POLITECNICO DI TORINO

Master of Science program in Building engineering

Master's degree Thesis

An Augmented Reality Framework for Geotechnical Monitoring

Candidate

Yeganeh Momen

Supervisors

Prof. Francesca Maria Ugliotti

Dr. Davide Aschieri

Academic Year2025-2026

Abstract:

Building Information Modelling (BIM) is now the standard framework for managing data in construction projects however, most of that data is on a desktop. It rarely reaches the people working on the actual site. This gap is especially clear for geotechnical and environmental data. Investigation campaigns produce detailed subsurface measurements, but these data end up in spreadsheets and reports with no real connection to the ground they describe.

This thesis addresses that gap by presenting GeoBIM Viewer, a mobile Augmented Reality (AR) application that brings geotechnical data from a BIM model straight to field professionals, at the location it describes.

The application is built in Unity 6 for Android, using AR Foundation and ARCore for tracking and Cesium for Unity for GPS-based positioning. The BIM model implemented in the app was developed in earlier research on the Città della Salute site in Turin, a former industrial brownfield now being redeveloped under the PSRI masterplan. The borehole geometry from that model is reconstructed as an interactive wireframe that users can explore either on-site or in an off-site preview mode. A data pipeline links this geometry to laboratory measurements, across fifty-seven contaminants, from eighty-one boreholes. On site, a user can select a borehole, check a specific depth, and see whether the measured concentration exceeds the legal threshold set by Italian regulation (D. Lgs 152/06). Spatial alignment between the digital model and the real site is achieved through the Cesium for Unity package, anchoring the model to known reference boreholes. The application also has an explore feature (offsite mode) using Cesium photorealistic terrain tiles which enables the user to explore the site by zooming and rotating with the context of the surrounding building and terrain.

The result is a working proof of concept, developed and tested on a single Android device. It shows that subsurface contamination data, otherwise locked inside specialist desktop software, can be anchored to its real-world location and put directly into the hands of the professionals working above it. By bringing together GeoBIM data, mobile AR, and GPS-based positioning in a single field tool, this work opens a practical path toward a use case that current research has yet to fully explore.

Keywords: Augmented reality (AR), Unity, Building Information Modeling (BIM), GeoBIM, Cesium for Unity

Table of Contents

Chapter 1 Introduction 4

1.1 Background and context of the problem 4

1.2 Research objectives and questions 5

1.3 Scope and limitations 6

1.4 Thesis structure overview 8

Chapter 2 Literature Review 9

2.1 GeoBIM for Subsurface and Geotechnical Data 9

2.2 Augmented Reality in Construction and Site Management 12

2.3 AR for Subsurface and Invisible Infrastructure Visualization 13

2.4 Research Gap and Positioning of This Work 17

Chapter 3 Case Study 19

3.1 Site Description and Historical Background 19

3.2 The Research Chain 21

Chapter 4 Methodology 26

4.1 Overview of the Approach 26

4.2 Inputs, Environment and Tools 27

4.2.1 Input data 27

4.2.2 Development environment 31

4.2.3 Project Configuration and Build Environment 32

4.3 System design 33

4.3.1 System architecture 33

4.4 Implementation of the Layers 39

4.4.1 AR Visualization and Interaction 39

4.4.2 Data Pipeline 42

4.4.3 Geospatial Positioning 44

4.4.4 Off-Site Exploration Mode 49

4.4.5 Build and deployment 50

Chapter 5 Results 51

5.1 Overview of the Implemented Application 51
5.2 AR Visualization and Interaction 54
5.3 Data Pipeline 55
5.4 Geospatial Positioning 58
5.5 Off-Site Exploration Mode 60
Chapter 6 Discussion 63
6.1 Integration of geometry and records 63
6.2 Alignment accuracy 63
6.3 Delivery of consumer hardware 65
6.4 Implications for site management 65
6.5 Limitations of the Study 66
Chapter 7 Conclusion 68
Appendix A: 70
References 71

Chapter 1 Introduction

1.1 Background and context of the problem

Over the past two decades, the construction industry has been transformed significantly with the progressive adoption of Building Information Modelling as the standard framework for project design, coordination, and data management. BIM moves beyond traditional two-dimensional drawings by embedding project information directly within three-dimensional geometric models, creating a unified digital representation that captures not only the physical form of a structure but also its associated data including material properties, structural characteristics, environmental measurements, schedules, and costs. In Italy, this shift was initially accelerated by Decree 560/2017, by which BIM adoption was progressively enforced for public projects based on decreasing financial thresholds. This framework has now been consolidated by the Code of Public Contracts of 2023, by which BIM is required as a mandatory standard for all public tenders exceeding one million euros starting from January 2025 (Koka *et al.*, 2025).

Despite this advancement, the practical utility of BIM remains largely confined to the office environment. The information embedded within a BIM model is typically created, reviewed, and updated by engineers and project managers working at desktop workstations, using software platforms that require significant technical expertise and hardware resources. On the construction site, where decisions are made in real time and conditions change daily, the information gap can potentially have immediate and severe consequences. Unfortunately, BIM data is rarely accessible in a direct and usable form on the sites. Site engineers, environmental consultants, safety inspectors, and field technicians must either carry printed extracts of relevant data, rely on verbal communication from office-based colleagues, or navigate complex software interfaces on devices not designed for field conditions. In all cases, the richness of the BIM model fails to reach the professional who needs it most, at the exact time and location where it is most relevant. To address this limitation, Augmented Reality (AR) technology has been successfully proposed as a visual bridge between digital models and physical construction environments. It has been demonstrated that complex multidimensional BIM data can be accurately superimposed onto real-world sites using standard portable devices, and these systems can operate effectively even in the absence of stable internet connectivity (Pan and Isnaeni, 2024).

This disconnect is particularly critical in geotechnical and environmental site investigation contexts. Subsurface investigation campaigns generate large volumes of laboratory data

including soil contamination measurements, stratigraphic profiles, groundwater analyses. These data are typically archived in spreadsheets and technical reports that bear no direct spatial relationship to the physical ground underfoot. When a field engineer stands at a specific location on a contaminated site, there is no immediate way to know what contamination data was collected at that precise point, at what depth, or whether the measured concentrations exceed regulatory limits. This information exists within the project's BIM model and associated databases, but the path from that data to the engineer's hands remains long and indirect.

The site of Città della Salute in Turin, Italy, provides a concrete and representative context for this problem. A large-scale public healthcare infrastructure project on a formerly industrialized brownfield site, has been the subject of extensive geotechnical and environmental investigation, generating a comprehensive dataset of subsurface contamination measurements across more than ninety borehole locations. This data, structured according to Italian regulatory standards under D.Lgs 152/06, defines the legal thresholds for soil contamination and forms the basis for remediation decisions that directly affect the safety and viability of the construction program. Making this data accessible, readable, and spatially meaningful to professionals working on the physical site represents both a practical necessity and a research opportunity.

1.2 Research objectives and questions

The primary objective of this research is the development of a mobile Augmented Reality (AR) application. This tool is designed so that geotechnical data can be delivered from a Building Information Model (BIM) directly to professionals on contaminated construction sites. To achieve this, four specific goals are pursued. First, a Revit-authored BIM model is integrated into a mobile AR environment at real-world scale. Second, a data pipeline is developed so that laboratory contamination measurements can be connected to individual boreholes within the AR model. Third, GPS-based spatial anchoring is implemented so the virtual model is accurately aligned with the physical site. Finally, a field-appropriate user interface is designed so that complex environmental data can be presented for immediate on-site decision-making.

The emergence of Augmented Reality as a mature technology offers a compelling solution to this challenge. Unlike Virtual Reality, spatial awareness is preserved because digital information is simply overlaid onto the real world. Furthermore, as ARCore-compatible smartphones become widespread, this capability is no longer confined to specialized hardware. Instead, it is readily available on everyday devices that are already carried across construction sites. Therefore, a practical bridge is formed between the BIM data and the physical reality of the site. Beyond the construction phase, these same principles can be applied to facility management. A field-

accessible AR interface is highly relevant for building managers, environmental officers, and maintenance engineers. Through this technology, subsurface data can continue to be interpreted long after the construction program is completed.

To guide this investigation, three primary research questions are addressed. What is the added value of contextualizing subsurface BIM and tabular contamination data through on-site Augmented Reality? And whether BIM-derived borehole geometry and associated environmental records can be integrated into a mobile AR application without data loss or mismatch. This is assessed through entity coverage, parameter coverage, and record-level agreement.

The second question asks what accuracy can be achieved outdoors when only the built-in sensors on a consumer device are utilized. This is evaluated by analyzing horizontal accuracy at the time of placement, the dispersion of positions across repeated trials, and the repeatability of the model's orientation.

The third question asks whether the complete dataset can be delivered on unmodified consumer hardware, and at what cost to performance. This is measured by the sustained frame rate, the time required from launch to placement, and the application's ability to function without a network connection.

These questions and their implications are discussed in Chapter 6. Finally, it should be noted that formal evaluation with intended users was outside the scope of the available time. Consequently, the practical assessment of whether this interface supports decision-making in the field is identified as future work.

1.3 Scope and limitations

This study focuses specifically on the development and technical evaluation of an AR application for Android smartphones, applied to the geotechnical investigation data of the Città della Salute site in Turin. By delivering this subsurface data directly to the field, building site management is significantly impacted. Specifically, environmental compliance is streamlined, and on-site decision-making is vastly improved. Field engineers are empowered to visualize hidden contamination hazards in real time, whereby communication delays between the office and the physical construction site are effectively reduced. To define the scope of this work, several specific boundaries are established.

In terms of platform, the application targets Android devices compatible with Google ARCore, reflecting the prevalence of Android smartphones on Italian construction sites and the maturity of ARCore as a mobile AR framework. iOS development falls outside the scope of this study, though the underlying Unity-based architecture would support cross-platform extension in future work.

In terms of data, the application visualizes soil contamination measurements collected during the site investigation campaign, structured according to the regulatory thresholds defined by D.Lgs 152/06. The dataset covers fifty-seven monitored contaminants across more than seven hundred sample records from approximately eighty-one borehole locations. Real-time data updates from ongoing laboratory analysis are not supported and the application operates on a static dataset loaded at startup, appropriate for field consultation but not for live monitoring workflows.

In terms of geometry, the three-dimensional model is limited to the Mass and Mass Floor elements exported from the Revit BIM model, representing the borehole investigation cells and their associated depth intervals. The full architectural and structural BIM model of the Città della Salute project is not included, as the focus of this study is the geotechnical investigation layer rather than the building design.

