

Augmented Reality for Subsurface Utility Engineering, Revisited

Hansen, Lasse Hedegaard; Fleck, Philipp; Stanner, Marco; Schmalstieg, Dieter; Arth, Clemens

Published in:

I E E E Transactions on Visualization and Computer Graphics

DOI (link to publication from Publisher):

[10.1109/TVCG.2021.3106479](https://doi.org/10.1109/TVCG.2021.3106479)

Publication date:

2021

Document Version

Accepted author manuscript, peer reviewed version

[Link to publication from Aalborg University](#)

Citation for published version (APA):

Hansen, L. H., Fleck, P., Stanner, M., Schmalstieg, D., & Arth, C. (2021). Augmented Reality for Subsurface Utility Engineering, Revisited. *I E E E Transactions on Visualization and Computer Graphics*, 27(11), 4119-4128. Article 21399443. <https://doi.org/10.1109/TVCG.2021.3106479>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal -

Take down policy

If you believe that this document breaches copyright please contact us at vbn@aub.aau.dk providing details, and we will remove access to the work immediately and investigate your claim.

Augmented Reality for Subsurface Utility Engineering, Revisited

Lasse H. Hansen, Philipp Fleck, Marco Stranner, Dieter Schmalstieg and Clemens Arth



Fig. 1. Virtual utility marking and daylighting in comparison to physical counterparts: (left) physical spray markings on the street surface, (mid left) virtual utility marking of similar information in AR, (mid right) worker examining a virtual excavation in Aalborg (Denmark), (right) AR view with embedded 3D reconstruction of said excavation.

Abstract— Civil engineering is a primary domain for new augmented reality technologies. In this work, the area of subsurface utility engineering is revisited, and new methods tackling well-known, yet unsolved problems are presented. We describe our solution to the outdoor localization problem, which is deemed one of the most critical issues in outdoor augmented reality, proposing a novel, lightweight hardware platform to generate highly accurate position and orientation estimates in a global context. Furthermore, we present new approaches to drastically improve realism of outdoor data visualizations. First, a novel method to replace physical spray markings by indistinguishable virtual counterparts is described. Second, the visualization of 3D reconstructions of real excavations is presented, fusing seamlessly with the view onto the real environment. We demonstrate the power of these new methods on a set of different outdoor scenarios.

Index Terms—Augmented Reality, Infrastructure, Computer Graphics, Localization

1 INTRODUCTION

In civil engineering, knowing the exact type and location of subsurface utilities, both old and new, is extremely important. This knowledge is required to efficiently plan construction activities, to avoid accidental damage during excavations, to communicate with other stakeholders in the construction process, and to document the assets for long-time archival purposes.

Excavation damage to underground infrastructure in particular is causing huge financial losses. Reports speak of GBP 270 million in the UK [6, 22] and USD 30 billion in the US [9] annually. Of course, *subsurface utility engineering* (SUE) does everything to prevent damages by following established procedures. Current best practices include:

1. Applying spray markings to annotate utility locations on the ground surface, based on information extracted from paper maps or geographic information services (GIS) running on handhelds
2. Using ground penetrating radar to scan for subsurface assets
3. So-called *daylighting*, i.e., slow and careful excavation work until the underground assets are exposed to daylight

Obviously, all these best practices incur additional cost, and measures such as radar scans or daylighting are only pursued if deemed necessary. Consequently, much reliance is placed on SUE documentation. According to a recent survey [1] questioning 477 excavation contractors, poor or outdated documentation of subsurface utilities is one of the

main causes of damages. Ordered by significance, the damages resulted from (1) lack of depth measurements in the documentation (i.e., no information about the level below surface where a utility is located), (2) street-level markings painted too far from the utilities (i.e., inaccurate surveying data created by rushed or unskilled surveyors), and (3) missing markings (which may have become eroded by weathering, or been removed along with the top-layer surface during construction).

Clearly, human error contributes most to damages and other problems in SUE. This is a clear indicator to consider *augmented reality* (AR) to improve upon traditional paper-based inspection, planning or maintenance procedures. Indeed, using AR for infrastructure visualization in SUE has already been proposed over 20 years ago by Spohrer [34]. At this time, it seemed plausible that emerging AR devices would soon deliver mobile 3D visualization of GIS data. Yet, to date, development and commercialization of such efforts has been modestly successful at best, as we will point out in the related work.

There are several technical, but also social reasons for this curb. In general, civil engineers operating under time pressure and with a tight budget tend to be very conservative and are reluctant to trade existing tools for technology considered immature. When it comes to AR, contemporary devices, be they handhelds (e.g., using ARKit/ARCore) or headsets (e.g., HoloLens), lack the ability for city-scale localization with the surveying-level accuracy expected by civil engineers. Even if accurate localization is solved, measurements in GIS databases are still ambiguous, since they are expressed *relative to the surface at site without recording the absolute height* of the actual surface.

Besides, while methods for creating high-quality 3D models exist (e.g., laser scanning after daylighting), there is no straight forward way to exploit them for on-site underground infrastructure visualization, other than through (non-AR) GIS viewers on handhelds.

In this paper, we revisit previous approaches to using AR for SUE, concentrating on lifting the aforementioned limitations. We feel that contemporary technologies, which were not feasible 10 or 20 years ago, warrant such a re-evaluation. We present the following contributions:

Wide-area localization. We describe a custom highly accurate sensor module, which enhances a tablet computer with wide-area lo-

- Lasse Hansen is with Aalborg University, Denmark. E-mail: lhha@build.aau.dk.
- Philipp Fleck, Marco Stranner and Dieter Schmalstieg are with Graz University of Technology, Austria. E-mail: {philipp.fleck, marco.stranner, dieter.schmalstieg}@icg.tugraz.at.
- Clemens Arth is with AR4 GmbH, Graz, Austria. E-mail: clemens@ar4.io.

Manuscript received xx xxx. 201x; accepted xx xxx. 201x. Date of Publication xx xxx. 201x; date of current version xx xxx. 201x. For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org.
Digital Object Identifier: xx.xxx/TVCG.201x.xxxxxxx



Fig. 2. Examples for commercial data visualization systems: (left) screenshot of vGIS, (right) screenshot of Trimble SiteVision. The usual visualization methods sufficiently convey basic location information for subsurface utilities; however, they fail to deliver on matching known and well-established marking schemes or plausible visual realism.

calization. We describe the hardware, which is currently in its second generation, and give details on a user-friendly procedure for accurate sensor-to-camera calibration. By fusing wide-area localization with SLAM, we can recover the surface elevation, which is necessary when interpreting measurements from GIS databases that lack explicit height.

Virtual utility marking. Based on accurate localization and an instant 3D scan of the user’s immediate surroundings (using a tablet with LIDAR sensor), we give the user the ability to place annotations in the form of virtual spray markings on the street surface. Not only does this approach resemble established work practices and is easy to understand, the visual appearance of simulated spray paint also ensures high contrast over arbitrary backgrounds. The latter is important when looking at a video-see through display with limited dynamic range in an outdoor setting (Fig. 1 left, mid left).

Virtual utility daylighting. In our experience, GIS data quality is often insufficient to synthesize a useful 3D model for AR visualization. However, it is becoming common to acquire 3D reconstructions (usually colored point clouds) of excavations created on site, before the excavation hole is covered up, and store them in a GIS. We show how these datasets can be prepared for on-site visualization with an AR system, for instance, as a prerequisite for planning or on-surface marking (Fig. 1 mid right, right).

