledge of road construction operations. Specifically, the method assumes that road surfaces are generally
286 not constructed with gaps or missing areas in the middle. The authors propose a neighbourhood consensus algorithm
287 to accomplish this task. Two input parameters, *neighbourDistance* and *consensusThreshold* are used. The
288 *neighbourDistance* parameter defines the size of the neighbourhood in which a branch classification is compared to
289 that of its neighbours. Figure 4 illustrates different construction cases for the branch comparison 'neighbourhood'.
290 Implementing the different construction cases ensures the comparison region size remains constant no matter where
291 the analysed branch lies on the *projectGrid*. The algorithm traverses the constructed neighbourhood, tallying branch
292 classifications by surface identification. If the ratio of any surface's tallied count to the total number of branches in
293 the neighbourhood exceeds the *consensusThreshold*, a consensus is declared in the neighbourhood. The algorithm
294 then compares and updates (if required) the analysed branch's classification and *leafVoxel* labels. The optimal

295 *consensusThreshold* value and size of the outlier correction window, defined by the *neighborDistance* parameter, are
296 explored in the experiments later in this paper.

297 *Construction and as-built data quality*

298 The proposed solution assumes an acceptable level of construction (ϵ_c) and as-built data measurement (ϵ_m) error. The
299 authors divide these errors into the horizontal ($\epsilon_{c,h}$, $\epsilon_{m,h}$) and vertical ($\epsilon_{c,v}$, $\epsilon_{m,v}$) directions to describe their impact on
300 the proposed method. Measurement errors also include noise, further divided into random noise, outliers, and
301 occlusions. The *BrickTree* structure implicitly limits the influence of outliers and occlusions by only considering
302 points within the *leafVoxel* regions. Points from low-level occlusions (e.g. equipment sitting on the road surface) could
303 still trigger detection errors, but this type of noise is difficult to control on an active site. Additionally, the outlier
304 correction step is designed to correct these errors as long as the occluded region is not much larger than the comparison
305 neighbourhood. Consequently the following analysis of required input data quality only considers random noise (ϵ_n),
306 which is assumed to be Gaussian and primarily affecting the vertical measurements.

307 Horizontal errors ($\epsilon_{c,h}$ and $\epsilon_{m,h}$) could result in detection errors, mostly along the boundaries of the road
308 corridor. The *gridCellSize* parameter defines the proposed solution's sensitivity to horizontal error, with the following
309 equation describing the relationship required to enable accurate detection:

$$310 \quad \epsilon_{c,h} + \epsilon_{m,h} < \textit{gridCellSize} \quad (4)$$

311 In practice, contract specifications dictate the acceptable level of construction error. The United Kingdom's Highways
312 Agency sets the tolerance for $\epsilon_{c,h}$ at 25 mm (*Manual of Contract Documents for Highway Works, Volume 1, Series*
313 *0700*). Using the 0.5 m *gridCellSize* implemented in this study and the maximum acceptable $\epsilon_{c,h}$, as-built data with
314 $\epsilon_{m,h}$ less than 47.5 cm is required, which is achievable using any of the technologies listed in Table 1.

315 Vertical errors ($\epsilon_{c,v}$, $\epsilon_{m,v}$, and ϵ_n) could result in detection errors within any branch of the *BrickTree*. The
316 *searchDistance* parameter drives the proposed solution's vertical error sensitivity, and is limited to a maximum
317 distance equal to the thickness of the thinnest road design layer to avoid overlapping *leafVoxels*. Equation 5 describes
318 the relationship between this parameter and the input data error required to enable detection:

$$319 \quad \epsilon_{c,v} + \epsilon_{m,v} + \epsilon_n < \textit{searchDistance} \quad (5)$$

320 Again, contract specifications drive the acceptable $\varepsilon_{c,v}$ levels, with the UK's Highway's Agency defining the limits at
321 +/- 6 mm for pavement, +/- 15 mm for base, and +/- 10 mm for subbase levels. Using the maximum allowable $\varepsilon_{c,v}$ of
322 15 mm and a notional 4 cm *searchDistance*, most points would have to be within a vertical measurement error ($\varepsilon_{m,v} +$
323 ε_n) of 2.5 cm to facilitate detection. In this case, unmanned aerial photogrammetry and low-level LiDAR are the only
324 feasible data collection methods per Table 1. The following section experimentally tests the viability of the proposed
325 solution and determines the optimal values for the various parameters and classification rules mentioned in this section.

326 **METHODOLOGY, EXPERIMENTS, AND RESULTS**

327 The authors generated simulated point cloud data to test the feasibility of the proposed solution prior to conducting
328 further verification experiments on real-world data. The simulated data allowed isolation and testing of the proposed
329 solution's performance in the presence of typical construction and as-built measurement errors, while controlling for
330 occlusions and clutter in the observed scene. Although occlusions are to be expected on a construction site, controlling
331 for them in these initial experiments allowed the authors to focus on the performance of the proposed solution under
332 the conditions in which it is designed to operate: when the measured surface is visible in the scene. Robustness to
333 expected levels of occlusion, and strategies for addressing errors caused by such occlusions will be addressed in further
334 research.

335 Unmanned aerial photogrammetry was selected as the data collection technology for this study, based on its
336 low cost, fast mobilization, and ability to meet required data quality thresholds. The authors developed a data
337 simulation module to produce the synthetic aerial photogrammetry point cloud, using a 3D surface model of an in-
338 progress road site as the input (Figure 5). The model, developed using AutoCAD Civil 3D, was a 500-meter long
339 undivided crowned road with -2% camber in the transverse direction to either side of the road's centreline, and 2%
340 grade in the longitudinal direction. Random construction errors were added to the surface layers within the allowable
341 limits discussed in the previous section.