In terms of spatial positioning, the application implements GPS-based georeferencing using a one anchor point and the orientation methodology with consumer-grade smartphone GPS. The positional accuracy achievable with this approach is constrained by the inherent limitations of smartphone GPS, typically three to five meters under open-sky conditions, which is sufficient for borehole-level navigation across a site of this scale but not for millimeter-precise alignment.

Several limitations follow from these established boundaries, which must be acknowledged to properly contextualize the research findings. First, the application was developed and tested exclusively on a single Android device. Consequently, the performance variability across different hardware specifications and processing capacities remains unexplored. Furthermore, a systematic usability evaluation with intended field professionals has not yet been conducted, meaning that the interface's practical efficiency under real-world site conditions is not fully proven. Second, while the GPS anchoring was calibrated using survey coordinates from the site dataset, it has not yet been validated through physical field measurements at the actual site. Therefore, potential environmental interferences, such as signal obstruction from heavy construction equipment, are not accounted for in this study. Finally, a simplified wireframe visualization approach was chosen for the borehole model. While this method is technically functional and computationally lightweight, the full visual richness of the original BIM model is not captured. This simplification might potentially affect the user's depth perception and spatial understanding when viewing the augmented environment.

1.4 Thesis structure overview

This thesis is organized into seven chapters.

Chapter 1 introduces the problems addressed by this research, the objectives and research questions, and the scope and limitations of the work.

Chapter 2 reviews the relevant literature on GeoBIM, Augmented Reality in construction, and mobile AR positioning, and identifies the gap this thesis addresses.

Chapter 3 presents the case study site and its industrial history and situates this work within the broader research chain at Politecnico di Torino.

Chapter 4 describes the methodology behind GeoBIM Viewer, covering the tools used, the system architecture, and its three main layers: AR visualization, the data pipeline, and geospatial positioning.

Chapter 5 presents the results, including a demonstration of the application, an evaluation of its data accuracy.

Chapter 6 discusses the significance of these results and their contribution to the field and the limitations encountered.

Chapter 7 concludes the thesis and outlines directions for future work.

Chapter 2 Literature Review

Building a mobile Augmented Reality application that makes subsurface data directly accessible to professionals on construction sites is multidisciplinary subject. It sits at the intersection of three areas of research that have largely developed in parallel: the extension of Building Information Modelling into geotechnical and underground contexts, the growing use of Augmented Reality as a practical tool on construction sites, and the technical foundations that make mobile AR work reliably in outdoor environments. Each of these areas has matured considerably in recent years. This chapter maps the relevant work across these three threads, traces what has been achieved and what remains unresolved, and closes by identifying the gap that the present research steps into.

2.1 GeoBIM for Subsurface and Geotechnical Data

Building Information Modelling, commonly known as BIM, is a digital approach to representing a building or infrastructure project. Instead of working with separate drawings and documents, BIM brings all project information into a single shared model that every member of the project team can access and update. This is reflected in the formal definition provided by ISO 19650-1 (2018), which describes BIM as the use of a shared digital representation of a built asset to facilitate design, construction and operation processes. In practice, however, what BIM means to an organization tends to shift as its experience grows. (Borkowski, 2023) observes that this evolution typically moves through three stages, from treating BIM as a technology for building three-dimensional models, to using it as a process for working more efficiently to finally understanding it as a methodology for delivering better outcomes across the full project lifecycle.

Despite its growing adoption, BIM has often been implemented in isolation from the broader information systems that surround a construction project. (Raitviir and Lill ,2024) argue that the industry's continued reliance on fragmented workflows results in significant information loss, inefficiencies, and errors across the project lifecycle. They propose that adopting open standards such as IFC and CityGML can bridge this gap, enabling continuous information flow from early planning through to demolition.

GeoBIM is the integration of Building Information Models with three-dimensional geospatial city models. BIM alone captures detailed information about a building but cannot handle city-scale geographic data, while GIS and 3D city models cover urban geography but lack the detailed building-level information that BIM provides. Bringing the two together enables applications that

neither can support independently, such as infrastructure design, detailed city analysis, building permit processing, and context-aware architectural design (Noardo et al., 2020). However, achieving this integration in practice is far from straightforward. The GeoBIM Benchmark 2019 tested available tools for georeferencing IFC models and converting data between IFC and CityGML and found significant gaps in both areas. The authors conclude that the difficulties do not stem solely from immature software but from a deeper misalignment between what the standards formally require and how those standards are implemented in available tools.

Geotechnical engineering is one field where the integration of BIM and GIS has been applied with convincing results. (Kassou et al. ,2025) demonstrate this through a bidirectional BIM-GIS workflow developed for the Tangier-Kenitra high-speed railway project in Morocco. Using Empirical Bayesian Kriging, they constructed three-dimensional subsoil models from pressuremeter test data, which were then dynamically linked to the BIM design environment through visual programming scripts. This connection allowed settlement calculations and the generative design of vertical drainage systems to update automatically as geotechnical inputs changed, eliminating the manual data re-entry that typically slows infrastructure design. The authors note that while the workflow significantly reduces errors and improves responsiveness to subsurface variability, updates still require re-running the visual programming scripts when new data becomes available, pointing to the need for more continuous and automated data integration in future work.

(Kim and Chung, 2026) push this approach further by proposing a full geotechnical characterization framework specifically designed to feed into GeoBIM applications. Working on a reclaimed coastal site, they apply Sequential Gaussian Simulation to high-density survey data to generate a high-resolution three-dimensional ground model with a 5×5×1 metre grid, which they term nanozonation. This level of detail allows the framework to assess multiple geohazards simultaneously, including settlement, liquefaction, and seismic amplification, producing hazard zonation maps that are directly embedded into GeoBIM for ground improvement planning and risk mitigation. The authors acknowledge however that the accuracy of the model depends heavily on the density and quality of the input survey data, and that geostatistical interpolation assumes spatial correlation in soil conditions that may not always hold in highly heterogeneous ground. The study nonetheless demonstrates that when sufficient data is available, GeoBIM can move beyond simple visualization and become a genuine decision-support tool for geotechnical design.

GeoBIM has also been applied to the management and monitoring of historic infrastructure. (Sanseverino and La Placa ,2026) present a semi-automated workflow for generating a GeoBIM model of the Naviglio water channel in Pavia, Italy, starting from nothing more than two-dimensional cartographic data and historical records. Using visual programming in Revit, the

workflow produces a georeferenced LOD 200 model in a matter of hours, which then serves as a digital ecosystem for linking inspection and monitoring data through unique IFC identifiers. The authors highlight that automation through visual programming makes the approach transferable and reproducible, not tied to a specific BIM authoring tool. While the model is limited to LOD 200 and the geometry cannot be directly edited once generated, it provides a flexible base that can be progressively enriched as more data becomes available. This case demonstrates that GeoBIM workflows can be designed to be lightweight and accessible, not only applicable to large-scale infrastructure projects with abundant survey data.

The application of BIM-GIS integration extends beyond geotechnical design into the broader management of underground urban infrastructure. (Sivakumar et al. ,2025) demonstrate this through a framework applied to stormwater drainage networks in Chennai, India, where traditional two-dimensional planning methods had proven inadequate for managing flood risks in a rapidly urbanizing environment. By combining BIM's detailed three-dimensional modelling of drainage components with GIS's geospatial context, the framework enabled improved visualization, simulation, and decision-making for infrastructure managers. The authors acknowledge however that reliance on historical two-dimensional data required simplifications during the modelling process, and that limited data availability affected the accuracy of the final model.

(Rajadurai and Vilventhan ,2024) take the visualization of underground infrastructure a step further by integrating BIM, GIS, and Augmented Reality into a single framework for managing multiple underground utilities in the field. Rather than limiting utility information to desktop environments, their framework delivers LOD 300 models with full geometric and geospatial attributes directly into the real-world environment through a mobile AR application, using OpenBIM data standards to eliminate information loss during integration. The authors acknowledge however that the framework relies on image-based markers to trigger augmentation, which becomes difficult to manage across larger areas, and suggest that future work could replace this approach with GPS or QR-based spatial anchoring. This study is particularly relevant to the present research as it demonstrates both the potential and the current limitations of delivering BIM-embedded subsurface data through mobile AR on construction sites.

The most direct precedent for the present research is provided by (Aschieri et al. ,2026), who evaluate four GeoBIM workflows for contaminated site management using real borehole and laboratory data from the PSRI redevelopment area in Turin-the exact site this thesis addresses. The study compares geometry-based geological modeling, geostatistical voxel approaches in GIS, and parametric BIM-native pipelines, ranking them through an Analytical Hierarchy Process against key performance indicators covering reliability, automation, interoperability, and

usability. The results reveal a fundamental trade-off: voxel-based GIS pipelines preserve spatial heterogeneity and perform best for remediation decision-support, while parametric BIM approaches score higher on automation and construction coordination. Across all four workflows however, interoperability emerges as the primary limiting factor, with georeferencing loss and weak environmental semantics reducing the reusability of outputs. Critically, none of the workflows evaluated address how the resulting models can be accessed and queried by professionals working directly on the site -a gap that the present research sets out to fill.

2.2 Augmented Reality in Construction and Site Management

Augmented Reality is a technology that overlays digital content onto the real world as perceived through a device camera or screen, allowing users to see virtual information superimposed on their physical surroundings. It sits on a spectrum between the fully real environment and fully virtual reality, with the key distinction being that the user remains grounded in the physical world rather than being immersed in a simulated one. This characteristic makes AR particularly suited to construction and operations contexts, where maintaining awareness of the physical environment is essential. (Amin et al. ,2023), through a systematic review of BIM-AR platforms, identify six key functions that any such system must address: positioning, interaction, visualisation, collaboration, automation, and integration. Of these, positioning is the most fundamental, as the accuracy of spatial anchoring directly affects the quality of all other functions, while automation remains the least mature, still considered to be in its early stages.

The application of AR across the Architecture, Engineering and Construction industry has grown substantially over the past decade, spanning all phases of the project lifecycle. (Schiavi et al. ,2022), in a systematic review of over 400 publications, identify construction management, quality control, progress monitoring, maintenance, and safety training as the principal areas of application. The construction phase accounts for the largest share of research, with mobile handheld devices emerging as the dominant delivery platform. The release of ARCore and ARKit frameworks in 2017 marked a turning point, enabling markerless AR on standard smartphones and significantly broadening both the research and practical adoption of the technology.