The resulting AR application supports SUE at a previously unseen level of quality, as we will demonstrate with several selected use cases.

2 RELATED WORK

Previous work on AR for SUE has been conducted by different interest groups. The AR research community has primarily focused on technical problems, such as outdoor localization and X-ray visualization techniques. Meanwhile, the civil engineering application investigated system integration and data management. Yet another relevant category are recent AR-like extensions of commercial surveying systems. This section gives an overview of these developments.

2.1 Outdoor localization

While early outdoor AR used bulky GPS sensors [12,25], the *Going Out* system was the first to introduce visual-inertial fusion for localization in a mobile form factor [27]. Around the same time, PTAM [19] showed how simultaneous localization and mapping (SLAM) can have practical use, paving the way for a huge wave of visual tracking for AR. Some of these works investigated outdoor localization [20, 41]. Wide-area localization may either leverage large-scale 3D reconstructions [3,5] or models synthesized from GIS-like maps of above-surface structures [4]. However, the issue remains that all these works remain dependent on maps or other data sources, which may be scarce or non-existent.

Meanwhile, the performance and versatility of GPS has increased silently, but significantly, since the 2010s. Most users frustrated by the underwhelming accuracy of GPS in their smartphones do not realize that centimeter-level localization via GPS not only exists, but is actually a reasonably inexpensive commodity. The main factor precluding the deployment of enhanced GPS in smartphones is that a slightly larger GPS chip and a significantly larger antenna cannot be comfortably

fitted into a smartphone case. We will pick up on this observation in our hybrid tracking system, which combines RTK-GPS, LIDAR and a camera in a low cost platform.

2.2 Infrastructure visualization

Soon after the first outdoor AR systems for navigation [12] and urban modeling [25] were introduced in the late 1990s, AR was adopted for SUE, first, using a 2D projection of utilities onto the street-level surface [28]. In the late 2000s, handhelds became available, and more complex SUE scenarios were addressed with X-ray visualization techniques [29–31]. This line of work considered more advanced forms of location sensing and combined them with 3D infrastructure visualizations generated from GIS data using a cut-away technique resembling a virtual excavation. We revisit this idea in our daylighting application.

Later work contributed enhanced X-Ray visualisation techniques, ghosting and shadow projections [44, 45]. The studies reported in these papers demonstrated that synthetic depth cues are essential for observers of SUE visualization techniques and that plain overlays alone are insufficient for understanding spatial relationships. A recent approach [7] for using AR in road-side maintenance has a focus on aligning worker motion and laser-range data.

Probably closest related to our work on virtual markings is the one of Cote and Mercier [8], who propose the use of elevation data for proper on-surface marking of subsurface utilities. Their approach uses an offline 3D reconstruction of a predetermined site, together with a manual registration procedure in order to recover the elevation changes of the street surface. Obviously, their approach has two significant restrictions: First, without a prior 3D reconstruction, their method cannot be used at unseen sites. Second, the current appearance of the surface at the given location cannot be incorporated. Our approach overcomes both restrictions by incorporating a live reconstruction obtained by SLAM. Moreover, our method not only considers surface geometry, but also surface *appearance* on site, ensuring the proper visual fusion of virtual and real parts of the scene.

2.3 Civil engineering applications

In the 2010s, AR applications have caught the interest of the civil engineering community, while the AR community seemed to loose interest in SUE. This new wave of implementations focused on general system design and data management [13, 21, 33, 36–38, 43]. Some more recent works in this continuity also investigated perceptual issues using contemporary AR technology. A recent study evaluated the vertical depth judgement performance on different X-Ray visualisation techniques [10], revealing that users perform better in estimating depth of pipes when using a cut-away technique compared to a simple overlay and edge-based ghosting technique. A similar survey on virtual environments confirmed that cut-away performed best for understanding the spatial placement of underground infrastructure [40]. A study based on informal interviews concluded that visualizing 3D reconstructions of previously captured utility excavation holes potentially benefits utility owners’ planning activities [17].



Fig. 3. Prototypical handheld visualization setup: We use a 12.9" Apple iPad Pro (4th gen.) equipped with a built-in rear-view LIDAR sensor. The center box contains the sensor platform used for accurate outdoor registration. The antenna is separately mounted on an adjustable bracket.

2.4 Commercial approaches

Several companies involved in construction and GIS data management developed commercial AR applications for SUE. *AugView*¹ were the first to deploy a commercial AR solution explicitly targeting underground infrastructure. More recently, *vGIS*² (Fig. 2, left) and *AVUS*³ started selling systems for "reality captures", *i.e.*, 3D reconstructions.

Common to these solutions is that they provide only software for bridging GIS with a visualization engine, while relying on third-party surveying equipment, such as laser theodolites, for localization. The only exception to this is *Trimble SiteVision*, an all-in-one solution for outdoor GIS data visualization (Fig. 2, right). All these commercial solutions follow a similar workflow, consisting of (i) GIS data import, (ii) localization using free-standing hardware (theodolite on a tripod) or other bulky hardware, (iii) visualization of simple 3D models using transparent overlays. With these capabilities, current commercial solutions improve upon traditional surveying or inspection work by presenting GIS data in a situated context, but without the interactivity or mobility typically associated with AR. Features such as "virtual spray paint" are not feasible with current commercial approaches.

3 PLATFORM

Both head-worn and handheld form factors are suitable for outdoor AR. We favored a handheld (tablet computer) form factor for several reasons. First, premium tablet computers can now compete with notebook computers in terms of performance, while AR headsets remain more constrained with respect to computational resources. Second, ruggedized tablet computers are widely used by civil engineers as mobile appliances, while headsets are considered too brittle. Third, compared to optical see-through headsets, presenting video-see through AR on a tablet computers can more easily achieve sufficient contrast in bright outdoor conditions. Fourth, tablet computers can be spontaneously shared by multiple users by jointly looking at the screen. Therefore, the relative benefits of AR headsets (primarily the hands-free operation) are diminished compared to handheld devices in the foreseeable future.

Consequently, we chose a tablet computer (Apple iPad Pro, 4th generation, Fig. 3) as our AR platform. Because of its built-in LIDAR sensor, it currently provides the best 3D scanning and SLAM performance for outdoor use. Our software framework runs as a Unity application and interfaces to the SLAM system through Unity AR-Foundation: The framework can also run on other platforms, such as Android/ARCore or Microsoft HoloLens/MRTK. However, all results reported here were obtained with our preferred platform, the iPad.

¹<https://www.augview.net/>

²<https://www.vgis.io/>

³<https://www.avus.tech/>




Fig. 4. Sensor board in transparent 3D-printed case: An optional GSM module can be mounted for fully autonomous operation.

3.1 Outdoor location sensing

SLAM systems using visual-inertial sensing provide 6DOF poses with high precision, but only relative to a starting point or, in vendor jargon, a *world anchor*. To present world-referenced content for SUE, it is necessary to place the world anchor in a global coordinate system. Establishing this global coordinate system is most easily done using GPS. However, consumer-grade GPS, as found in smartphones, is only accurate to tens of meters and not considered useful in civil engineering.

High quality GPS is available in surveying equipment, such as theodolites and total stations, but at a high cost and poor mobility. The actual sensor components are relatively inexpensive, but not ready for plug and play on mobile computers. Moreover, certain components, such as magnetometers and antennae, are susceptible to electromagnetic interference and must be placed at a minimum distance from other electronic components.