342 *Simulated Aerial Photogrammetry Data*

343 The aerial data simulation module aims to produce point clouds that are qualitatively similar to real-world aerial
344 photogrammetry data in measurement error profile and point density. Error in photogrammetric surveying products
345 depends on a wide range of factors, to include sensor quality, image resolution, range (i.e. how close the sensor is to

346 the scene), focal length, angle of incidence, percent overlap between adjacent images, level of texture in the scene,
 347 lens distortion, and the quality of camera calibration (Dai et al., 2014). A recent study (Slocum & Parrish, 2017)
 348 proposed an in-depth computer graphics workflow for generating simulated aerial photogrammetry data; a method
 349 that shows promise for isolating and analysing the contributing factors in photogrammetric error. However, the
 350 complexity of this approach was deemed unnecessary for the purposes of this study, and the authors opted instead for
 351 a simpler approach that models photogrammetric error as a function of Ground Sample Distance (GSD); a descriptor
 352 that accounts for sensor quality, resolution, focal length, and range. Focusing on these factors is a reasonable
 353 simplification, as a number of other contributing factors can be controlled using effective data collection planning and
 354 established sensor calibration routines. GSD describes the size of an image pixel projected onto the observed ground
 355 surface, and is calculated as (Pix4d, 2017):

$$356 \quad GSD = \frac{w_s \cdot h}{w_l \cdot f} \quad (6)$$

357 where w_s is the width of the sensor in meters, w_l is the width of the image in pixels, h is the data collection height in
 358 meters, and f is the focal length in meters. A recent white paper on a state-of-the-art aerial photogrammetry system
 359 reported horizontal root-mean-square errors (RMSE) in the 1 – 7 pixel range ($\mu = 2.7$, $\sigma = 1.5$) and vertical RMSE in
 360 the 1 – 4.3 pixel range ($\mu = 2.4$, $\sigma = 1.1$) after analysing data produced at various heights and under differing
 361 weather/lighting conditions (Pauly, 2016). Using this as a guide, the aerial data simulation module generates each
 362 noisy point ($p_{i,\varepsilon}$) in the following manner:

$$363 \quad p_{i,\varepsilon} = p_i + \varepsilon_{xy} \hat{v}_{xy} + \varepsilon_z \hat{v}_z \quad (7)$$

364 where $p_i \in P$, $\varepsilon_{xy} \sim N(0, 2.7 \cdot GSD)$, and $\varepsilon_z \sim N(0, 2.4 \cdot GSD)$. Here, P is the set of points in the sampled cloud,
 365 \hat{v}_{xy} is a randomly-generated normalized vector in the xy -plane ($\langle x, y, 0 \rangle$) and \hat{v}_z is the vector $\langle 0, 0, 1 \rangle$.

366 To produce P , the algorithm randomly samples points on the input triangulated surface model (e.g. Figure 5)
 367 according to the user-specified density profile, defined by an input mean and standard deviation. This random
 368 sampling aims to mimic variations in point density attributable to changes in image texture while allowing the user to
 369 control the mode of these variations. This approach relies on the assumption that the as-built site contains sufficient
 370 texture for photogrammetric reconstruction, and that the density variations can be adequately described by a single
 371 normal distribution. Multimodal variations could be modelled by splitting the input surface into regions where

372 different materials are expected and assigning separate input parameters to each prior to simulation. Since this study
373 is primarily focused on detecting asphalt road design layers, the uniform image texture within these regions should
374 result in point densities that can be described by a unimodal normal distribution. The next section examines the
375 veracity of this assumption. The simulated data produced for the purposes of this study was generated using a GSD of
376 2 cm with a density distribution of $\mu = 200$ pts/m² and $\sigma = 1.5$.

377 *Real World Data*

378 The authors collected as-planned data on two separate days during the construction of a residential road for a new
379 development in Cambridge, UK. The asphalt road design included five distinct surfaces: (1) the formation (bottom of
380 excavation), (2) a 520 mm thick sub-base layer, (3) a 125 mm thick base course, (4) a 65 mm thick asphalt binder
381 course, and (5) a 40 mm thick asphalt wearing course. Autodesk's *Civil 3D* application was used to generate a 3D
382 corridor model from the 2D design information provided by the contractor. The wearing course thickness limits the
383 *searchDistance* parameter to a maximum of 4 cm, which requires most points to have a vertical measurement error of
384 less than 2.5 cm to enable accurate detection.

385 Trimble's UX5 fixed wing UAS conducted the aerial surveys, the details of which are provided in Figure 6.
386 The authors used the following steps to plan the data collection: (1) Select a desired GSD to achieve the required
387 accuracy, considering typical error ranges discussed in the previous section, (2) Select a desired data collection height
388 using Equation 6, taking into consideration local aviation authority regulations and on-site structures, equipment, and
389 obstructions, (3) Define the desired degree of overlap between adjacent images, (4) Plan the flight route to achieve
390 the parameters defined in the previous three steps, and (5) Establish ground control point (GCP) targets throughout
391 the site to enable accurate geo-referencing during post-processing. The authors used Trimble's UX5 flight planning
392 software to complete step 4, using a desired elevation of 70 m and target image overlap of 80% in both the longitudinal
393 and transverse directions. Eight GNSS-surveyed GCP targets were distributed throughout the site. The authors
394 performed photogrammetric post-processing of the collected images in Trimble's Business Center (TBC) aerial
395 photogrammetry application. The process involved: (1) importing the images and synchronized on-board sensor data,
396 (2) refining the sensor-defined camera positions using manual target observations in at least 3 images per GCP, and
397 (3) conducting a dense photogrammetric reconstruction to produce the final survey products. One issue that can affect
398 photogrammetric reconstruction quality on road sites is the potential for visual similarity between multiple regions

399 along the construction corridor, leading to errors in feature matching and calculation of external camera parameters.
400 The employed method accounts for this by logging the UX5's GPS, altimeter, and inertial measurement unit (IMU)
401 sensor readings at each image location, and then enforcing a feature-matching constraint that declares positive matches
402 can only occur between images with overlapping frames measured from their sensor-defined positions. The quality of
403 the collected data is analysed and compared to the synthetic data in the following section. Figure 7 depicts the point
404 clouds and ground truth conditions at the time of data collection.