A concrete example of mobile BIM-AR deployed on a construction site is presented by (Hsieh et al. ,2023), who develops a progress management system using SLAM-based positioning on iOS handheld devices. The system overlays a color-coded BIM model onto the real construction environment, allowing site engineers to compare planned and actual progress in real time. However, the authors identify significant limitations in the positioning component: the monotonous textures and rapidly changing conditions typical of construction sites degrade SLAM

accuracy, requiring careful path planning and frequent rescanning. This highlights a persistent challenge in on-site AR - achieving reliable spatial anchoring under real field conditions.

(Barrile et al. ,2024) explored the integration of BIM and GIS with VR/AR/MR technologies for urban planning and building management. Their study demonstrated how Cesium, Unity, and HoloLens can support the geospatial visualization of BIM models, highlighting the potential of these technologies for sustainable urban development and digital twin applications. The results suggest that combining BIM, GIS, and immersive visualization can support more informed planning, inspection, and management decisions and could contribute to the development of digital twins for sustainable urban development (Figure 1); however, the authors also acknowledge that the proposed system remains primarily focused on visualization and information access rather than advanced analytical functions.

Figure 1 BIM and AR-based visualization of the building within its urban context (Barrile et al., 2024).

While the reviewed studies demonstrate that AR has become a genuinely useful tool across the construction lifecycle, the evidence consistently points to spatial anchoring as the critical unsolved challenge - a limitation that becomes particularly acute when the information to be visualized lies not on the surface of the built environment but beneath it, in the subsurface conditions that are invisible to the naked eye.

2.3 AR for Subsurface and Invisible Infrastructure Visualization

Underground infrastructure presents a peculiar problem: the people who most need information about it is standing on top of it, unable to see it. Pipelines, utilities, and subsurface conditions are invisible from the surface, yet the decisions that depend on knowing their exact location and condition must be made in the field, often under time pressure. (Li et al. ,2023) point out that while GIS and BIM have become standard tools for managing this kind of information, they keep it locked inside virtual environments on a screen, far removed from the physical reality of the

site. Augmented Reality offers a direct response to this problem, bringing digital representations of what lies underground into the real world by overlaying them onto the environment as seen through a mobile device. Their study demonstrates this through a visualization framework (Figure 2) for urban gas pipe networks that uses a smartphone linked to a lightweight RTK device to achieve a positioning accuracy of 0.028 metres in real field conditions.

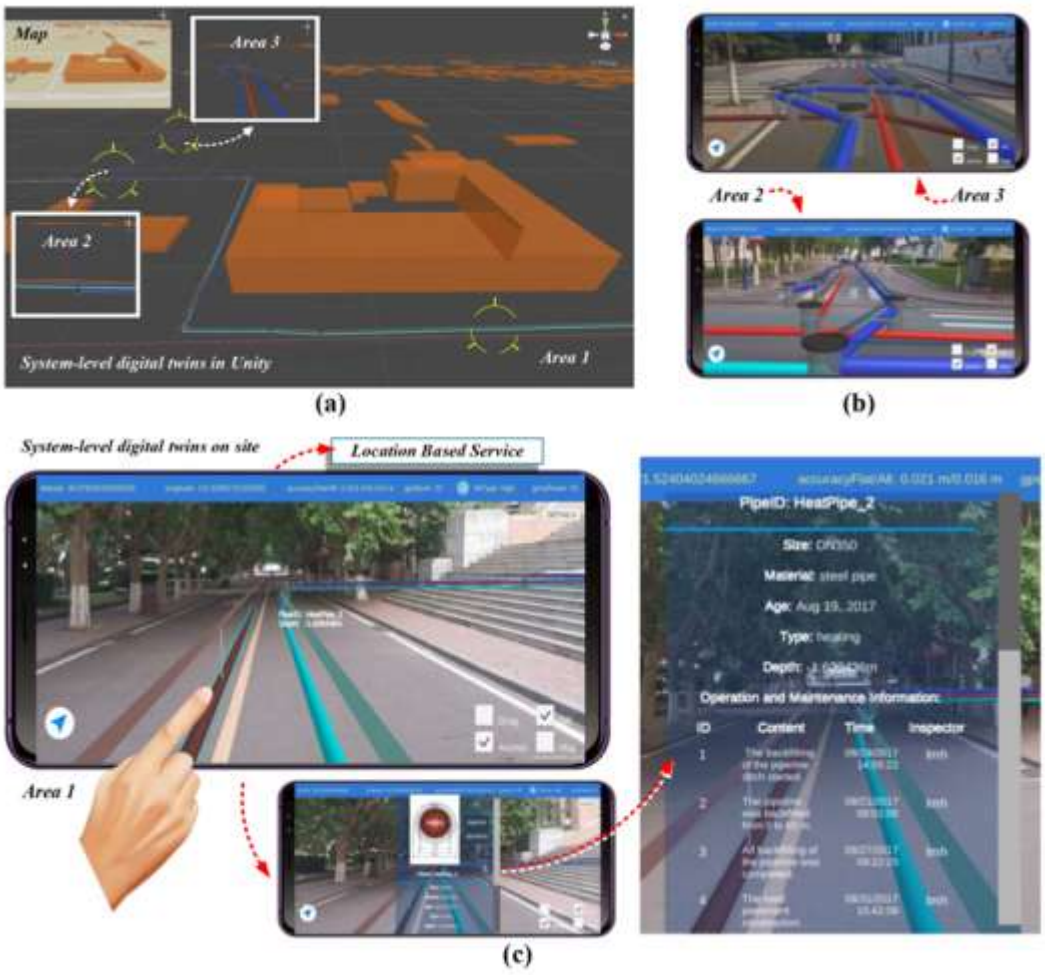

Figure 2 Visualizing underground utility networks using a smartphone and lightweight RTK device (Li et al., 2023).

(Rajadurai and Vilventhan ,2024) build on this direction by integrating BIM, GIS and AR through OpenBIM data standards, delivering LOD 300 utility models with full geometric, non-geometric and geospatial attributes directly into the field through a mobile application. Both frameworks confirm that field delivery of subsurface data through mobile AR is technically achievable, though both rely on specific anchoring mechanisms that introduce practical constraints for deployment across larger or more dynamic site areas.

While most AR studies for underground infrastructure focus on technical development, Khorrami (Shad et al. ,2024) take a different approach by testing two AR prototypes directly with 31

excavation professionals in New Zealand, providing a rare industry perspective on how AR is perceived and used in the field. The study presented participants with the same AR application delivered through two different devices — an optical see-through headset and a video see-through tablet — and asked them to evaluate both in the context of real excavation tasks. The findings reveal that participants consistently preferred AR over traditional paper-based drawings, and a quantitative analysis using the NASA Task Load Index showed a statistically significant reduction in perceived workload for specific tasks such as utility location. Among the two devices, participants strongly favored the tablet, citing familiarity and ease of use as the primary reasons, while a smaller group acknowledged the long-term potential of hands-free headsets as the technology matures. The study also identifies a critical technical gap that professionals themselves flagged as the most significant barrier: the lack of accurate vertical information about underground utilities. In current practice, operators can often locate a utility in plan but cannot reliably determine its depth, and this limitation persisted even with AR assistance. This finding is significant because depth accuracy — knowing precisely how far below the surface something lies — is one of the most demanding requirements for any AR system operating in a subsurface context, and one that the existing literature has not yet fully resolved.

The studies discussed so far assume that subsurface data already exists in a digital model and that the challenge is simply one of delivering it to the field. (Safaan et al. ,2025) push the boundary further by combining Ground Penetrating Radar with deep learning and Augmented Reality to detect and visualize subsurface utilities that have not yet been modelled. Their framework automates the interpretation of GPR data using a Mask R-CNN deep learning model, converts the detected utilities into a spatial point cloud, and visualizes the results through a HoloLens mixed reality application, creating a system that moves from raw sensor data to on-site AR visualization in a largely automated pipeline. The authors acknowledge however that the system is constrained by the operating frequencies of the GPR antennas, which limit detectable depth and object size, and that false positive detections remain a challenge in complex subsurface conditions. Despite these limitations, the study represents the most technically ambitious direction in this space and points toward a future where AR does not merely display what is already known but actively helps reveal what is not.

Of the six key functions that define a BIM-AR platform, positioning is the most fundamental and the most technically demanding (Amin et al., 2023). (Skabek et al. ,2026), who benchmark four AR localization technologies — Vuforia, Immersal, MultiSet, and the ARCore Geospatial API — across indoor and outdoor environments, observe that indoor positioning relies on visual SLAM, fiducial markers, and Bluetooth beacons, while outdoor positioning depends primarily on GPS, often enhanced through sensor fusion or geospatial visual data. Their findings reveal meaningful differences between platforms in initialization speed, accuracy, battery consumption, and tolerance to adverse conditions. For outdoor positioning specifically, the ARCore Geospatial API,

which fuses GPS with Google's Visual Positioning System, emerges as a practical solution for large-scale outdoor AR, enabling seamless transitions between indoor and outdoor anchoring. Battery consumption, however, remains consistently high across all tested platforms — a practical concern for field deployment where devices may be in continuous use.

Exploring alternative markerless approaches, (Mužík and Běloch , 2026) evaluate the use of Visual Positioning Systems (VPS) and ARCore/ARKit to anchor virtual reconstructions in urban environments without relying on GNSS (Figure 3). Their findings reveal that consumer smartphone localization depends heavily on ambient lighting and available environmental feature points, achieving over 70% localization success during daylight but failing entirely at night. This underscores a vital research imperative: determining exactly what spatial accuracy, dispersion, and placement repeatability can be achieved outdoors when utilizing only the built-in sensors of a consumer device under real-world conditions.

Figure 3 Mobile AR application for marker less outdoor visualization (Mužík and Běloch, 2026).

One approach to outdoor AR positioning that has gained traction in construction contexts is the integration of GPS location data with computer vision-based object tracking. Htet et al. (2026) develop an AR framework for real-time information retrieval about dynamic objects on construction sites, using GPS coordinates as the spatial anchor that links what the device camera sees to what is stored in the digital model. By incorporating object motion direction as an additional matching constraint, the framework achieves stable and reliable associations between image entities and their GPS records in real time, extending the approach beyond static BIM elements to cover moving workers and equipment. The authors acknowledge however that real-time implementation remains affected by sensor inaccuracies and synchronization delays

between visual and location data, particularly in crowded or rapidly changing site conditions, pointing to the need for improved sensor fusion techniques to enhance reliability in practical deployment.