For these reasons, we designed a self-contained sensor box in order to solve the outdoor localization problem. This platform is an evolution of the work of Stranner *et al.* [35] on integrating all required components on a single board with a small form factor. The commercially available⁴ device version (Fig. 4) contains exactly the same electrical components with the same accuracy and precision. Its main components include (i) a differential GPS-RTK sensor, (ii) a highly accurate smart IMU with integrated sensor fusion, and (iii) a performance pressure sensor with altimetry. The board is powered over USB with an integrated charging circuit, which optionally enables wireless operation on battery. The board is equipped with a low-power, dual-core 32-bit SoC microcontroller, which controls sensors over I2C and manages communication to external devices over Wifi or Bluetooth using the MQTT protocol [18]. A multi-color LED array is used to signal the current status of the sensor platform.

During regular operation, the host is put into Wifi hotspot mode to enable direct communication with the sensor box. While using an external Wifi network is also possible, the Wifi hotspot is simpler and alleviates additional latency in sensor data transport incurred from Wifi contention with other devices on the same network. Moreover, it establishes a tight connection between the host and the sensor box, allowing the latter to directly leverage the host's Internet up-link to retrieve differential GPS correction data from an NTRIP server. The sensor box streams absolute position and rotation measurements (*i.e.*, GPS coordinates and north-aligned quaternions) to a predefined MQTT broker, which is usually residing on the host. Any visualization software (AR or non-AR) can connect to the MQTT broker to consume the actual sensor data at high frame rates.

3.2 Runtime operation

For good rendering precision, the outdoor AR application operates in local coordinates. Due to the limited numeric precision of datatypes used by rendering engines, any content (SUE data in our case) needs to

⁴AR4 GmbH: <https://www.ar4.io>



Fig. 5. Procedure to acquire calibration data: After assembling a rigid sensor setup, the user acquires the dataset. It is essential to perform a set of rotational and translational movements, as well as displacements of the sensors with respect to a calibration pattern in order to capture enough variation for the calibration method to perform properly.




Fig. 6. Importance of IMU-to-camera calibration: (left) The building model is rendered with a proper calibration. (right) Through using an inaccurate calibration, the rotational error around all axes accumulates to about 6° , which leads to a clear displacement of the virtual model from reality.

be transformed into this local coordinate frame before rendering. For tracking position and orientation of a user, two options are available:

1. On devices with on-board SLAM features, the local frame is automatically created and maintained by SLAM. To establish a global reference, we require only a single 6-DOF position and orientation measurement right at the beginning, *i.e.*, a GPS measurement and a north-aligned orientation measurement from the IMU. During further operation, we assume that re-localization and map refinement of the SLAM system ensure drift-free operation.
2. If no SLAM is available on the host computer (*e.g.*, when using a tablet running Windows), the sensor box can be used at full frame rate as the sole source of tracking. In other words, the application resorts to purely non-visual tracking served by the sensor box, performing global pose estimation and registration on a frame-to-frame basis.

In both cases, the IMU-to-camera calibration plays a crucial role to properly align the coordinate frame in terms of global rotation. On the one hand, in SLAM-enabled systems, the gravity vector is usually known, but alignment to north is needed at least once at the beginning. On the other hand, when using pure inertial tracking, the need for proper mapping of the full 6-DOF pose estimate from the sensor box to the camera at frame rate is apparent. Even small errors accumulate and lead to a clear misregistration of virtual data (see Fig. 6 for example).

3.3 IMU-to-camera calibration

Because the overall tracking quality strongly depends on it, we must ensure proper IMU-to-camera calibration. We need a proper mapping of axes between the IMU and the camera, such that an axial rotation measured by the IMU is transformed into a proper rotation matrix matching the motion perceived by the camera. The underlying problem is a very common and extensively studied topic in robotics [23], as it is fundamental in visual-inertial SLAM [26]. We use Kalibr⁵ to perform the required computations. The algorithms of Kalibr [15, 24] were not modified; our contribution is a framework that actually makes it easy for untrained users to perform the required calibration quickly without requiring special knowledge.

Ease of use is essential for our application area, as devices cannot be calibrated ahead of time. In contrast to rigid sensor arrangements, *e.g.*, in drones, which are calibrated at design time or during manufacturing, the calibration of our external sensor box to the host device's camera is only possible after assembly and must be performed by the user. Calibration depends on the form factor of the device used, the available mounting options, or may even vary with the application. The proper transformation matrix between the IMU and the camera as delivered by the server will remain valid for any number of runs or applications, as

⁵<https://github.com/ethz-asl/kalibr>




Fig. 7. Ray casters are placed along the geo-spatial utility data at a fixed height, shooting rays downwards to check for collisions with a 3D mesh reconstruction of the ground surface.



Fig. 8. Example of overlay blend mode used to blend the spray paint texture on the top layer with different surface textures on the base layer

long as the rigid configuration is not changed. If the configuration is modified, the recording and the calibration have to be repeated.

We propose a method comprised of two stages: A client component collects all relevant information on a device paired with the sensor box, and a server component in the cloud runs the computationally expensive part of the algorithm. The client component records a video file and stores timestamps for individual frames in a local database. At the same time, sensor information (*i.e.*, the absolute rotation and acceleration of the IMU recorded at 120 Hz) is retrieved through MQTT and stored in the database as well. Upon finishing the recording, the database and video file are uploaded to a cloud server (hosted on Amazon Web Services), which exposes an API to Kalibr via Node-RED⁶. Upon finishing a calibration run, the user is automatically notified by email of the intrinsic camera matrix and the extrinsic transformation describing the relative position and orientation between IMU and camera center.

A regular camera calibration typically requires 50-100 frames, but the timing of these frames is not important, provided the scene does not change. In contrast, calibrating the IMU to the camera requires accurate timestamps for both data sources (the IMU and the camera frames) with thousands of samples in order to infer the proper geometric constellation between the two entities. It is therefore important to capture sufficient translational and rotational motion during the recording of the calibration dataset in a time-synchronized fashion (Fig. 5).

In a typical calibration run, the user captures a dataset with respect to a calibration target with known properties for a length of 2-3 minutes, properly stimulating the IMU gyroscope and accelerometer along and around all axes through alternating slow and fast motion.

Our method calibrates only IMU to camera and does not recover the geometric relationship of the GPS antenna to the rest of the setup. Calibrating the GPS antenna position directly with respect to the camera center would be extremely challenging and is better resolved by relying on real physical measurements taken from the antenna mount.

4 VIRTUAL UTILITY MARKING

Civil engineers and construction workers use spray paint marking on the ground surface to indicate the type and location of underground infrastructure. These markings stand out from the environment because of their bright colors and, simultaneously, use conventions and visual encodings (*e.g.*, glyph shapes) familiar to the stakeholders on location. Commercial solutions display markings by simple overlay of renderings created from the respective CAD model, which lack realism and do not match the appearance of the site. Plausible display of virtual spray paint requires that

1. the placement of markings follows the ground surface;

⁶<https://nodered.org/>

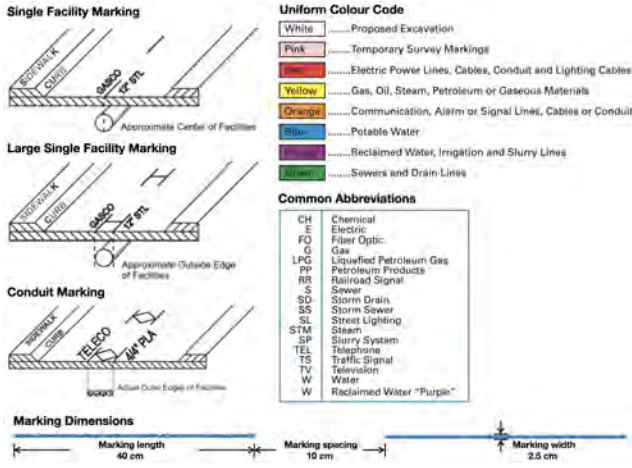


Fig. 9. Examples of common utility marking designs, colors, abbreviations and dimensions sourced from uniform marking guidelines.