405 *Comparing the Simulated and Real World Datasets*

406 The authors verified the simulated data's suitability by comparing it to the real world data, specifically examining the
407 point density and point-to-ground-truth-surface distance statistics. This was done by constructing histograms of the
408 data for each metric; normal distributions were then fit to the histograms to describe and visualize the results (Figure
409 8). Isolating photogrammetric measurement error in this case would require accurately modelling the construction
410 error throughout the road corridor, which would in-turn require measurement and modelling techniques that introduce
411 their own errors. To account for this, and make the following analysis an 'apples-to-apples' comparison, the authors
412 computed the simulated data's point-to-surface distances relative to the errorless design surfaces. Consequently, the
413 point-to-surface distributions reported below are composite values that include both measurement and construction
414 vertical errors along the length of the corridor. Points attributable to occlusions in the real-world data were manually
415 cropped prior to performing the analysis in order to focus only on the quality of the road surface measurements. This
416 comparison confirms that the simulated data reasonably approximates the real-world aerial photogrammetric data. The
417 real-world vertical errors were within $\sigma_{Day1} = 1.5 px$ and $\sigma_{Day2} = 1.3 px$ respectively, both of which are within the
418 expected ranges discussed above.

419 *Development Platform and Performance Measures*

420 The authors developed the proposed solution using an in-house coding platform named Gygax that allows for
421 processing and visualization of both images and PCD, and incorporates the open-source *Emgu CV* and *Point Cloud*
422 *Library* code libraries. The solution uses a combination of C++ and C# code written and compiled in Microsoft Visual
423 Studio 2015. All experiments utilized a computer with 4.0 GHz Intel i7 processor, 32 GB RAM, a dedicated 1,280-
424 core GPU with 2 GB memory, and Windows 10 64-bit operating system.

425 This study used average precision (p_μ), average recall (r_μ), and the F score (F_1) to measure experimental
 426 performance by comparing the binary *leafVoxel* detection labels to the ground truth. The ground truth was developed
 427 by overlaying the project grid with the as-built data and manually labelling each branch with the observed design layer
 428 name. Equations 6-8 define these metrics for a model with l design surfaces, where true positive (TP_c) is the number
 429 of *leafVoxels* correctly labelled as 'detected' in surface c , false positive (FP_c) is the number incorrectly labelled as
 430 'detected', true negative (TN_c) is the number correctly labelled as 'not detected', and false negative (FN_c) is the number
 431 incorrectly labelled as 'not detected'.

$$432 \quad p_\mu = \frac{\sum_{c=1}^l TP_c}{\sum_{c=1}^l TP_c + FP_c} \quad (8)$$

$$433 \quad r_\mu = \frac{\sum_{c=1}^l TP_c}{\sum_{c=1}^l TP_c + FN_c} \quad (9)$$

$$434 \quad F_1 = \frac{2 \times p_\mu \times r_\mu}{p_\mu + r_\mu} \quad (10)$$

435 p_μ measures the reliability of the positive-detection decision as the portion of *leafVoxels* labelled 'detected' that were
 436 detectable in the as-built scene. r_μ measures the proposed solution's positive-detection effectiveness as the portion of
 437 detectable *leafVoxels* that were labelled as 'detected'. Each of these metrics describes a different component of the
 438 proposed solution's overall effectiveness and optimal performance is achieved when each is maximized. In practice,
 439 this is difficult because there are often trade-offs between precision and recall. F_1 is the harmonic mean of p_μ and r_μ ,
 440 and is used to measure overall performance in the following experiments.

441 *Experiments on Synthetic Data*

442 The first experiments conducted on the synthetic data aimed to determine which branch detection deconfliction rule
 443 produces the best results, and how variations in the voxel height (*searchDistance*) affect those results. The authors
 444 conducted 20 trials for each decision rule tested (60 in total), varying the *searchDistance* between 0.1 and 2.0 cm in
 445 0.1 cm increments. As discussed previously, the wearing course's 40 mm design thickness drove selection of the
 446 *searchDistance* values tested. Figure 9 shows the F1 Score achieved in each of these trials. The maximum point
 447 decision rule resulted in the highest F1 Score in each case. Additionally, the authors concluded that the minimum
 448 distance rule is not reliable for deconflicting branch detection decisions, as it demonstrated no ability to improve
 449 detection performance. Finally, the results in Figure 9 show a preference for the maximum-possible voxel height,

450 using a search distance of 2 cm in this case. Using the optimal decision rule (maximum point) and search distance (2
451 cm), the authors recorded an average F₁ Score of 91.0% with the misclassifications occurring almost exclusively in
452 the upper-two design layers (77.1% and 80.8% F₁ Scores respectively for the Binder and Asphalt Wearing Surface
453 layers). This indicates that the normally-distributed measurement errors cause confusion in the classification decision
454 for thinner layers. The following experiments aimed to correct these errors using the outlier correction step described
455 in the previous section.

456 To examine the influence of the design parameters on the outlier correction step's performance, the authors
457 conducted further trials iterating through *neighbourDistance* values from 1 to 10 while varying the
458 *consensusThreshold* between 0.5 and 1.0 in increments of 0.05. The range of *neighbourDistance* values tested was
459 determined experimentally to ensure an observable peak in F₁ Score. The authors chose to use 0.5 as the minimum
460 *consensusThreshold* assuming that the most likely scenario would involve a neighbourhood with just two surface
461 layers under consideration at a time. Figure 10 shows the results of these experiments. Peak performance was achieved
462 using a neighbour distance of 2 and consensus threshold of 60%, which resulted in an average F₁ Score of 99.7%. The
463 results confirm the utility of the proposed solution under ideal conditions, when the constructed surfaces are clearly
464 visible. The following section examines the proposed solution's performance on the real-world datasets.