2.4 Research Gap and Positioning of This Work

The literature reviewed across this chapter traces three converging lines of development. In the domain of GeoBIM, significant progress has been made in constructing three-dimensional models of subsurface conditions, integrating geostatistical interpolation with BIM and GIS workflows to represent stratigraphy and contamination data with increasing accuracy and detail. In the domain of Augmented Reality for construction, technology has matured from experimental prototypes to practical field tools, with established frameworks for positioning, visualization, and interaction across the project lifecycle. In the specific area of AR for subsurface and underground infrastructure, a growing body of work has demonstrated that making invisible conditions visible through mobile AR is not only technically feasible but genuinely valued by field professionals.

Yet, these lines of development have not converged into a solution capable of supporting active site remediation. The research gap is not merely that few such systems exist, but that prior studies do not jointly provide the functional dimensions necessary for on-the-ground decision-making. Specifically, existing literature lacks a framework that simultaneously integrates:

1. **Georeferenced subsurface BIM:** Moving beyond standard utility geometry (pipes and cables) to spatially anchor complex, interpolated environmental data.
2. **Record-level contamination querying:** Permitting field professionals to interrogate environmental data by specific borehole and depth interval.
3. **Automatic regulatory-threshold classification:** Providing an immediate, automated visual assessment of measured parameters against legal limits.
4. **Consumer-device outdoor AR:** Delivering this complex geometry and data pipeline entirely offline on unmodified, standard mobile hardware.

This research addresses that gap directly as a testable claim. Building on the GeoBIM model developed for the PSRI remediation site in Turin and the BIM reconstruction (Aschieri et al 2025), this work develops a prototype that actively implements all four of these dimensions. It translates information that has traditionally existed only on specialist desktop software into a field-ready application where laboratory contamination measurements are dynamically linked to individual borehole elements and automatically evaluated against regulatory thresholds.

However, while this work successfully implements the data pipeline, spatial anchoring, and hardware deployment required to close this gap, certain operational aspects remain unvalidated. Because testing was restricted to a proxy location, the absolute positioning accuracy achievable under the specific environmental constraints of the physical PSRI site is not yet proven. Furthermore, while the prototype successfully delivers classified subsurface records to a consumer device, its practical usefulness, usability, and impact on daily site management remain theoretical until evaluated through formal testing with end-users in the field.

Chapter 3 Case Study

3.1 Site Description and Historical Background

The site subject to this research is in the Lingotto district of Turin, Italy, in the northwestern portion of a former industrial area bounded to the north by the Lingotto Congress Centre, to the east by Via Nizza, to the south by the Piedmont Regional headquarters, and to the west by the Oval Lingotto. The area, designated as Lot 1 within the broader PSRI masterplan, extends across approximately 118,000 square meters and reaches a maximum excavation depth of 7.5 meters below ground level, between elevations 234 m and 227 m above sea level. The location of LOT 1 is shown in Figure 4.

Figure 4 The location of the current study (image taken from Google earth)

The site is currently being transformed into the Parco della Salute, della Ricerca e dell'Innovazione, a large-scale urban regeneration project that will host a new hospital, research facilities, and innovation infrastructure for the city of Turin and the Piedmont region. The main devision of PSRI masterplan is demonstrated in Figure 5.

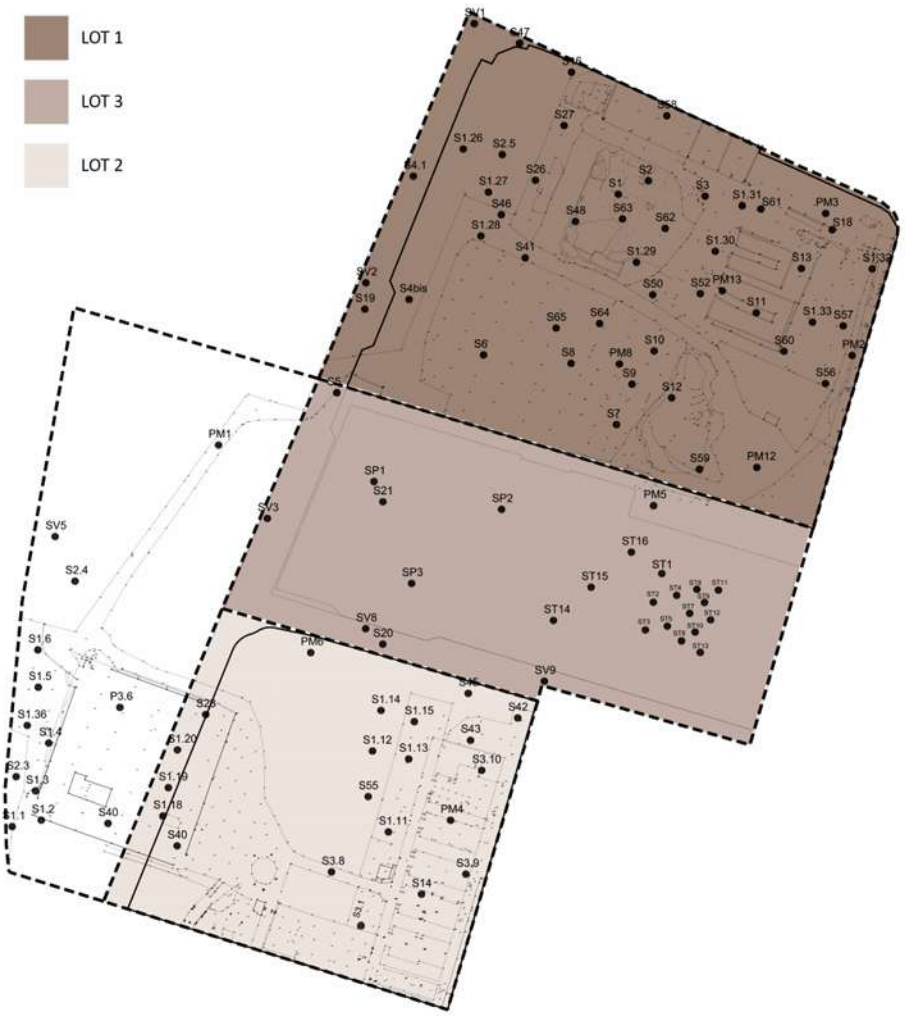

Figure 5 The location of three lots of parco della salute project with Boreholes positioning (Fonsati et al., 2023))

The complexity of this transformation is compounded by the significant environmental legacy of the site, which required extensive geotechnical characterization and remediation before construction could begin.

The industrial history of the site spans nearly two centuries and has left a profound imprint on the composition of its subsurface. The earliest documented productive use dates to 1837–1838, when a chemical plant to produce artificial fertilizers was established on the land. This configuration remained largely unchanged through the first half of the twentieth century, until a petroleum and gasoline storage facility owned by SIAP was added in 1907.

In the post-World War II period, Fiat-Gestione Lubrificanti reorganized the area into two functional zones. One was a petroleum storage and distribution sector, and the other was lubricants processing sector. This layout persisted until 1969, when several structures including pumping stations, lubricant storage buildings, paint shops, and tyre warehouses were demolished.

The most significant industrial phase began in 1976 with the construction of the Fiat Aviazione S.p.A. complex, which occupied the full site area and was dedicated to aeronautical component manufacturing. Operations included surface treatment processes such as galvanic plating, thermal treatment, and painting. In 1982, Fiat Avio expanded into the northwestern portion of the site, relocating painting facilities and plasma deposition units and decommissioning the site's thermal power plant.

From 2003 onwards, the site underwent progressive demolition, with the removal of buildings and partial stripping of surface pavements. Historical records also document excavation, remediation, and backfilling operations in portions of the northern sector of the site.

This layered industrial history produced a highly heterogeneous subsurface, with contamination from multiple source types distributed unevenly across both depth and plan surface. Environmental investigations confirmed the presence of heavy and light hydrocarbons, chlorinated organic solvents, and heavy metals including lead, copper, chromium, nickel, tin, and arsenic, at concentrations exceeding the regulatory thresholds established by Italian environmental legislation (D.Lgs. 152/2006) and, for chromium and nickel, the natural background values defined specifically for the Piedmont region by the City of Turin's Determination n. 161/2019.

3.2 The Research Chain

The environmental and geotechnical complexity of the PSRI site has become the subject of a progressive sequence of research studies at the Politecnico di Torino. Each building upon previous findings to advance the digital representation and practical usability of subsurface data.

To effectively bridge the gap between subsurface geological data and above-ground design, recent research by Aschieri et al. (2026) presents a comprehensive methodological framework for GeoBIM integration. Because this study establishes the core technical foundation for the application developed in this thesis, it is crucial to examine its approach in detail. The researchers evaluated multiple workflows to transform raw environmental and geotechnical data into a unified, navigable repository tailored for cross-domain coordination. Their technical

pipeline, illustrated in Figure 6, is structured into three sequential phases: initial data pre-processing, workflow implementation, and final model federation.