2. the markings blend into the ground surface;
3. the design of the markings follows the styles, symbols and colors set forth in industry guidelines.

These requirements have only been partly addressed in previous work, which used simplistic marking placement and texture (Fig. 2). In this section, we describe our visualisation technique, *virtual utility markings*, which aims to address these requirements in full. As a prerequisite, we assume that SUE data is already imported and has been converted to *marking primitives*, i.e., 2D or 3D graphical primitives, such as poly-lines or tubes. Moreover, we assume the AR system provides a *ground mesh*, i.e., a 3D mesh reconstruction of the ground surface or, at least, a ground plane estimate, in real-time.

4.1 On-surface placement

To obtain the on-surface effect, we place ray casters along and above the marking primitives to check at which location the vertical rays intersect with the ground mesh. This is illustrated in Fig. 7. The ray origins are placed at fixed height above the utility line with a downwards casting direction. We use the farthest intersection of a ray with the reconstructed mesh, i.e., the lowest surface point, for placing the markings. In case any physical objects above the ground surface are captured by the mesh reconstruction process, for example, the arm of an excavator or other construction site equipment, this strategy ensures that the visualizations are always placed on the ground surface.

From the position and orientation of the surface point hit by the ray, "paint" particles, i.e., semi-transparent billboards with a size of 16×16 pixels, are instantiated. We use a resolution of 30 particles per meter to ensure that they appear as a complete line. Placing dense particles mimics the way real spray paint behaves, making the marking follow the curvature of the ground surface closely.

4.2 Spray paint texture

We use an *overlay blend mode* as known from image editing software to create a realistic appearance of semi-transparent paint. We derive a blended color by combining per-pixels luminance values of the ground layer, denoted by $g \in [0, 1]$, and the top layer, denoted by $t \in [0, 1]$, using the formula

$$f(g, t) = \begin{cases} 2gt, & \text{if } t < 0.5 \\ 1 - 2(1 - g)(1 - t), & \text{otherwise,} \end{cases} \quad (1)$$

where f is computed in a pixel shader, g is obtained by projecting the video texture of the AR device to the ground mesh, and t is obtained by rendering the particles into an off-screen framebuffer (Fig. 8).

4.3 Standardized marking templates

To facilitate clear and consistent communication, SUE professionals have developed industry guidelines for marking, standardizing symbols, styles, abbreviations, colors and dimensions. Inspired by these

guidelines, we have implemented procedural visualization templates of some common marking designs (Fig. 9). The type of marking is determined from the attributes of the SUE data, and the appropriate template is chosen. The templates can include plain outlines, simple glyphs of vectorized outlines, e.g., a diamond-shaped conduit, and text. Thereby, the virtual spray paint is not only applied to indicate simple lines following real utility tracks, but also to place such signs and glyphs at locations indicated by the SUE data attributes.

5 VIRTUAL DAYLIGHTING

While on-surface markings can quickly convey standard SUE information, virtual daylighting allows a detailed view of subsurface infrastructure, which may be important in ambiguous or dangerous situations. Automatically acquired laser scans of excavation sites are often large, in the order of 1-10 million points per site, which is challenging for a mobile GPU. Since we want to support ad-hoc work practices where timely preprocessing (e.g., geometric simplification) is not really an option, an alternative strategy is required to cope with the significant amount of data.

Our rendering solution is based on Potree⁷, an open-source point cloud renderer for the web [32]. We have integrated the renderer directly into Unity using the method of Fraiss [14], such that it can run entirely on the device without involving cloud rendering and inducing communication latency. The cloud service is therefore only responsible for preprocessing, to serve the point cloud data, and to create occlusion information created, as described below.

As demonstrated by Zollmann *et al.* [45], contradictions within the virtual and the real world are confusing to users. Unfortunately, since we rely on 3D point clouds without explicit surface information, occlusion cannot be resolved by simple backface culling. To tackle this small but important problem, we developed a method to automatically create occluders from scanned data sets, which tightly fit the excavation holes, enabling ad-hoc use of the scanned models in AR. Our goal is to find the contour of the excavation hole at street level. We assume that all walls of the excavation hole are vertical, so that a polygonal extrusion of the street-level contour downwards to the lowest point of the scan delimits the excavation in a way that is suitable for occlusion rendering. The challenge is to find the street-level contour without knowing the exact street level elevation in the scanned data.

Our method accepts as input either unstructured point clouds (as acquired with LIDAR) or 3D reconstructions created from photo sets via photogrammetry. First, we compute point normals for each point using a disk-based Poisson distribution. The normals will be used to vote for the street level: While points inside the hole belong to the walls and tend to have horizontal normal, the points on and above street level tend to have vertical normals.

Starting at the lowest point of the excavation, we sweep a horizontal plane upwards and statistically evaluate all scanned points inside slices with a depth 2ϵ , where ϵ was empirically set to $0.2mm$ as a compromise between the number of slices and the average number of points per slice. When we observe a sudden switch from horizontal to vertical normals, the sweep stops, and the points just below the street level are selected to form the contour of the excavation hole.

The result is applied as a stencil drawn on a plane at ground level, which is served alongside the actual point cloud to the end-user application, masking the area outside the contour during rendering of the scanned model. The result are realistic-looking occlusions of a virtual excavation hole (Fig. 10).

6 EXPERIMENTS

In order to show the impact of our work over existing approaches, we conducted several experiments. We first show some results for the proposed calibration method, followed by some use cases in outdoor AR applications and an expert study to assess the plausibility of our new visualization methods.

⁷<https://github.com/potree/potree>




Fig. 10. Plane-sweeping algorithm applied to excavation data: (top left) estimated normals of points, where green denotes orthogonality to street level, (top right) individual contours estimated from bottom to top, (bottom left) intermediate contour overlaid on colorized point cloud, (bottom right) top-down view of final occluder applied on colorized point cloud

Rotation [°]			Translation [mm]		
	mean	std. dev.		mean	std. dev.
r_x	-2.5391	0.9908	t_x	51.9249	4.6588
r_y	0.6609	0.3704	t_y	207.9650	7.6064
r_z	-178.7168	0.4641	t_z	35.7136	7.8353

Table 1. Results for our 11 calibration runs. Both rotation and translation estimates are plausible as compared to the setup shown in Fig. 11.

6.1 Calibration results

The accuracy of the differential GPS-RTK sensor in our sensor box falls into the range of 2 – 4cm for position measurements. For the IMU, the error in orientation measurements is well below 1°. Thus, the global pose estimates delivered by the sensor box already reach the required accuracy for geodesy applications. For more detailed evaluation results, the reader is referred to Stranner *et al.* [35]. In the following experiment, we are concerned with the accuracy of our calibration method, which maps the global pose estimates into the camera coordinate system, ultimately solving the misalignment problem depicted in Fig. 6.