465 *Experiments on Real-World Data*

466 The data collection surface on Day 1 was subject to a number of occlusions caused by a parked vehicle, a roller
467 compacter, and multiple safety barriers associated with work related to other schedule tasks. Additionally, the
468 contractor intentionally built three regions of the binder material to a lower-than-designed elevation. These regions
469 corresponded to areas where connecting roads were awaiting their asphalt binder installation, during which time the
470 contractor would correct the intentionally-low regions to provide a smooth connection. The initial real-world-data
471 experiments eliminated these regions and occluded areas, in both datasets, from the analysis in order to isolate the
472 method's performance on regions where the ground truth surfaces are observable in the data. Strategies for dealing
473 with these troubling regions are then discussed in subsequent paragraphs, and are the focus of continuing research.
474 Figure 11 illustrates the location and type of some of the occlusions. The authors implemented the proposed solution
475 on the Day 1 and Day 2 datasets using the optimal parameters from synthetic data testing. Table 2 summarizes the
476 results of these experiments. The overall results show reasonably good performance, but further improvement is

477 needed. The reduced performance compared to the synthetic data is due to localized error regions in the ground-truth
478 Asphalt Binder and Wearing Surface areas (Figure 12 highlights some typical examples). These localized errors could
479 be due to unplanned deviations in surface elevation (i.e. construction errors) or systematic errors in the point cloud
480 data. Regardless of the cause, a progress monitoring approach should be able to recognize which design layers these
481 regions are supposed to belong to. This also holds true for the regions that are occluded or otherwise deviate from the
482 design layer elevations. The following experiments examined how such regions could be accounted for using the
483 proposed outlier correction method.

484 *Adjusting for error regions and occlusions*

485 The authors experimented with increasing the outlier correction method's neighbour distance parameter to compensate
486 for the error and occluded regions. The combinations of outlier correction parameters producing an F1 Score above
487 95% in the synthetic data (as shown in Figure 10), were implemented on the Day 1 and Day 2 datasets for search
488 distances greater than 2. The best overall performance resulted from a search distance of 7 with a consensus threshold
489 of 55%. Table 3 shows the results for this combination, both with and without the occlusions accounted for in the
490 performance calculations. Figure 13 illustrates the results.

491 *Processing Time*

492 The proposed solution took approximately 11 seconds to generate the incremental surface detection decisions for the
493 200 m section of road analysed above, using a 5.5-million-point input cloud. Note that this does not include the
494 *BrickTree* build time (~40 seconds), as this step only needs to be completed at project initialization or when the
495 design is changed. The total cycle time (including *BrickTree* initialization, data collection, and post-processing) was
496 approximately 3 hours, 35 minutes on Day 1, and 2 hours, 29 minutes on Day 2. The increased efficiency on Day 2
497 was due mostly to better familiarity with the post-processing software and field equipment. The authors tested the
498 proposed solution on additional synthetic data in order to evaluate how changes in the size of the as-built and as-
499 planned data affect its processing time, the results of which are summarized in Figure 14. Note that *BrickTree*
500 initialization time (not shown) also increased linearly with projects size, with the 10-Lane, 500 m long road
501 requiring 2 minutes, 8 seconds to initialize the data structure.

502 **SUMMARY AND CONCLUSION**

503 Effectively monitoring construction progress can enable timely and effective project control decisions aimed at
504 keeping a project on track. Unfortunately, current transportation project monitoring practice is manual, error-prone,
505 inefficient, and ultimately contributes to the annual \$114.3B in global cost overrun. Automating this process can
506 improve overall project performance, and detecting design surfaces in digital as-built data is a key requirement for
507 realizing such an automated approach. A number of recent studies proposed methods for conducting this detection on
508 building structural components. Unfortunately, these methods are not well-suited for larger and more complex layered
509 road design surfaces. The only related transportation-focused study required labelled as-built data, entirely skipping
510 the automated surface detection process.

511 This paper marks the first study specifically aimed at automatically detecting layered road design surfaces in
512 unlabelled as-built point cloud data. The authors proposed and evaluated a novel model-guided space partitioning
513 hierarchical tree structure, termed the *BrickTree*, for accomplishing this task. A simulated aerial photogrammetry point
514 cloud was generated and used to test the performance of the proposed solution under ideal conditions not subject to
515 occlusions or construction errors beyond acceptable levels. These initial experiments culminated in an average F_1
516 Score of 99.7%, while providing perspective on the best decision rule and outlier correction parameters to use for
517 further testing. The solution was then tested on two real-world datasets, ultimately resulting in an average F_1 Score of
518 95.2%.

519 It is important to note that the proposed solution does have limitations. This study focuses on progress
520 monitoring, and as such does not aim to identify or classify construction errors. It requires the constructed surfaces to
521 be mostly within established quality limits. The authors acknowledge the natural link between progress and quality,
522 and note that large-scale quality measurement is inherent in the proposed *leafVoxel* populating and *rasterBranch*
523 classifying steps, since as-built points lying outside the model-hugging voxel regions will not be considered. In this
524 manner, road layers that are built with errors exceeding the acceptable thresholds in large areas will not be detected
525 by the method, and this lack of detection could be used to highlight potential construction errors for further
526 investigation by the project management team. Another limitation is that the proposed solution requires up-to-date
527 design information that reflects field-level construction procedures. For example, consider the regions discussed
528 earlier that were intentionally left lower than the design elevations to account for connection with intersecting roads

529 that had yet to be paved. Modelling this level of detail is possible, but requires close coordination between the design
530 and construction teams. This study's main contributions are:

531 (1) The *BrickTree* space partitioning hierarchical data structure, which has numerous advantages. First, the
532 sparse nature of the hierarchical tree structure, utilizing a fixed tree depth where voxels are only generated at known
533 surface geometry locations, reduces the complexity of tree traversal and search operations while simultaneously
534 reducing memory usage compared to equivalent non-sparse voxelization approaches (Laine & Karras, 2011). Second,
535 the uniform fit of each voxel to its parent surface geometry, thanks to the model-guided construction topology, ensures
536 a consistent detection framework across the entire surface. This improves the consistency of point-to-surface
537 attribution and eliminates the need for an as-planned traversal of the tree structure (because it's already known which
538 design surfaces belong to which voxels). Third, the ability to adjust voxel shape allows for point searches in tighter
539 vertical tolerances while maintaining broader horizontal ranges that enable consideration of sparser as-built point
540 clouds. Finally, the layered vertical branch orientation enables logical reasoning about the presence of a design surface
541 based on the status of the other voxels within the branch and its' neighbours.