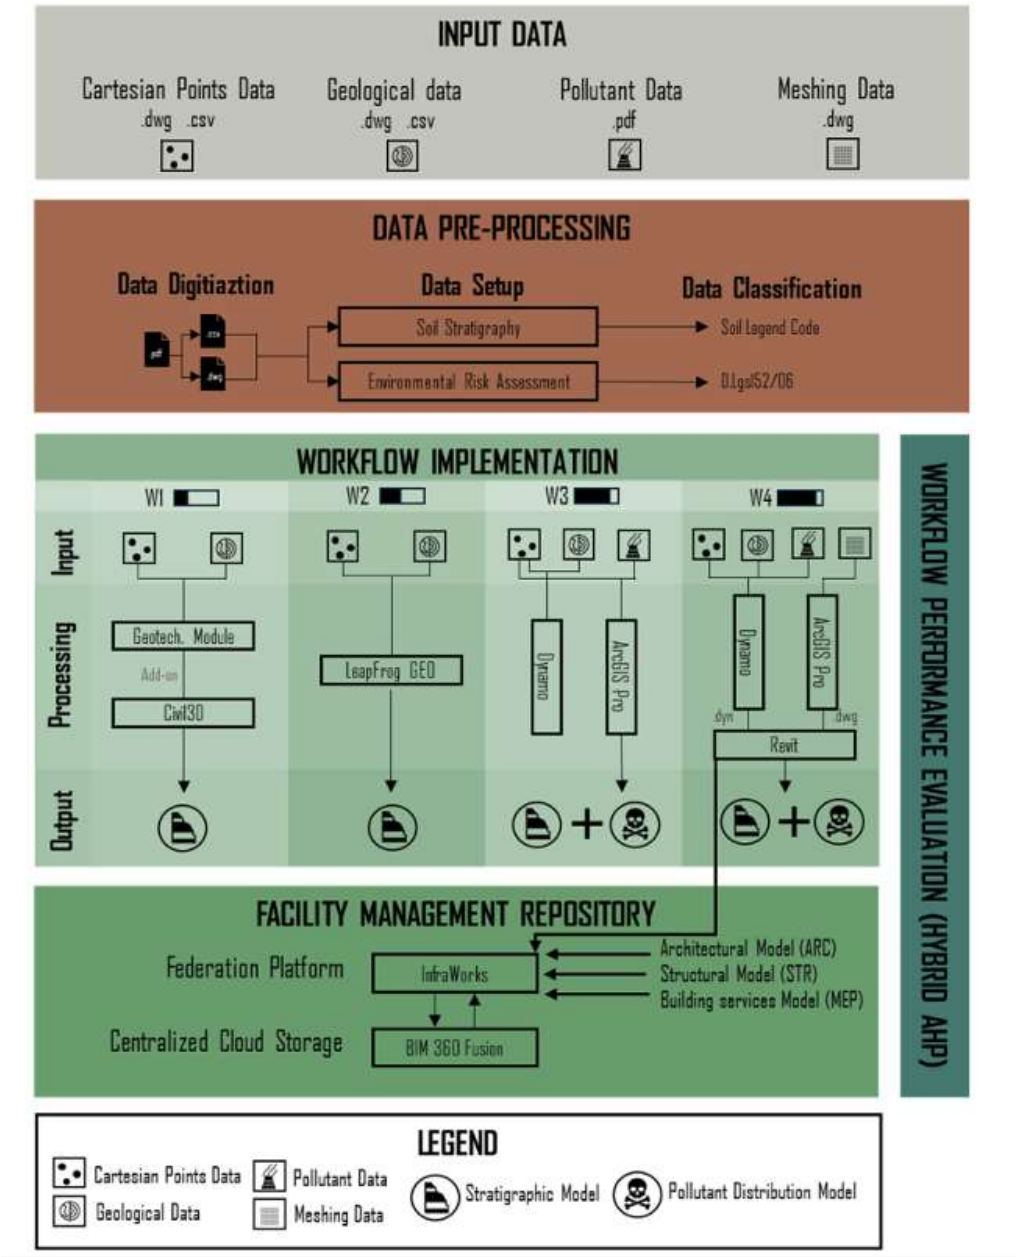

Figure 6 General methodological framework for GeoBIM modelling and evaluation, adapted from (Aschieri et al., 2026)

The first phase focused on standardizing heterogeneous source materials—such as borehole logs, CAD drawings, and spreadsheets—into computable and fully interoperable formats. This involved digitizing raw data, organizing it into relational tables for both stratigraphy and pollutant concentrations, and classifying the information against regulatory environmental thresholds.

Following data preparation, the study implemented and evaluated four distinct GeoBIM modeling workflows.

A critical distinction within these tested workflows is the difference between continuous modeling and voxel modeling. The initial workflows utilized continuous modeling paradigms. By employing explicit CAD-BIM tools and implicit geological interpolation, these approaches generated continuous stratigraphic boundaries and closed volumetric units, as seen in Figure 7. While continuous models excel at representing cohesive geological layers and maintaining structural continuity between boreholes, they provide limited support for volumetric contaminant depiction and struggle to represent highly localized environmental variables.

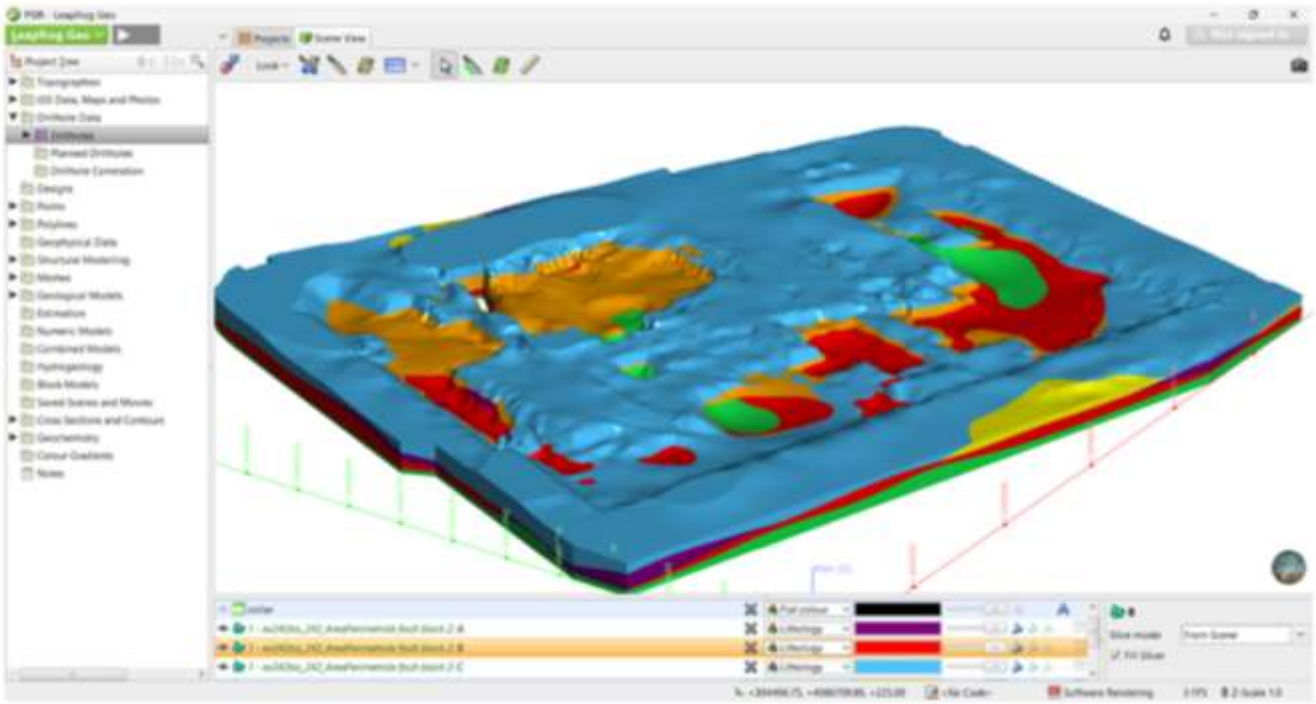

Figure 7 The continuous geological surfaces (Aschieri et al., 2026)

To address the limitations of continuous surfaces, the study also explored voxel-based and discretized grid modeling. Voxel models discretize the subsurface into a three-dimensional grid of volumetric cells, where each individual block acts as an independent data container storing spatial, lithological, and geochemical attributes. As demonstrated in Figure 8, workflows employing these paradigms translate subsurface data into a discrete, voxel-like grid of parametric elements to map contaminant distributions. While the Revit–Dynamo workflow offered strong performance in automation, interoperability, and modeling time, it was acknowledged to have

lower geometric precision than the original Leapfrog model and limited simultaneous multi-parameter visualization.

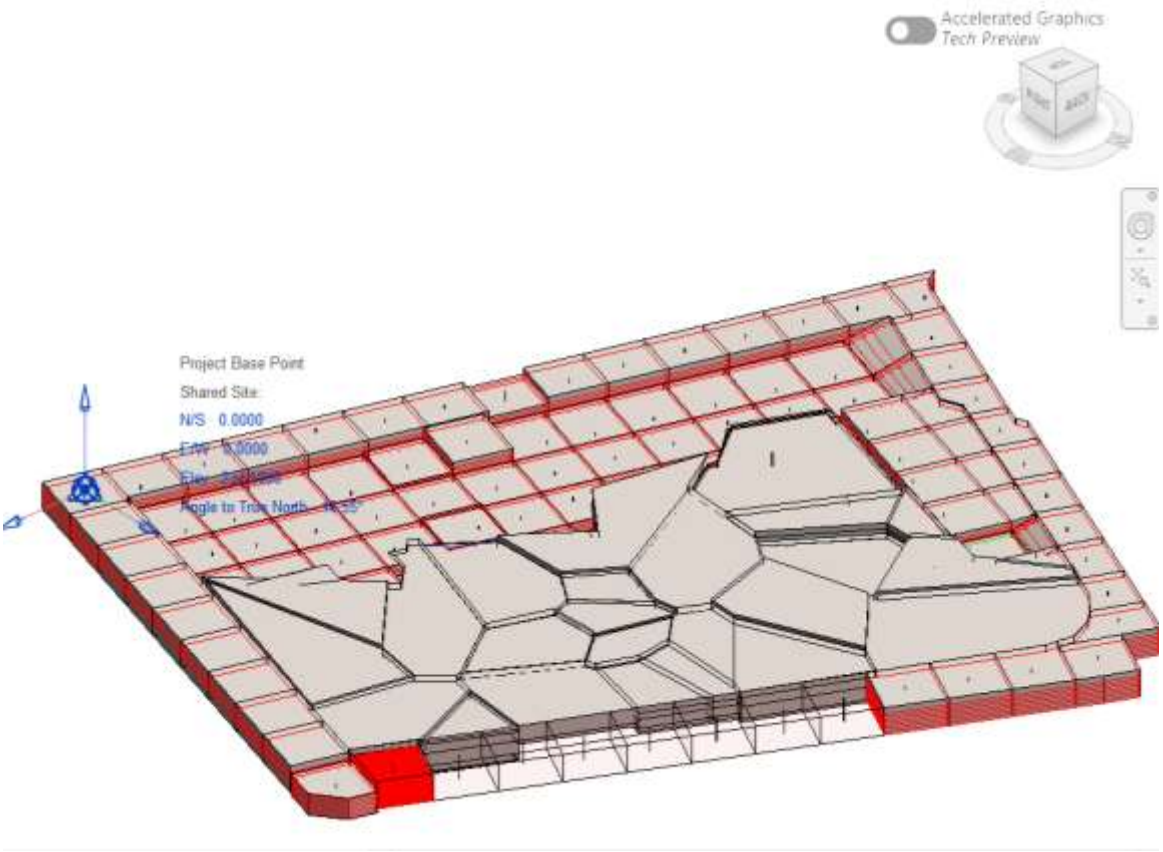

Figure 8 The discretized/voxel-based model of LOT 1 using Revit + Dynamo

The present research, which is described in a simple way in Figure 9, represents the next step in this chain. Taking the BIM model developed by Khosravi as its input, this work develops a mobile Augmented Reality application that delivers the subsurface data directly to field professionals at the physical site. By integrating GPS-based spatial anchoring, a data pipeline connecting laboratory contamination measurements to individual borehole elements, and a field-appropriate user interface, the application aims to close the gap between the digital model and the physical reality of the construction site, making invisible subsurface conditions visible, and actionable in the field.