Based on our distributed setup, we performed a total of 11 calibration runs with a rigid configuration of the iPad and our sensor box in the sample configuration shown in Figure 11, together with the camera coordinate system and the estimated IMU coordinate systems for said runs. The data points clusters consistently around the real position of the IMU. On the right of Fig. 11, a side view is shown, revealing the offset in z-direction. In Tab. 1, the mean and standard deviation of rotation and translation is shown. The results are accurate, coming at the ease of an almost fully automatic approach.

During our developments, we had two insights that were not immediately apparent from the documentation and literature of Kalibr. First, it is crucial to consider the accuracy of the manual measurements (*i.e.*, the size and spacing of individual tags) taken from the physical calibration target. This information needs to be provided to the application before uploading recordings in order to introduce metric scale. Even an error of less than 1mm in the target causes the calibration to deviate by 1cm or more in translation. Second, in order to arrive at an accurate result, the calibration run needs to incorporate significant translational and rotational (spontaneous) motion (Fig. 5). Insufficient motion results in inaccurate and useless transformation estimates. It is of utmost importance that the calibration recordings are taken with due diligence.

6.2 Accuracy and repeatability of virtual utility markings

We performed an evaluation of initial accuracy and repeatability of the localization delivered by the fusion of our sensor box with SLAM. We omit a dedicated evaluation of pure sensor-based tracking (*i.e.* when

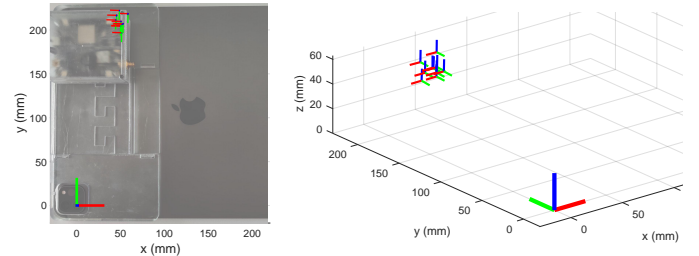


Fig. 11. Calibration results for a prototypical sample configuration: (left) camera coordinate system and individual coordinate systems overlaid on a top-view of the calibration setup. The individual IMU estimates cluster around the real IMU location (axes color coded as x=red, y=green, z=blue), (right) side-view revealing the offset in z-dimension.



Fig. 12. Setup for our marking repeatability and accuracy test: (left) A tape measure is placed perpendicular to measure displacements of virtual spray marks. We measure (middle) the offset at the reference point and (right) the displacement at the 1m, 2m and 3m marking.

SLAM is not available), as it can be reduced to the raw sensor accuracy. Consequently, results only depend on the quality of the GPS fix, *i.e.* if a fixed-integer fix is reached, the accuracy is always < 2cm at that particular point.

We sprayed four virtual concentric circles around a geodesically verified reference point (RP) near our laboratory at a radius of 10cm, 1m, 2m and 3m with a tape measure laid out for reference (see Fig. 12). We collected measurements using the following procedure: (i) The user starts at a random position within 4m of the RP. (ii) The user walks around the RP to obtain a SLAM map. (iii) The user takes screen snapshots along the tape measure at distances of 10cm, 1m, 2m, and 3m. Snapshots are automatically annotated with timestamp, initial GPS position and snapshot position. This procedure was repeated on three different days, obtaining a total of 15 measurements.

In the snapshots, we can read the displacement of the virtual markings on the tape measure and calculate the initial distance Δ of the device to the RP, the displacement from the center (ϵ) and the deviation from meter markings (ϵ_{1m} , ϵ_{2m} , ϵ_{3m}), which are summarized in Tab. 2.

Date	RTK	Δ	ϵ	ϵ_{1m}	ϵ_{2m}	ϵ_{3m}
31/5/21, 12:35	float	47	12	3	1	1
31/5/21, 12:50	fixed	85	10	11	8	10
31/5/21, 12:45	fixed	287	2	8	10	11
31/5/21, 12:38	fixed	291	1	1	3	9
31/5/21, 12:53	fixed	207	1	1	4	5
1/6/21, 15:54	float	91	30	23	21	19
1/6/21, 15:27	float	365	25	13	16	20
1/6/21, 15:34	fixed	199	9	6	10	12
2/6/21, 11:28	float	122	17	20	27	35
2/6/21, 11:24	float	264	15	21	30	33
2/6/21, 11:31	float	157	15	10	5	2
2/6/21, 11:34	fixed	59	10	11	9	7
2/6/21, 11:18	fixed	66	10	2	2	2
2/6/21, 11:14	fixed	281	6	2	2	4
2/6/21, 11:07	fixed	165	1	3	3	4

Mean (μ) fixed only	5.56	5.00	5.67	7.11
StdDev (σ) fixed only	4.28	4.12	3.50	3.55
Mean (μ) float only	19.0	15.0	16.67	18.33
StdDev (σ) float only	6.96	7.72	11.71	14.58
Mean (μ) fixed and float combined	10.93	9.0	10.07	11.6
StdDev (σ) fixed and float combined	8.61	7.53	9.33	10.75

Table 2. Virtual utility marking evaluation results: All measurements are taken in centimeters. Δ refers to the initial GPS distance to the reference point. ϵ is the measured virtual marking error at 0m, 1m, 2m and 3m.

The use of SLAM accumulates drift with increasing tracking distance. With GPS in "floating" mode (DGPS, *i.e.*, only corrections, but no phase information), we expect at least an initialization error of 30cm, while, in "fixed" mode (RTK-GPS), of at least 2cm, as shown by the evaluation data. For the intended use case, namely running "fixed" mode GPS while being within 3m distance of the RP, the errors are within a margin of 10cm.

A closer look into the expected precision of hardware components reveals the following characteristics: (i) GPS-RTK gives < 2cm of error, (ii) errors concerning compass based heading strongly depend on the environment and fall within a range of 1° – 5°, and (iii) local tracking drift [42] over a 4m distance is $\leq 9cm$ on tarmac. The measurements taken in our evaluation, as shown in Figure 13, confirm this behavior and initialization errors from up to 10cm. We observed higher SLAM drifting when facing strong direct sunlight, as it causes more interference at the test location. Overall, the results deliver the precision required for SUE applications.

In general, the number of observed satellites is neither directly related to the quality of the GPS fix, nor is it reported by our sensor. In contrast, the quality of the receiving antenna has the largest impact on the fix. However, depending on the satellite constellation at a particular site, a cheap antenna can also deliver the same results at the cost of a longer initialization time. The reader is referred to reports on GNSS antennae [39] and receiver performance [16].

6.3 SUE information and datasets

SUE data generally refers to all kinds of information about buried or hardly accessible subsurface infrastructure, with a considerably varying degree of accuracy, completeness and, more importantly, availability.

A SUE database usually contains information about cable or pipe specifications, purpose, owner, and location. The latter is represented in 2D only, *e.g.*, 2D points or line strips, while 3D information about the sub-surface depth or utility width is scarce, often based on manual measurements and subject to regional documentation practices. The location accuracy may vary significantly in the centimeter range, notwithstanding additional sources of error induced through the conversion between coordinate systems (*e.g.*, Gauss-Krüger to UTM) during documentation and later remapping to reality. While these circumstances make it difficult or impossible to specify a single ground truth for localization accuracy of geo-referenced data, it is accepted practice that the utility should be discoverable within $\pm 30cm$ from the spot where it is marked in the GIS data.