542 (2) Empirical evidence supporting the optimal *searchDistance* parameter and branch classification decision
543 rule to use when implementing the proposed solution. The authors noted a preference in the results for use of the
544 largest *searchDistance* possible based on the thinnest design surface layer. However, further research is needed to
545 identify how well these conclusions generalize when applied to different as-built and as-planned datasets.

546 (3) A workflow for generating synthetic as-built road construction point cloud data that is qualitatively similar
547 to real-world aerial photogrammetric data. The developed application allows the user to specify parameters controlling
548 the synthetic point cloud's noise/error and density distribution profiles while also modelling construction errors. This
549 can be used to test the viability of a proposed method and/or its sensitivity to fluctuations in data quality. It can also
550 be used to determine flight parameters and equipment requirements (using ground sample distance) for achieving the
551 level of data quality needed for a given method.

552 While the results reported here are promising, there remains room for improvement. The authors divide the
553 remaining errors, as visible in the Day 1 results at the top of Figure 13, into two different categories. The first is
554 confusion between the Asphalt Binder and Wearing Surface layers. These surfaces are separated by the thinnest
555 margin, i.e. the thickness of the Wearing Surface layer, and are thus more susceptible to misclassification errors due

556 to noise in the PCD and construction errors near the acceptable limits. While more accurate and precise as-built data
557 collection methods (e.g. laser scanning) could ameliorate this situation, UAS photogrammetry offers several
558 advantages that warrant further investigation into methods for dealing with such misclassifications in the aerial PCD.
559 The biggest advantages of UAS photogrammetry include the lower cost and ability to collect the data remotely. The
560 latter advantage improves site safety by not exposing further personnel to the hazards inherent on a construction site,
561 and ensures the data collection will not impede ongoing work.

562 The second error category is confusion at the boundary between adjacent layers. Note that this issue did not
563 emerge while testing the synthetic data using a smaller neighbour distance during outlier correction. However, as the
564 authors increased the neighbour distance to accommodate larger error regions than modelled in the synthetic data, the
565 boundaries between adjacent layers became less well defined. For example, consider Figure 13 which shows that the
566 staggered boundary between the Base Course and Asphalt Binder layers in the Day 1 data is lost using the larger
567 neighbour distance. A more sophisticated outlier correction method could potentially be used to account for initial
568 classification errors while maintaining the boundary conditions in the as-built data. Both of these error categories will
569 be further examined in the authors' future research.

570 **ACKNOWLEDGMENTS**

571 This research is made possible through funding from the United States Air Force and the Cambridge Commonwealth
572 and International Trust. The authors express gratitude to the Trimble Corporation for their support in lending
573 equipment and expertise to the data collection operation. The views expressed in this paper are those of the authors
574 and do not necessarily reflect the official policy or position of the Air Force, the Department of Defense or the U.S.
575 Government. All supporting data is included in this paper and the provided references.

576 **REFERENCES**

- 577 Bosché, F. (2010). "Automated recognition of 3D CAD model objects in laser scans and calculation of as-built
578 dimensions for dimensional compliance control in construction." *Advanced Engineering Informatics*, 24 (1), 107–
579 118.
- 580 Braun, A., Tuttas, S., Borrmann, A. & Stilla, U. (2015). "Automated progress monitoring based on photogrammetric
581 point clouds and precedence relationship graphs." *Proc., 32nd International Symposium on Automation and
582 Robotics in Construction and Mining*, 274–280.
- 583 Cantarelli, C.C., Flyvbjerg, B., Molin, E.J.E. & van Wee, B. (2010). "Cost Overruns in Large-Scale Transportation

584 Infrastructure Projects: Explanations and Their Theoretical Embeddedness." *European Journal of Transport and*
585 *Infrastructure Research*, 10 (1), 5–18.

586 Chan, D.W. & Kumaraswamy, M.M. (1997). "A comparative study of causes of time overruns in Hong Kong
587 construction projects." *International Journal of Project Management*, 15 (1), 55–63.

588 Cheng, M.Y. & Chen, J.C. (2002). "Integrating barcode and GIS for monitoring construction progress." *Automation*
589 *in Construction*, 11 (1), 23–33.

590 Dai, F., Feng, Y. & Hough, R. (2014). "Photogrammetric error sources and impacts on modeling and surveying in
591 construction engineering applications." *Visualization in Engineering*, 2 (1), 1–14.

592 Del Pico, W.J. (2013). *Project Control: Integrating Cost and Schedule in Construction*. Hoboken, NJ: John Wiley &
593 Sons.

594 DJI (2017). "DJI Site Survey Solution". <http://store.dji.com/product/site-survey-solutions-single-perpetual> (5
595 March 2017).

596 Dobbs, R., Pohl, H., Lin, D.-Y., Mischke, J., Garemo, N., Hexter, J., Matzinger, S., Palter, R. & Nanavatty, R.
597 (2013). *Infrastructure productivity: how to save \$1 trillion a year*. McKinsey Global Institute.

598 Faro (2017). "Laser Scanner Tech Sheets". <http://www.faro.com/en-gb/resources/tech-sheets/> (13 November
599 2017).

600 Golparvar-Fard, M., Peña-Mora, F. & Savarese, S. (2015). "Automated Progress Monitoring Using Unordered Daily
601 Construction Photographs and IFC-Based Building Information Models." *Journal of Computing in Civil*
602 *Engineering*, 10.1061/(ASCE)CP.1943-5487.0000205, 04014025.