Figure 9 Simplified procedure of the present research

Chapter 4 Methodology

4.1 Overview of the Approach

The Geo BIM Viewer -an Augmented Reality (AR) application- is designed to be used on construction sites to provide geotechnical information visually, overlaying a 3D model of the site on the real world through the device camera.

The application allows users to interact with individual borehole cells within the site model, accessing contamination data, depth information, and downloadable reports directly in the field. In other words, the application brings that information onto the site. A user standing on the ground can hold up a phone and see the soil model in place beneath them, positioned at its real coordinates rather than simply displayed on the screen. Selecting a borehole opens the measurements recorded there, arranged by depth and by contaminant, with each value compared automatically against its regulatory limit. The original field documentation can be downloaded from the application. Simply said, the existing information becomes reachable during the site visit with a smart phone.

The application also provides a second mode for use away from the site. In this mode, the same model is displayed within a three-dimensional representation of the real surroundings, navigated freely with touch controls. This makes it possible to review the model, prepare for a visit, or present the work without being physically present, and it also allows the whole site to be seen at once from the top, which is not possible from within it.

The application itself was built and organized into three layers, one presenting the model and handling interaction, one placing it correctly in the real world, and one loading and delivering the borehole data. These layers are distributed across two modes, described in the sections that follow.Figure 10 shows an overview of how GeoBIM viewer application works.

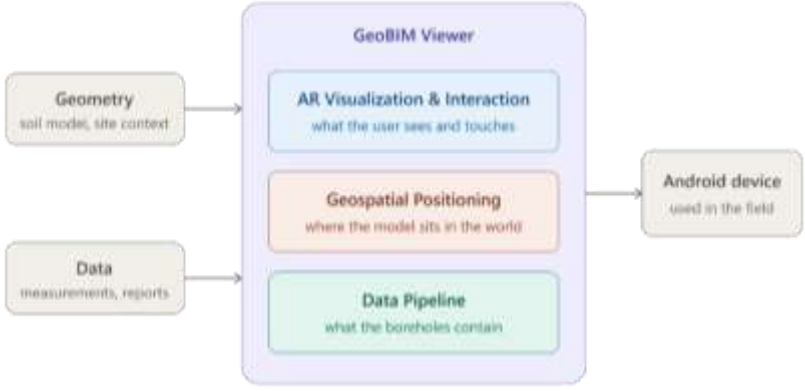

Figure 10 An overview of the GeoBIM viewer Application

4.2 Inputs, Environment and Tools

4.2.1 Input data

The input assets utilized in this project comprise two main categories: geometric models representing the physical and subsurface environment, and analytical data records provided in tabular and document formats (CSV and PDF).

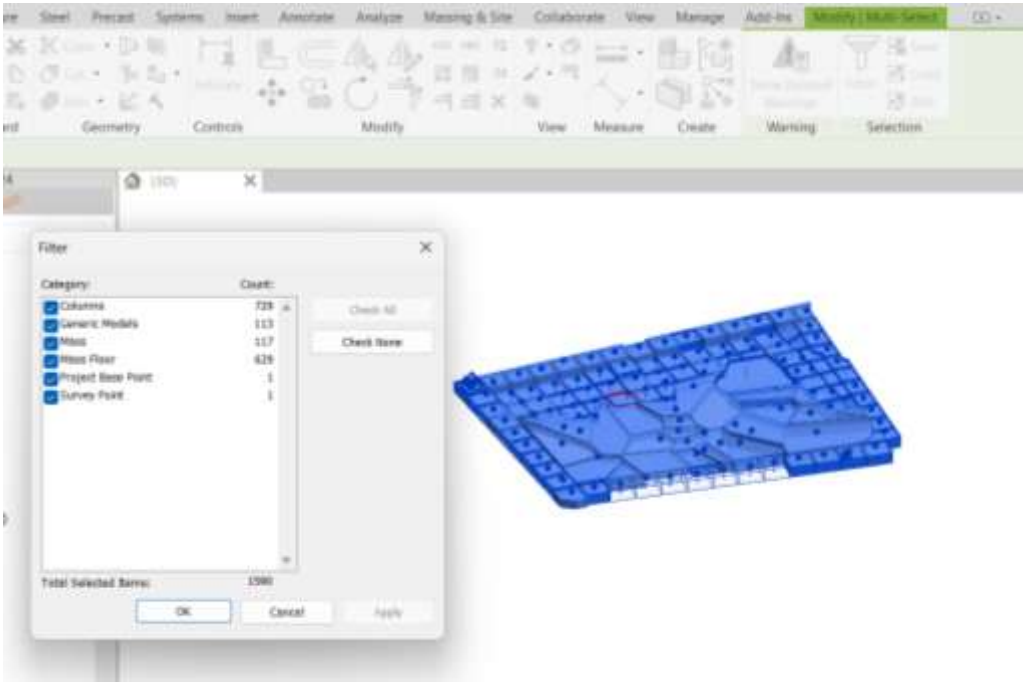

Figure 11 Revit selection filter showing the category breakdown and component counts of the subsurface model

As shown in Figure 11, the main geometric subsurface asset was generated in Autodesk Revit as an assembly of 1,590 discrete elements structured to organize spatial and contamination data across the site. The layout is partitioned into 117 conceptual mass volumes, each representing a grid cell that encloses a single borehole. These masses are subdivided by 629 mass floors to define vertical depth intervals. The boreholes are explicitly modeled using 729 column elements, which represent the vertical sampling intervals along each borehole depth. Each column segment stores the laboratory-measured contamination parameters for that specific depth. In addition, 113 generic model elements instantiated as adaptive point families mark the surveyed collar positions on the surface, while the Project Base Point and Survey Point establish the internal origin and real-world georeferencing.

When exporting the subsurface model from Autodesk Revit via FBX format, only the raw geometric entities—such as masses, mass floors, and column meshes—are transferred, while

embedded BIM parameters are not retained. As a result, the exported file serves strictly as a spatial and volumetric framework, requiring the analytical contamination data and depth parameters to be linked back to the geometry within the application at runtime.

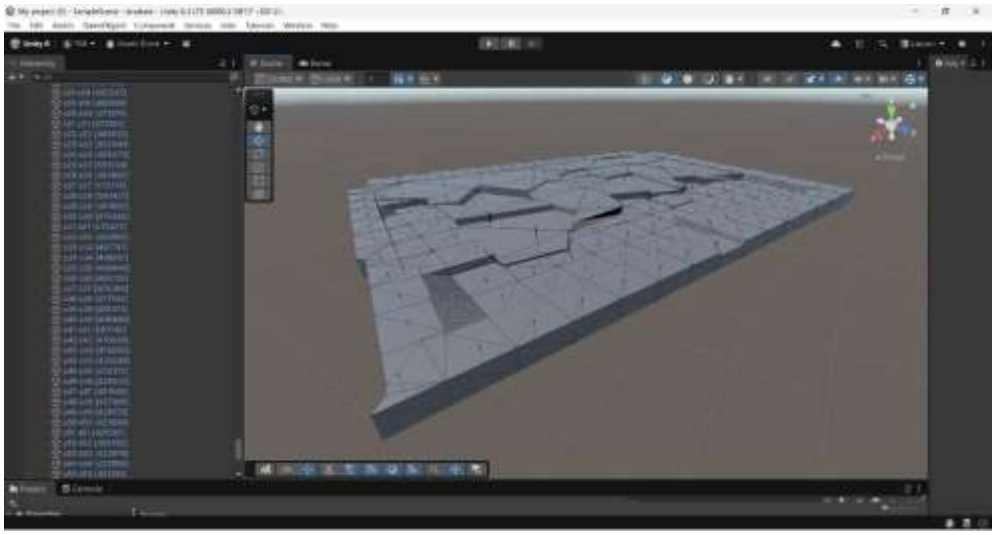

Figure 12 Subsurface geometry imported into the development environment

As shown in Figure 12, importing the model into the development environment preserves the object hierarchy and unique entity labels (such as individual borehole identifiers along the scene tree). While the internal BIM property sets are absent, this preserved structural tree and naming scheme provide the geometric handles necessary for the application's runtime scripts to associate each 3D element with its corresponding tabular measurement records.

The second geometric asset consists of the surrounding contextual environment exported from Autodesk InfraWorks as an OBJ file, functioning as a real-world visual reference to verify orientation and spatial positioning during on-site mode. Using the geographic selection tools within InfraWorks, the model was cropped around the exact site footprint to include existing surface terrain and surrounding urban structures. When integrated into Unity, this reference geometry provides observable landmarks that allow the user to visually corroborate the alignment between the digital model and physical surroundings in the field.

The non-geometric assets integrated into the application provide the analytical measurements, regulatory criteria, and archival documentation necessary for subsurface risk evaluation. These inputs consist of two structured data tables and a collection of borehole documentation files packaged directly within the application’s streaming assets directory.

The primary tabular dataset, shown in Figure 13, organizes the laboratory chemical analyses across the site. Originally derived from comprehensive laboratory reports, this table correlates each location identifier with its spatial coordinates (X, Y, Z), the sampled interval boundaries (Depth Top and Depth Base), and the quantitative concentrations of individual chemical analytes, such as heavy metals and polycyclic aromatic hydrocarbons (PAHs). The application parses this structured table dynamically to query the specific contamination type and concentration present at any given depth interval of a selected borehole.

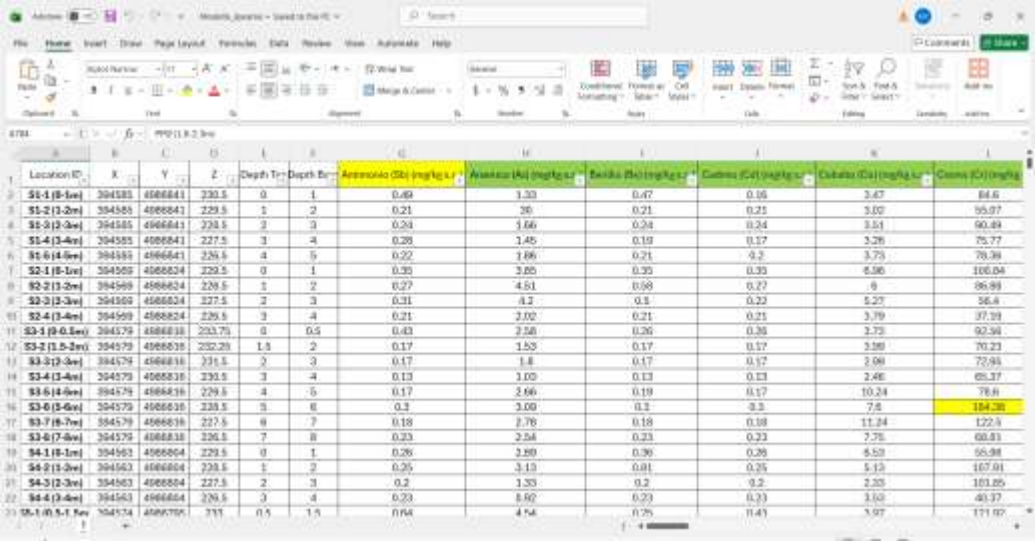

Figure 13 Structured measurement table (Modello_dynamo) mapping borehole locations, depth intervals, and contaminant concentrations

To evaluate environmental compliance on-site, a reference threshold table is loaded in parallel. This table pairs each chemical analyte with its permissible regulatory limit according to the applicable standards of D.Lgs. 152/06 (considering natural background values where defined. The application compares the measured concentration retrieved from the primary dataset against these reference limits at runtime, enabling automated identification of threshold exceedances directly within the mobile interface.