For our experiments, we used two datasets. The first dataset includes geo-spatial subsurface utility data from a suburban area in Copenhagen,




Fig. 13. Position specific initialization error: The origin depicts the reference point (RP), and the position of the mark shows absolute distance from the RP. The color encodes the displacement error between the virtual spray mark center and the RP. Lower is better. A less accurate GPS-RTK signal ("RTK float", triangles) results in worse results.

Denmark. The dataset consists of points and poly-lines representing underground facilities like manholes, wells, pipes and cables. The geo-spatial data is further divided into utility supply types such as electricity, gas, water, telecom, *etc.* The geo-spatial data also has attributes such as material type, thickness, cross-section shape, accuracy tolerance zone, *etc.* The dataset is provided by the *Danish register of underground cable owners* as a test dataset to educate others on how geo-spatial utility data should be structured following their GML data model. For productive use, the same information can be directly retrieved from a live GIS database using proper credentials.

The second dataset consists of geo-located 3D reconstructions of utility excavation holes gathered from a nearby area. The 3D reconstructions are represented as point clouds and are captured by the local water utility company. The company uses these 3D reconstructions as a supplementary as-built documentation for their water utility facilities. Using 3D reconstructions techniques as a means of documentation during excavation work was introduced in 2019. Currently the goal is to explore which value creation and benefits these new technologies have in practice.

6.4 Outdoor virtual utility marking and daylighting

In order to demonstrate the visual appeal of our spray paint effect, we present selected results from multiple scenarios, including showcases illustrating virtual daylighting.

Fig. 14 (top left) shows a surface reconstruction obtained with LIDAR, which is used to draw the markings directly onto the surface. Because the reconstruction closely follows the real surface, the markings align neatly across the curb (see top mid of Fig. 14). Fig. 14 (top right) shows two examples of multiple lines with different colors, demonstrating the impact of real-world shadows on the appearance of the lines. The shadows do not harm the visual effects at all, and the spray markings appear coherently, because the spray effect directly incorporates the real surface appearance captured by the camera. In Fig. 14 (bottom left), two results are shown comparing our approach to a traditional shadow-projection visualisation, which only targets planar surfaces that are estimated at a fixed height. Because the surfaces (sink and curb) violate the planarity assumption, the shadow projections deviate from the real location of the subsurface utilities. Leveraging real data from GIS, Fig. 14 (bottom mid) depicts the results of our spray paint visualization. The markings look very realistic and are placed at their correct location. This is especially noticeable at surfaces with darker-colors and rich textures, for instance, on a wet pavement. In case the surface has been reconstructed properly, temporary occlusions can be resolved [2] (see Fig. 14 bottom right).

The previous examples show the use of our system on closed road surfaces. To demonstrate the accuracy with respect to real buried infrastructure, Fig. 15 shows an overlay of virtual utility markings on top of an actual excavation. As can be seen, the overlays deviate from the uncovered assets by an offset of a few centimeters. Nonetheless, finding and identifying the assets during excavations is successfully guided by the overlays.



Fig. 14. Outdoor virtual utility marking using real GIS data: (top left) The mesh generation from LIDAR closely follows the real surface. (top mid) Virtual spray markings showing glyphs which respect the elevation change. (top right) Rendering with dynamic illumination. As the spray effect directly incorporates the shadows as seen by the camera, the spray markings fuse with reality in an indistinguishable way. (bottom left and inset) The orange dashed lines are showing the location of the subsurface utility as a naive overlay on the video, while the blue lines show the spray following the reconstructed surface. (bottom mid) Real subsurface utility data overlaid on wet surfaces. (bottom right) Using a feature to detect people, dynamic occlusion is correctly resolved.

In analogy to GIS data from our database, we use the geo-located 3D point clouds to show excavations at site. In Fig. 16, a scenario is shown with both the unaltered view of the user and the virtual excavation. Because our automatically generated occluders carve out the reconstructed areas surrounding the ditches of interest, the visualizations align with existing structures at the contours of the respective 3D excavation holes.

6.5 Semi-structured interviews with expert users

We performed semi-structured interviews with a group of domain experts in order to assess the plausibility and applicability of our virtual marking and daylighting approaches. Based on a previous series of interviews [17], it was expected that AR could potentially help prevent excavation damage. The interviews were conducted on site within the first half of 2021, comparing a live try-out of our AR prototype to commercial AR visualizations as shown in Fig. 2. Six experts from two companies with an average experience in the field of 20 years were interviewed, belonging to the contractor, utility and surveying sector. We put the focus on three topics: (i) current SUE and excavation practice, (ii) assessment of visualisation approaches and (iii) applicability during SUE and excavation work.

Current SUE and excavation practice The interviews exposed major issues and pitfalls, such as damaging existing infrastructure through the absence or the poor quality of documentation, particularly concerning soft cables in electric and telecom utilities. Statements such as *"Yesterday we just hit such a cable, and it's just so frustrating."* or *"You can not see from a utility drawing if it is in one or the other sidewalk tile."* are symptomatic for situations, which one expert summarized as: *"So, as a starting point, there is a presumption that there is something, but depth and width and such – we have to see what it is like as we continue to dig."*

Concerning spray markings, experts explained that the main issue is their application in practice: *"Spray markings also disappear. If it is wet, then you cannot spray. If there is dust, then you cannot spray. Therefore, utility marking with spray has a limited use, and, in addition, the road authority would like it removed again."* One expert stated, *"Yes, I think we should be able to avoid excavation damage more often."*

So if it was marked more often, then people in the field will probably be more attentive." The cost of excavation damage was further elaborated, and one expert stated that just avoiding 10% of their insurance cases would already be a significant saving.

Assessment of visualisation approaches All experts were very positive towards the virtual utility marking and daylighting approaches, as it provided a more comprehensive understanding of the underground beneath the excavation site in terms of *what is there* and *where it is*. When asked about how understandable the visualizations were, one expert said regarding markings, *"It was almost better than in reality and it seems very useful."* Another one said *"It gave you a lifelike feel. Sometimes I asked myself, is it real spray markings?"* A third expert elaborated further, *"It was very comprehensible and somehow more pedagogical – especially, because it also said if it was a gas line or an electric line or something else and because the virtual lines had the colors that they [real spray marking] usually have."* We conclude from these statements that the experts could relate to the virtual markings and they created a sense of familiarity with their experience from real spray markings.

When compared to commercial AR visualisations, as shown in Fig. 2, the experts felt that our virtual markings were more suited for their work tasks. One expert described why she did not prefer the commercial AR visualizations by saying, *"It is very messy to look at, and the 3D effect is not intuitive to understand. It is as if you yourself have to move the 3D models parallel below the surface."* She followed up by saying, *"I like the other one [virtual markings] better. It was more understandable and more direct for our use."* Another expert said, *"It was floating-like."* This reinforced our assumptions that it is essential to overcome the lack of parallax in commercial AR visualizations, which causes the sub-surface models to appear as if they were floating above the surface.

Regarding virtual daylighting, one expert said, *"In our industry, we have always wanted some kind of X-ray glasses, and this is the closest we come; this is the closest I have seen to actual X-ray vision."* The majority of the experts said that having depth information about the utility lines was the most valuable aspect provided by virtual daylighting. One expert further elaborated that he believes the comprehensible overview



Fig. 15. Virtual utility markings overlaid onto real excavations. Because our algorithm is able to mark lines even during excavation works, it is possible to validate the accuracy of GIS data with respect to the actual physical location of utilities. Markings deviate by a few centimeters, in particular from soft cables, but the AR overlays successfully guide the daylighting.