603 GOM Inspect (2017). "GOM Inspect Software". <http://www.gom.com/3d-software/gom-inspect.html> (29 April
604 2017).

605 González, P., González, V., Molenaar, K. & Orozco, F. (2014). "Analysis of Causes of Delay and Time Performance
606 in Construction Projects." *Journal of Construction Engineering and Management*, 10.1061/(ASCE)CO.1943-
607 7862.0000721, 04013027.

608 Han, K.K., Degol, J. & Golparvar-Fard, M. (2018). "Geometry- and Appearance-Based Reasoning of Construction
609 Progress Monitoring." *Journal of Construction Engineering and Management*, 10.1061/(ASCE) CO.1943-
610 7862.0001428, 04017110.

611 Highways Agency, UK (2014). "Manual of Contract Documents for Highway Works, Volume 1, Series 0700",
612 London, UK.

613 Kalasapudi, V.S., Tang, P. & Turkan, Y. (2017). "Computationally efficient change analysis of piece-wise
614 cylindrical building elements for proactive project control." *Automation in Construction*, 81, 300–312.

615 Kim, C., Son, H. & Kim, C. (2013). "Automated construction progress measurement using a 4D building
616 information model and 3D data." *Automation in Construction*, 31, 75–82.

617 Kivimaki, T. & Heikkila, R. (2015). "Infra BIM based Real-time Quality Control of Infrastructure Construction
618 Projects." *Proc., 32nd International Symposium on Automation and Robotics in Construction and Mining*, 877–882.

619 Laine, S. & Karras, T. (2011). "Efficient sparse voxel octrees." *IEEE Transactions on Visualization and Computer
620 Graphics*, 17 (8), 1048–1059.

621 Leica (2017). "Leica Pegasus:Two Mobile Reality Capture." <http://www.leica-geosystems.co.uk> (13 November
622 2017).

623 Memon, A.H., Abdul Rahman, I. & Abdul Aziz, A.A. (2012). "The cause factors of large project's cost overrun: a
624 survey in the southern part of Peninsular Malaysia." *International Journal of Real Estate Studies (INTREST)*, 7 (2),
625 1–15.

626 Navon, R. & Sacks, R. (2007). Assessing research issues in Automated Project Performance Control (APPC).
627 *Automation in Construction*, 16 (4), 474–484.

628 NCHRP (2013). "Synthesis 446: Use of Geospatial Data, Tools, Technologies, and Information in Department of
629 Transportation Projects." Washington, D.C.

630 Olsen, M.J., Roe, G. V., Glennie, C., Persi, F., Reedy, M., Hurwitz, D., Williams, K., Tuss, H., Squellati, A. &
631 Knodler, M. (2013). "NCHRP Report 748: Guidelines for the Use of Mobile LIDAR in Transportation
632 Applications." Washington, D.C.

633 Pauly, K. (2016). "White Paper: Trimble UX5 - Increasing Your Productivity." Trimble Corporation, Belgium.

634 Pix4d (2017). "Pix4D Drone Mapping." <https://pix4d.com/> (7 July 2015).

635 Razavi, S.N. & Haas, C.T. (2010). "Multisensor data fusion for on-site materials tracking in construction."
636 *Automation in Construction*, 19 (8), 1037–1046.

637 Razavi, S.N. & Moselhi, O. (2012). "GPS-less indoor construction location sensing." *Automation in Construction*,
638 28, 128–136.

639 Salling, K.B. & Leleur, S. (2015). "Accounting for the inaccuracies in demand forecasts and construction cost
640 estimations in transport project evaluation." *Transport Policy*, 38, 8–18.

641 Schneider, C. (2013). "Techbrief: 3D, 4D, and 5D Engineered Models for Construction, an Executive Summary."
642 US Department of Transportation, Federal Highway Administration.

643 Shahi, A., Aryan, A., West, J.S., Haas, C.T. & Haas, R.C.G. (2012). "Deterioration of UWB positioning during
644 construction." *Automation in Construction*, 24, 72–80.

645 Siemiatycki, M. (2009). "Academics and Auditors: Comparing Perspectives on Transportation Project Cost
646 Overruns." *Journal of Planning Education and Research*, 29 (2), 142–156.

647 Slocum, R. & Parrish, C. (2017). "Simulated Imagery Rendering Workflow for UAS-Based Photogrammetric 3D
648 Reconstruction Accuracy Assessments." *Remote Sensing*, 9 (4), 396.

649 The Survey Association (2015). "Client Guide to Aerial LiDAR Surveys." Newark-on-Trent, UK.

650 The Survey Association (2016). "Client Guide to Small Unmanned Aircraft Surveys." Newark-on-Trent, UK.

651 TransMagic (2017). "MagicCheck – CAD Model Comparison." <[https://transmagic.com/magiccheck-cad-model-
652 comparison/](https://transmagic.com/magiccheck-cad-model-comparison/)> (29 April 2017).

653 Trimble (2015). "Trimble TX8 Laser Scanner Datasheet." [Online]. <[http://www.trimble.com/3d-laser-
654 scanning/index.aspx](http://www.trimble.com/3d-laser-scanning/index.aspx)> (16 July 2015).

655 Trimble (2017). "Trimble Business Center Aerial Photogrammetry Module." [Online]. <[http://uas.trimble.com/tbc-
656 am](http://uas.trimble.com/tbc-am)> (29 April 2017).

657 Turkan, Y., Bosché, F., Haas, C.T. & Haas, R. (2014). "Tracking of secondary and temporary objects in structural
658 concrete work." *Construction Innovation*, 14 (2), 145–167.

659 Tuttas, S., Braun, A., Borrmann, A. & Stilla, U. (2014). "Comparison of Photogrammetric Point Clouds with BIM
660 Building Elements for Construction Progress Monitoring." *The International Archives of Photogrammetry, Remote
661 Sensing and Spatial Information Sciences*, XL-3, 341–345.

662 U.S. Army Corps of Engineers (2013). *Engineering and Construction Bulletin 2013-18*. (ECB 2013-18).
663 Washington, D.C.