In addition to the structured tables parsed directly by the application code, official geotechnical and environmental reports are stored in PDF format inside the application's file directory to provide original documentation in the field. The first asset is the comprehensive laboratory report shown in Figure 14, which compiles the complete analytical results and statutory conformity indicators for the entire site into a single reference document accessible across all borehole cells. The second documentary asset comprises individual borehole stratigraphic reports, exemplified in Figure 15, which provide detailed soil profiles, technical logging parameters, and core sample photographs for specific drill locations. By embedding these files

directly within the project package, the interface allows users to review and download the official technical records during field operations.

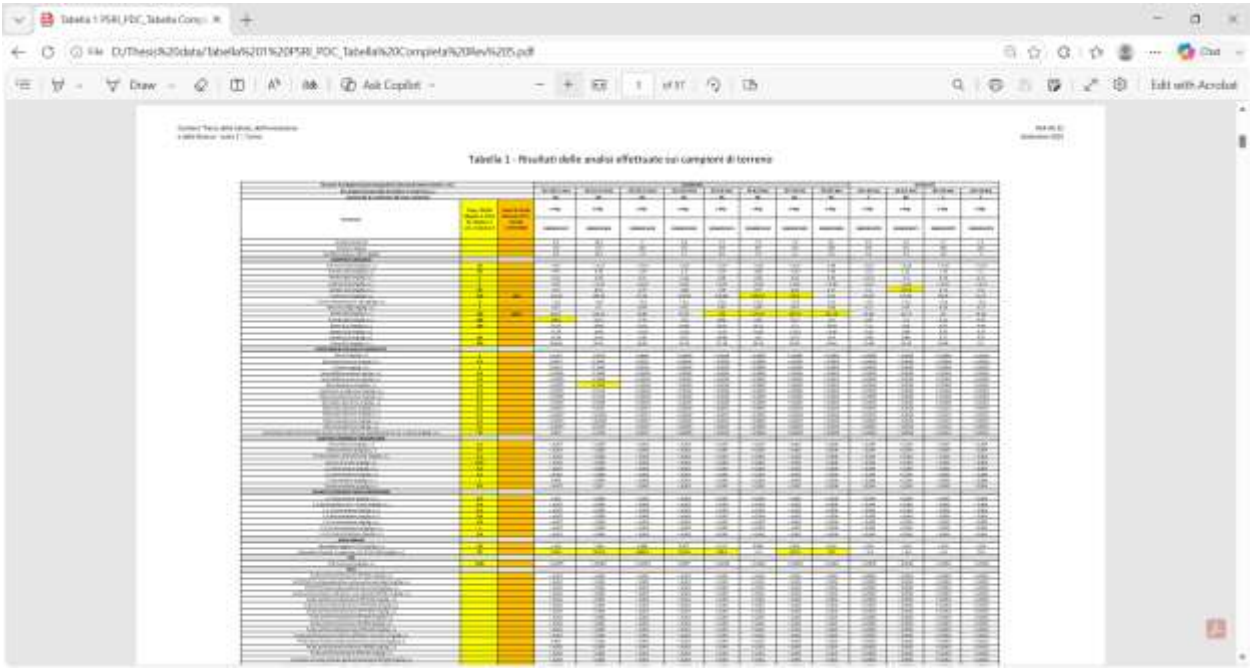

Figure 14 Comprehensive analytical summary report (Tabella 1 PSRI_PDC) compiled as a universal PDF asset

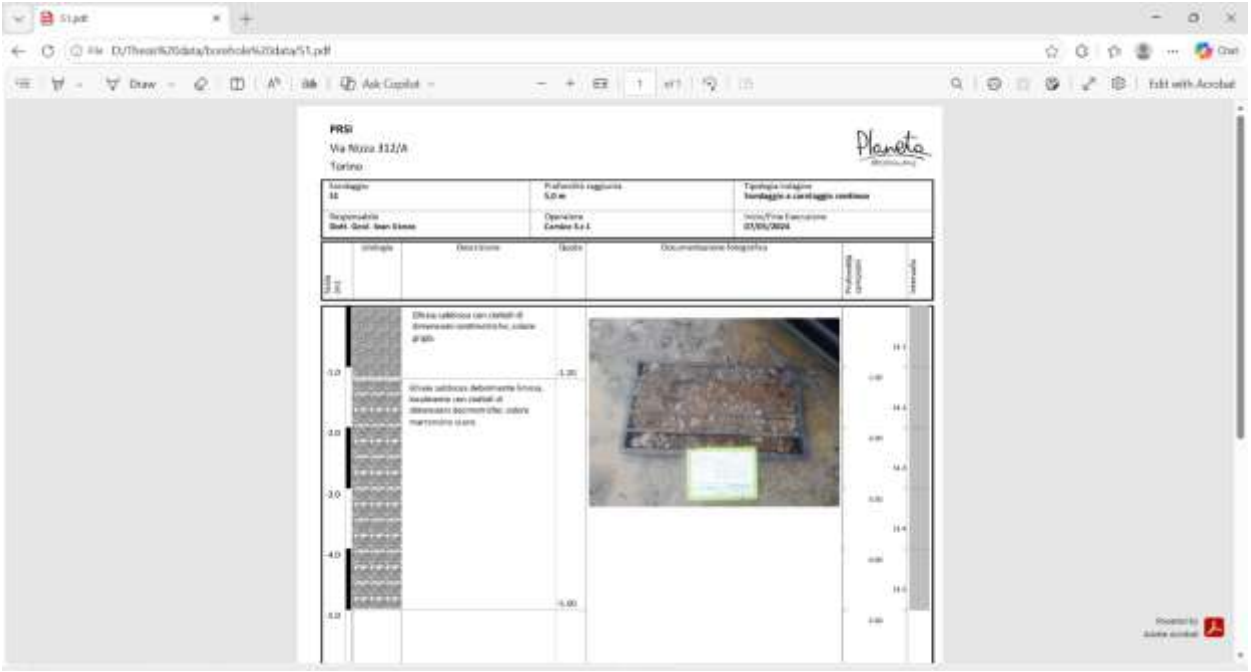

Figure 15 Representative borehole PDF report (e.g., S1.pdf) showing stratigraphy, technical metadata, and core sample photographs

4.2.2 Development environment

The technical implementation required selecting an integrated development ecosystem capable of reconciling complex 3D engineering geometries, geospatial positioning, runtime tabular querying, and mobile hardware constraints. To determine the most viable development engine for the reasearch goal, an initial multi-criteria evaluation was conducted (Table 1)across three leading paradigms: Native mobile development via Google ARCore, Unreal Engine, and Unity.

Table 1 Comparison between three development engines

Evaluation Criteria	Native ARCore (Android SDK)	Unreal Engine (v5.x)	Unity (Unity 6 LIS)
Primary Language	Java / Kotlin	C++ / Blueprints	C#
Licensing Model	Open Source (Apache 2.0)	Free tier (Royalty above threshold)	Free (Personal tier under revenue limit)
3D Asset Ingestion	Manual parsing / OpenGL / Sceneform (Deprecated)	Native FBX/OBJ via Datasmith	Native FBX/OBJ ingestion with mesh optimization
Geospatial & BIM Support	Minimal; requires third-party coordinate transformation	Cesium for Unreal (high computational overhead)	Cesium for Unity (optimized WGS84 tile slicing)
Cross-Platform Deploy	Android exclusive	Cross-platform (large package footprint)	Cross-platform via AR Foundation abstraction
Development Overhead	High (low-level graphics pipeline authoring)	Moderate to High (steep learning curve, strict hardware)	Moderate (rapid iterative prototyping, managed memory)

Native ARCore development was eliminated during the preliminary phase. While it provides direct, low-level access to device sensors and visual-inertial odometry (VIO) routines, it lacks an integrated real-time graphics authoring pipeline. Importing, rendering, and interactively querying multi-element BIM geometries would have necessitated writing custom rendering engines and collision-detection algorithms from scratch, making it unfeasible within the operational constraints of the thesis.

Unreal Engine provides enterprise-grade rendering capabilities and supports geospatial pipelines via Cesium for Unreal. However, its core architecture is optimized for desktop and high-end console fidelity rather than lightweight mobile execution. The substantial storage footprint of

exported APK packages, demanding memory overhead, and complex C++ compilation cycles offered negligible advantage for an application whose primary purpose is positional precision and structured database retrieval rather than photorealistic illumination.

Finally, Unity 6 was selected as the operational environment. Unity balances high-performance real-time rendering with low hardware overhead, ensuring robust execution on mid-tier mobile hardware. Its component-based architecture and managed C# scripting environment facilitate rapid development of spatial recasting logic, dynamic file I/O operations, and asynchronous parsing of tabular datasets.

The technical foundation of the mobile application was established within Unity 6.3 LTS (Editor version 6000.3.10f1). The initial project configuration was instantiated using the official **AR Mobile Template**, providing a baseline scene hierarchy preconfigured for augmented reality workflows, including the root AR Session, XR Origin (AR Rig), and an interactive user interface canvas.

4.2.3 Project Configuration and Build Environment

Following the selection of the development engine, the operational project environment was configured to target mid-range Android mobile hardware. The implementation was established using the Universal Render Pipeline (URP) within Unity 6.3 LTS (version 6000.3.10f1), deployed under the package identifier “com.polito.geobimviewer”.

The build pipeline and runtime graphics settings were structured to optimize GPU throughput, manage system memory constraints, and ensure reliable execution of spatial tracking routines on the physical device.