Fig. 16. Outdoor virtual utility daylighting: (left) The real world as seen by the user. (right) The virtual models from the excavations are overlaid at the respective locations. Due to the occluding geometry wrapping around the excavations, the embedded 3D reconstructions appear very realistic and deliver the impression of looking into real holes.

of the utility excavation is the biggest benefit. With virtual daylighting, the utility information is no longer just, as this expert phrased it, “*small blue and red lines in a drawing.*”

Applicability during SUE and excavation work The experts deemed virtual markings useful, with one expert stating, “*I could really see this in the hands of our people in the field. It represents utility information that they are used to look at on drawings.*” When asked how it would be useful, one experts summarizes as follows, “*This will provide direct value when planning and digging holes. It would reduce excavation damage and work environment risks associated with the excavation activities we do in the field.*” Experts also valued that the virtual markings would always be visible, whereas real markings may disappear. It was also perceived as valuable that virtual markings are not permanently painted onto the surface, as physical paint would. One expert said: “*I could really see it in use, because you do not have to remove it afterwards.*”

The experts were similarly positive towards the usefulness of virtual daylighting. One aspect compared to virtual markings that they found useful was the added depth information, which is often lacking in existing utility documentation. In this regard, all experts agreed that any tool that gives more insight to an excavation site reduces cost and efforts (time, damage and repair). One expert emphasized that a high level of realism in 3D-captured models is beneficial, “*As a contractor, you want to see it all.*”

Overall, the experts had surprisingly little concerns. One expert went as far as saying, “*I can’t imagine why using this [virtual markings] would be a bad idea. Using this won’t limit you to also use traditional printed or digital utility drawings.*” One remaining concern mentioned by several experts comes from the inaccurate and incomplete SUE data sourced from existing utility drawings and GIS databases. The expert saw a danger that the virtual utility markings would convey a false sense of trustworthiness. One expert suggested, “*A kind of disclaimer must be added that the utility information that is displayed may be inaccurate and incomplete.*” However, he did not think it would be a

deal-breaker, as the information is still the same information that field workers would normally read from a drawing. Availability of AR may actually be an incentive to improve SUE data quality.

7 DISCUSSION AND CONCLUSION

In this work, we revisit the topic of subsurface utility engineering and refresh the perception that AR is important for this domain using contemporary technology. In particular, we described a hardware/software solution for highly accurate outdoor localization and use it to drive two novel techniques for visualizing subsurface infrastructure, namely virtual utility marking and virtual daylighting. We demonstrate how these visualizations come very close to reality as currently used in the field. The expert user interviews confirmed that our results are appealing and closer to practical usage than previous approaches.

Although we consider our methods mature enough for practical use, poor data quality in GIS databases remains a problem. Intervals between physical daylighting of subsurface utilities may be many years long, and updates to utility documentation are tardy. Many entries in current utility databases were originally imported from outdated or erroneous paper plans, often missing accurate location information. Nonetheless, the situation is improving. In an effort by the European commission to build *smart cities* [11], new regulations were put in place to collect infrastructure data in federated databases. We expect that such legislation, combined with the wider spread of GIS, will steadily improve data quality over time.

While standards for AR are still lacking, the availability of richer geo-located content for other purposes, such as urban planning or map services, will enhance the quality of AR experiences as well. Our work is an attempt to demonstrate the long way AR has come towards reaching a useful state for SUE and other outdoor applications. We believe the quality is now sufficient to convince even practitioners in conservative industries such as civil engineering.

ACKNOWLEDGMENTS

This work was supported in part by a grant from FFG (# 859208).