664 Watt, A. (2000). *3D Computer Graphics*. Third. Essex, UK: Pearson Education.

665 Zhang, C. & Ardit, D. (2013). "Automated progress control using laser scanning technology." *Automation in
666 Construction*, 36, 108–116.

667

668

669

670


671

672

673

TABLES

674

Table 1. Summary of accurate and dense spatial data collection technologies. 


Technology	Cost	Accuracy	Density	Strengths	Limitations
Manned aerial LiDAR ^{a,c,d}	\$\$	cm	10s - 100s	Lighting invariance, large coverage area, non-intrusive	Cost, mobilization time
Unmanned aerial LiDAR ^{a,c,d}	\$	mm - cm	100s	Lighting invariance, non-intrusive	Battery life
Unmanned aerial photogrammetry ^{a,b,e}	< \$	cm	100s	Cost, non-intrusive	Battery life, daylight collection only
Mobile LiDAR ^{a,c,f}	\$ - \$\$	cm	100s - 1,000s	Lighting invariance, coverage area	Cost, intrusive, ground view more sensitive to occlusions
Terrestrial LiDAR ^{a,g,h}	\$	mm	100s - 10,000s	Accuracy, density, Lighting invariance	Coverage, intrusive, ground view more sensitive to occlusions

675 Note: Cost ranges are up-front, \$ indicates costs in the \$10,000s and \$\$ indicates costs in the \$100,000s. Density
 676 ranges are in pts/m². The intrusive/non-intrusive assessment considers whether or not the technology requires on-site
 677 collection in the proximity of ongoing work, which has implications for safety and work productivity.

678 Sources: ^aNCHRP (2013), ^bThe Survey Association (2016), ^cOlsen et al. (2013), ^dThe Survey Association (2015b),
 679 ^eDJI (2017), ^fLeica (2017), ^gTrimble (2015), ^hFaro (2017)

680

681


Table 2. Summary of initial results on the Day 1 and Day 2 datasets 

Dataset	p_{μ}	r_{μ}	FI
Day 1	83.6%	82.8%	83.2%
Day 2	84.7%	84.7%	84.7%
Overall	84.2%	83.8%	84.0%

682

683

684

Table 3. Summary of final results on the Day 1 and Day 2 datasets without (and *with*) occlusions considered 

Dataset	p_{μ}	r_{μ}	FI
Day 1	90.2% (83.6%)	89.8% (83.2%)	90.0% (83.4%)
Day 2	99.9% (99.9%)	99.9% (99.8%)	99.9% (99.9%)
Overall	95.3% (91.7%)	95.1% (91.5%)	95.2% (91.6%)

685

Figure Captions

Figure 1. Longitudinal road design section showing 2D projection of 10 cm voxel grid

Figure 2. Proposed Surface Detection Process

Figure 3. The BrickTree Structure

Figure 4. Neighbourhood construction example for *neighbourDistance* = 1

Figure 5. Simulated In-Progress Road Construction Site

Figure 6. UAS Flight Details

Figure 7. As built point clouds and ground truth conditions

Figure 8. Comparison of the real world and simulated datasets

Figure 9. Performance for the tested search distances and decision rules

Figure 10. Performance for each combination of outlier correction parameters tested

Figure 11. Occlusions and deviations in the Day 1 data (ground truth surfaces shown)

Figure 12. Localized error regions (circled) in the Day 1 dataset

Figure 13. Day 1 and Day 2 surface detection results after increasing the *searchDistance* parameter

Figure 14. Complexity of the proposed solution


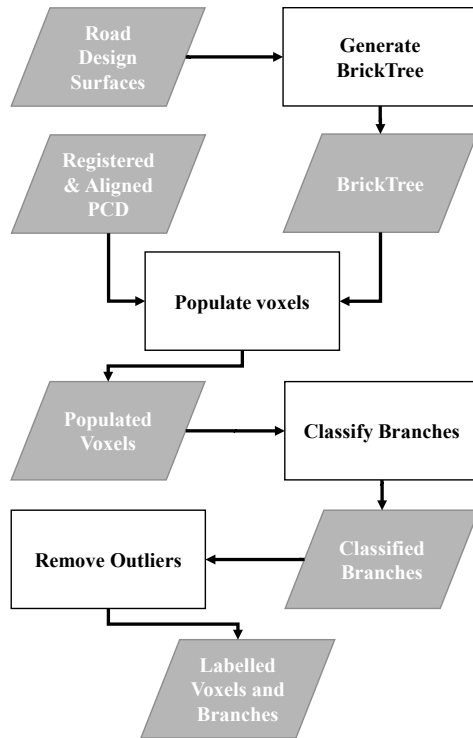


Figure 1. Longitudinal road design section showing 2D projection of 10 cm voxel grid





Legend:



 Process Inputs/Outputs  Process

Figure 2. Proposed Surface Detection Process



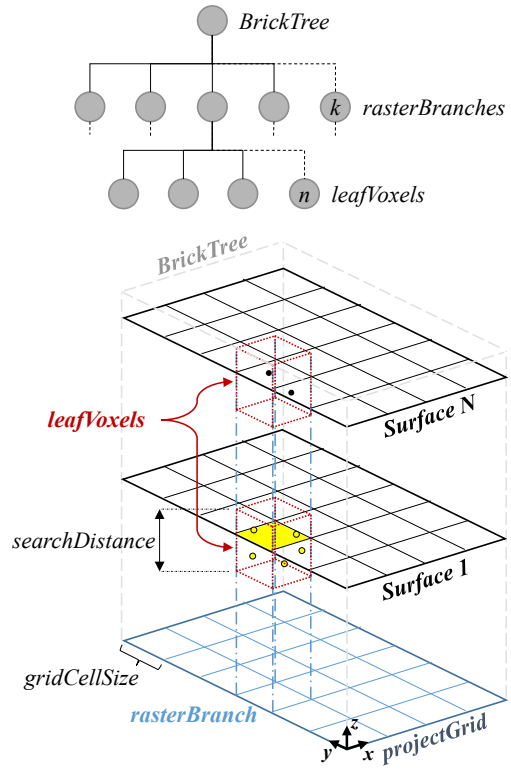


Figure 3. The BrickTree Structure






Figure 4. Neighbourhood construction example for neighbourDistance = 1






Figure 5. Simulated In-Progress Road Construction Site





Flight Information	Day 1	Day 2
Avg. Flying Height AGL (m):	70.3	70.7
No. of Images:	224	251
Avg. Photo Scale:	1 : 4799	1 : 4883
Avg. GSD (cm):	2.29	2.33
Pre-Flight Setup (min)	45	27
Flight Duration (min)	11	12
Post Processing (min)	158	110