REFERENCES

- [1] A. J. Al-Bayati and L. Panzer. Reducing Damage to Underground Utilities: Lessons Learned from Damage Data and Excavators in North Carolina. *Journal of Construction Engineering and Management*, 145(12):1–8, 2019. doi: 10.1061/(ASCE)CO.1943-7862.0001724
- [2] Apple. Arkit people occlusion, 2021. Visited online, 3 Mar 2021, https://developer.apple.com/documentation/arkit/camera_lighting_and_effects/occluding_virtual_content_with_people.
- [3] C. Arth, M. Klopschitz, G. Reitmayr, and D. Schmalstieg. Real-Time Self-Localization from Panoramic Images on Mobile Devices. In *ISMAR*, pp. 37–46, 2011. doi: 10.1109/ISMAR.2011.6092368
- [4] C. Arth, C. Pirchheim, J. Ventura, D. Schmalstieg, and V. Lepetit. Instant outdoor localization and SLAM initialization from 2.5d maps. *TVCG*, 21(11):1309–1318, 2015. doi: 10.1109/TVCG.2015.2459772
- [5] C. Arth, D. Wagner, M. Klopschitz, A. Irschara, and D. Schmalstieg. Wide Area Localization on Mobile Phones. In *ISMAR*, pp. 73–82, 2009. doi: 10.1109/ISMAR.2009.5336494
- [6] R. Broome, P. Cornforth, S. Crossland, N. Metje, and A. Rhoades. Utility strike damage report, 2019. Visited online 3 Mar 2021, https://www.utilitystrikeavoidancegroup.org/uploads/1/3/6/6/13667105/usag_2019_strike_damages_report.pdf.
- [7] C.-T. Chang, R. Ichikari, K. Makita, T. Okuma, and T. Kurata. Road maintenance mr system using lrf and pdr. In *ISMAR*, pp. 204–205, 2015. doi: 10.1109/ISMAR.2015.66
- [8] S. Côté and A. Mercier. Augmentation of road surfaces with subsurface utility model projections. In *VR-adj.*, pp. 535–536, 2018. doi: 10.1109/VR.2018.8446545
- [9] DIRT. Dirt annual report of the common ground alliance, 2019. Visited online, 3 Mar 2021, <https://commongroundalliance.com/Portals/0/Library/2020/DIRT%20Reports/2019%20DIRT%20Report%20FINAL.pdf?ver=2020-10-14-185343-180>.
- [10] M. T. Eren and S. Balcisoy. Evaluation of x-ray visualization techniques for vertical depth judgments in underground exploration. *The Visual Computer: International Journal of Computer Graphics*, 34(3):405–416, Mar. 2018. doi: 10.1007/s00371-016-1346-5
- [11] European-Union. Smart cities: Cities using technical solutions to improve the management and efficiency of the urban environment, 2021. Visited online, 3 Mar 2021, https://ec.europa.eu/info/eu-regional-and-urban-development/topics/cities-and-urban-development/city-initiatives/smart-cities_en.
- [12] S. Feiner, B. MacIntyre, T. Höllerer, and A. Webster. A touring machine: Prototyping 3d mobile augmented reality systems for exploring the urban environment. *Personal Technologies*, 1(4):208–217, 1997.
- [13] A. Fenais, S. Ariaratnam, S. Ayer, and N. Smilovsky. Integrating geographic information systems and augmented reality for mapping underground utilities. *Infrastructures*, 4:60, 09 2019. doi: 10.3390/infrastructures4040060
- [14] S. M. Fraiss. Rendering Large Point Clouds in Unity. Master's thesis, Vienna University of Technology, 2017.
- [15] P. Furgale, T. D. Barfoot, and G. Sibley. Continuous-time batch estimation using temporal basis functions. In *ICRA*, pp. 2088–2095, May 2012. doi: 10.1109/ICRA.2012.6225005
- [16] D. Gebre-Egziabher. Evaluation of low-cost, centimeter-level accuracy oem gnss receivers. <https://www.dot.state.mn.us/research/reports/2018/201810.pdf>, 2018. Visited Jun 28, 2021.
- [17] L. Hedegaard Hansen, S. Swanström Wyke, and E. Kjems. Combining reality capture and augmented reality to visualise subsurface utilities in the field. In *ISARC*, pp. 703–710. Kitakyushu, Japan, October 2020. doi: 10.22260/ISARC2020/0098
- [18] G. Hillar. *MQTT Essentials - A Lightweight IoT Protocol*. Packt, 2017.
- [19] G. Klein and D. Murray. Parallel tracking and mapping for small AR workspaces. In *ISMAR*. Nara, Japan, November 2007.
- [20] G. Klein and D. Murray. Parallel tracking and mapping on a camera phone. In *ISMAR*, pp. 83–86, 2009. doi: 10.1109/ISMAR.2009.5336495
- [21] W. Li, Y. Han, Y. Liu, C. Zhu, Y. Ren, Y. Wang, and G. Chen. Real-time location-based rendering of urban underground pipelines. *ISPRS International Journal of Geo-Information*, 7(1), 2018. doi: 10.3390/ijgi7010032
- [22] L. O. Makana, N. Metje, I. Jefferson, M. Sackey, and C. D. Rogers. Cost Estimation of Utility Strikes: Towards Proactive Management of Street Works. *Infrastructure Asset Management*, pp. 1–34, 2018. doi: 10.1680/jinam.17.00033
- [23] F. M. Mirzaei and S. I. Roumeliotis. A Kalman Filter-based Algorithm for IMU-Camera Calibration: Observ. Analysis and Perf. Eval. *IEEE Trans. on Robotics*, 24(5):1143–1156, 2008. doi: 10.1109/TRO.2008.2004486
- [24] R. S. Paul Furgale, Joern Rehder. Unified temporal and spatial calibration for multi-sensor systems. In *IROS*, pp. 1280–1286. Springer International Publishing, Tokyo, Japan, 2013.
- [25] W. Piekarski and B. H. Thomas. Tinmith-metro: new outdoor techniques for creating city models with an augmented reality wearable computer. In *ISWC*, pp. 31–38, 2001. doi: 10.1109/ISWC.2001.962093
- [26] T. Qin and S. Shen. Online temporal calibration for monocular visual-inertial systems. In *IROS*, pp. 3662–3669. IEEE, 2018.
- [27] G. Reitmayr and T. Drummond. Going out: Robust model-based tracking for outdoor augmented reality. In *ISMAR*, pp. 109–118, 10 2006. doi: 10.1109/ISMAR.2006.297801
- [28] G. Roberts, A. Evans, A. Dodson, B. Denby, S. Cooper, and R. Hollands. The use of augmented reality, gps and ins for subsurface data visualisation. In *XXII FIG Congress*, 04 2002.
- [29] G. Schall, E. Méndez, E. Kruijff, E. E. Veas, S. Junghanns, B. Reitinger, and D. Schmalstieg. Handheld augmented reality for underground infrastructure visualization. *Personal and Ubiquitous Computing*, 13(4):281–291, 2009. doi: 10.1007/s00779-008-0204-5
- [30] G. Schall, E. Mendez, and D. Schmalstieg. Virtual redlining for civil engineering in real environments. In *ISMAR*, pp. 95–98, 2008. doi: 10.1109/ISMAR.2008.4637332
- [31] G. Schall, D. Wagner, G. Reitmayr, E. Taichmann, M. Wieser, D. Schmalstieg, and B. Hofmann-Wellenhof. Global pose estimation using multi-sensor fusion for outdoor augmented reality. In *ISMAR*, pp. 153–162. IEEE Computer Society, 2009. doi: 10.1109/ISMAR.2009.5336489
- [32] M. Schuetz. Potree: Rendering large point clouds in web browsers. Master's thesis, Institute of Computer Graphics and Algorithms, Vienna University of Technology, Sept. 2016.
- [33] G. Soria, L. Ortega, and F. Feito. Augmented and virtual reality for underground facilities management. *Journal of Computing and Information Science in Engineering*, 18:041008, 07 2018. doi: 10.1115/1.4040460
- [34] J. C. Spohrer. Information in places. *IBM Systems Journal*, 38(4):602–628, 1999. doi: 10.1147/sj.384.0602
- [35] M. Stranner, P. Fleck, D. Schmalstieg, and C. Arth. A high-precision localization device for outdoor augmented reality. In *ISMAR-adj.*, 2019.
- [36] E. Stylianidis, E. Valari, K. Smagas, A. Pagani, J. Henriques, A. Garca, E. Jimeno, I. Carrillo, P. Patias, C. Georgiadis, A. Kounoudes, and K. Michail. Lara: A location-based and augmented reality assistive system for underground utilities' networks through gnss. In *VSMM*, pp. 1–9, 10 2016. doi: 10.1109/VSMM.2016.7863170
- [37] X. Su, S. Talmaki, H. Cai, and V. Kamat. Uncertainty-aware visualization and proximity monitoring in urban excavation: a geospatial augmented reality approach. *Visualization in Engineering*, 1, 12 2013. doi: 10.1186/2213-7459-1-2
- [38] S. A. Talmaki, S. Dong, and V. R. Kamat. Geospatial databases and augmented reality visualization for improving safety in urban excavation operations. In *Proceedings Construction Research Congress*, pp. 91–101, 2010. doi: 10.1061/41109(373)10
- [39] Trimble. Antenna list. <http://tr1.trimble.com/docushare/dsweb/Get/Document-882667>, 2020. Visited Jun 28, 2021.
- [40] S. E. O. Trujillo, J. Wendel, J. M. Santana, S. M. Murshed, I. Boates, A. R. T. Pino, A. Nichersu, and J. P. S. Rivero. Making the invisible visible - strategies for visualizing underground infrastructures in immersive environments. *ISPRS International Journal of Geo-Information*, 8, 2019. doi: 10.3390/ijgi8030152
- [41] J. Ventura and T. Höllerer. Wide-area scene mapping for mobile visual tracking. In *ISMAR*, pp. 3–12, 2012. doi: 10.1109/ISMAR.2012.6402531
- [42] vGIS. 2020 ipad pro: Does the lidar sensor improve spatial tracking? <https://www.vgis.io/2020/04/23/2020-ipad-pro-does-the-lidar-sensor-improve-spatial-tracking/>, 2020. Visited Jun 20, 2021.
- [43] X. Zhang, Y. Han, D. Hao, and Z. Lv. Argis-based outdoor underground pipeline information system. *Journal of Visual Communication and Image Representation*, 40, 07 2016. doi: 10.1016/j.jvcir.2016.07.011
- [44] S. Zollmann, R. Grasset, G. Reitmayr, and T. Langlotz. Image-based x-ray visualization techniques for spatial understanding in outdoor augmented reality. In *Proc. ACM OzCHI*, 12 2014. doi: 10.1145/2686612.2686642
- [45] S. Zollmann, D. Kalkofen, E. Mendez, and G. Reitmayr. Image-based ghostings for single layer occlusions in augmented reality. In *ISMAR*, pp. 19–26, 11 2010. doi: 10.1109/ISMAR.2010.5643546