Figure 6. UAS Flight Details






Figure 7. As built point clouds and ground truth conditions






Figure 8. Comparison of the real world and simulated datasets



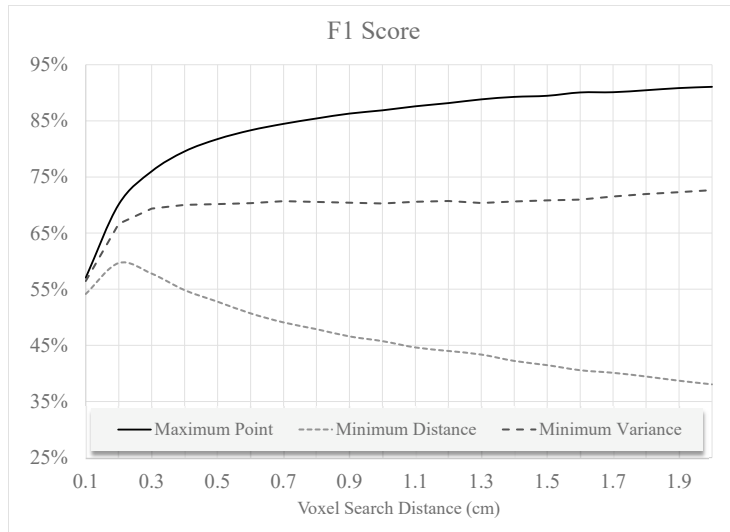


Figure 9. Performance for the tested search distances and decision rules



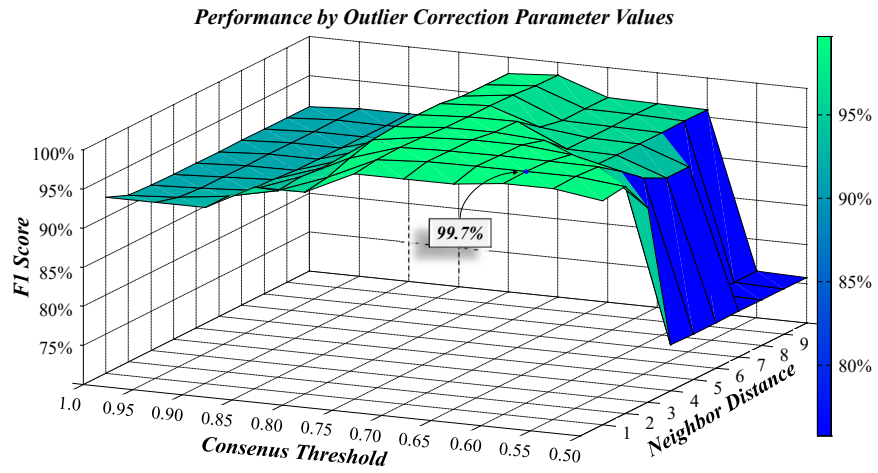



Figure 10. Performance for each combination of outlier correction parameters tested





Ground-truth layers:

Base Binder

Figure 11. Occlusions and deviations in the Day 1 data (ground truth surfaces shown)








Figure 12. Localized error regions (circled) in the Day 1 dataset



Day 1:



Day 2:



Legend:

Sub Base Base Binder Wearing Not Detected

Figure 13. Day 1 and Day 2 surface detection results after increasing the searchDistance parameter



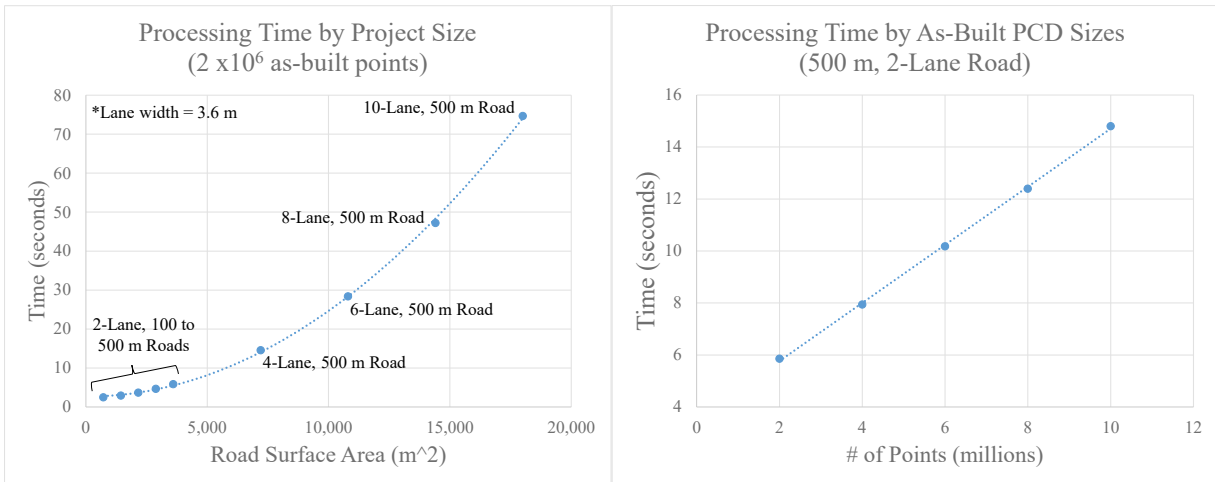


Figure 14. Complexity of the proposed solution