Table 2 Unity project build settings, target platform configurations, and technical justifications for Android deployment.

Parameter	Configuration Setting	Technical Justification
Minimum Operating System	Android 10.0 (API Level 29)	Ensures compatibility with modern ARCore sensor-fusion routines and runtime permission models.
Target Architecture	ARM64	Mandatory 64-bit execution requirement for modern mobile processors, ensuring compatibility with native AR plugins.
Scripting Backend	IL2CPP	Compiles intermediate C# code directly into optimized native C++ binaries, improving CPU execution speed for real-time mathematical operations.
Graphics API	OpenGL ES 3.0	Lightweight mobile rendering backend selected to reduce graphics driver overhead and stabilize frame rates on mid-tier GPUs.
Active Input Handling	Input System Package (New)	Provides a unified event-driven input interface for touch gestures, multi-touch raycasting, and screen-space interactions.

To support augmented reality routines and global georeferencing, the engine was provisioned with the software packages via the package management system. Core spatial computing capabilities were established by integrating AR Foundation alongside the Google ARCore XR Plugin (version 6.3.4), providing the communication bridge to the device's camera stream and inertial measurement unit (IMU).

Planetary-scale coordinate transformations and geospatial streaming were enabled by incorporating the Cesium for Unity runtime package (version 1.14.0). Utilizing this platform required creating a Cesium ion developer account and generating a dedicated access token, which authenticates the Unity editor to securely access cloud-hosted geospatial services. Within the development environment, Cesium was deployed across two distinct operational modes. First in the **Off-Site Exploration,** Cesium ion was used to stream Google Photorealistic 3D Tiles into the scene. This cloud-streamed terrain and building mesh provides a detailed urban context, allowing the user to review the subsurface model relative to real-world topography and surrounding structures prior to field deployment. The other use of cesium was in **On-Site scene**, in this mode Spatial tracking relies on the CesiumGeoreference component rather than streaming heavy tile assets over mobile networks. This component anchors the local CAD/BIM Cartesian vectors directly to the Earth-Centered, Earth-Fixed (ECEF) reference frame using WGS84 geographic coordinates (latitude, longitude, and ellipsoidal height), eliminating floating-point precision loss and aligning the digital borehole geometry with its physical coordinates on site.

Together, these build configurations, AR tracking plugins, and geospatial coordination libraries establish the baseline software environment upon which the runtime scene architecture and tracking scripts are deployed.

4.3 System design

4.3.1 System architecture

The comprehensive software architecture of the GeoBIM Viewer is illustrated in Figure 17. The system is organized into a modular pipeline that ingests external project inputs, routes execution through dedicated operational scenes, and interfaces with a unified data pipeline executing on the Android runtime.

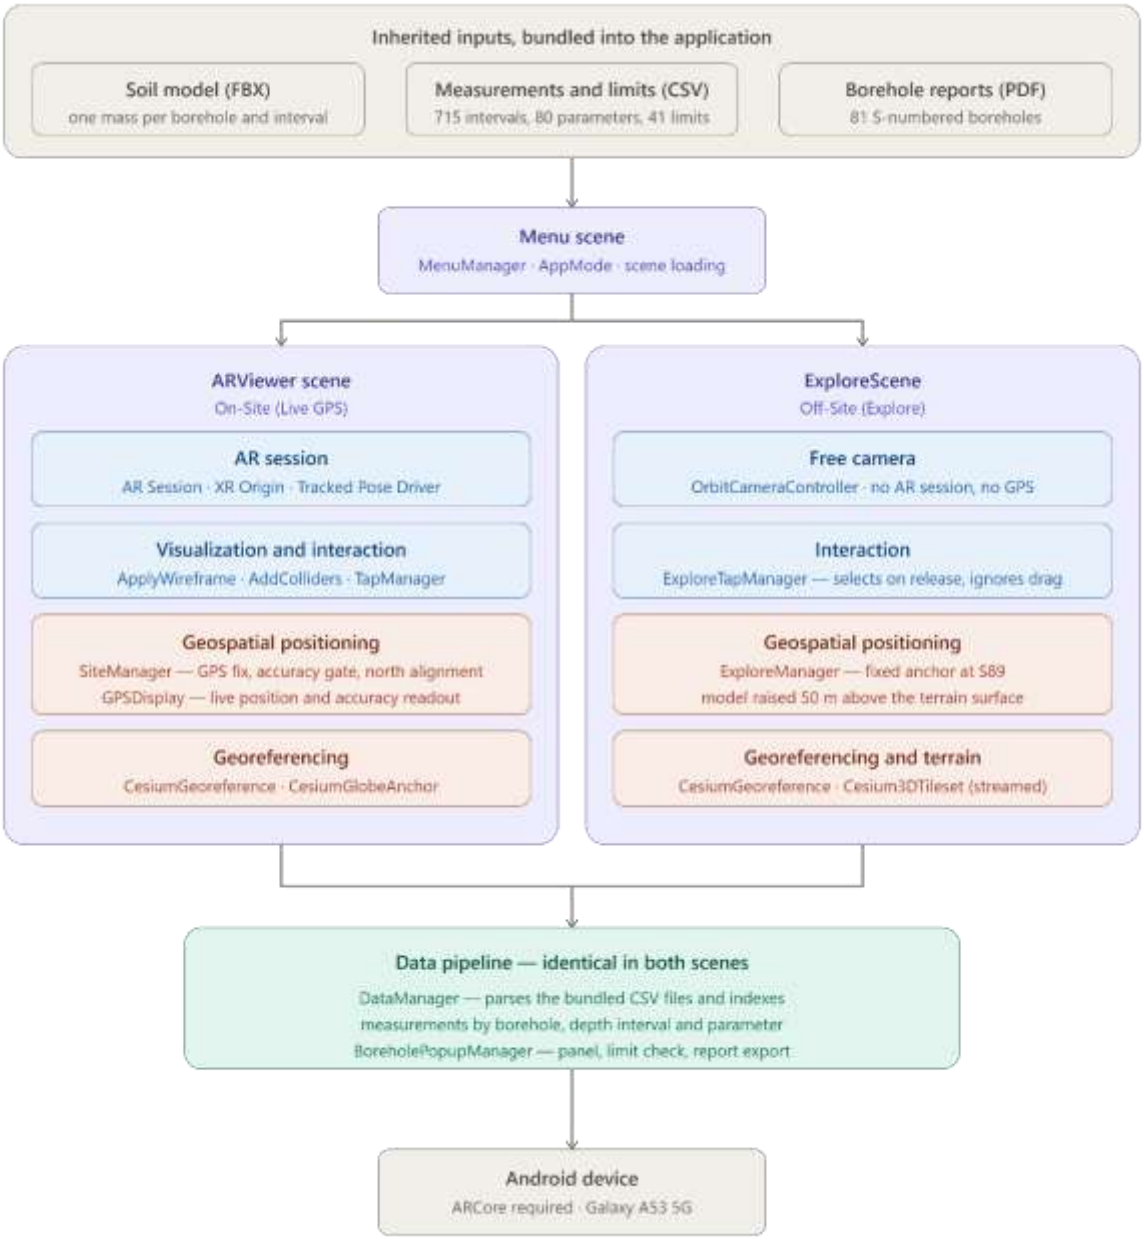

Figure 16 Comprehensive system architecture of the GeoBIM Viewer, illustrating input bundling, operational scene segregation, functional subsystem layering, and the unified data core

The application was deployed and evaluated on a Samsung Galaxy A53 5G. This device was selected as the experimental benchmark because it represents the global mid-tier Android smartphone segment, which constitutes the largest volume of enterprise mobile deployments in civil and environmental fieldwork according to industry shipment metrics (e.g., Counterpoint Research data on the Samsung Galaxy A-series). Verifying performance on mid-range hardware

demonstrates that the GeoBIM AR system does not require high-end flagship or enterprise LiDAR hardware, confirming its practical viability on standard site equipment.

To ensure reproducibility across both technical and academic implementations, the exact experimental configuration—encompassing physical device hardware, software runtimes, tracking layers, and compiled package metrics—is detailed in Table 3.

Table 3 System reproducibility specification: Hardware parameters, software runtimes, and build metrics for the test device

Layer / Domain	Component / Metric	Technical Specification / Value
Physical Hardware	Device Model	Samsung Galaxy A53 5G (SM-A536B)
	System-on-Chip (SoC)	Samsung Exynos 1280 (5 nm process)
	CPU	Octa-core (2×2.4 GHz Cortex-A78 & 6×2.0 GHz Cortex-A55)
	GPU	ARM Mali-G68 MP4
	Memory (RAM) / Storage	6 GB LPDDR4X / 128 GB UFS 2.2
	Display	6.5” Super AMOLED (1080×2400 px, 120 Hz)
Sensors & Tracking	Positioning (GNSS)	Multi-constellation: GPS, GLONASS, Galileo, BDS, QZSS
	Inertial Sensors	Tri-axial Gyroscope & Accelerometer (6DoF IMU)
	Directional Sensor	3-axis Geomagnetic Sensor (Magnetometer)
	Optical Sensor	64 MP main sensor (f/1.8, OIS) configured for AR video feed
Software Runtime	Operating System	Android 14 (One UI 6.0; Target API Level 34, Min API 29)
	Spatial Computing Framework	Unity AR Foundation 6.3.4 & Google ARCore XR Plugin 6.3.4
	Geospatial Runtime	Cesium for Unity (v1.14.0)
	Development Engine	Unity 6.3 LTS (Editor version 6000.3.10f1)
	Scripting & Graphics Backend	IL2CPP (ARM64) / OpenGL ES 3.0
Model & Package Metrics	Subsurface Model Assets	1,590 elements (117 masses, 629 mass floors, 729 columns)
	Polygon / Vertex Count	≈245,000 triangles / ≈180,000 vertices
	Application Package (APK) Size	≈85 MB (inclusive of streaming PDFs and tabular CSVs)

Sensor Integration and Spatial Positioning Chain

The system coordinates the device’s physical sensor array to achieve stable spatial tracking, true geographic orientation, and accurate borehole overlay on site. Table 4 outlines the functional mapping between the phone's physical hardware interfaces, the intermediary processing engines, and their system-level roles.

Table 4 Functional mapping of mobile hardware capabilities and integrated software frameworks.

Hardware Subsystem	Integrated Framework	Functional Role in GeoBIM Viewer
RGB Camera	ARCore / AR Foundation	Captures continuous environmental video feed for AR background rendering and optical feature-point tracking.
IMU (Gyro + Accel)	Google ARCore	Measures high-frequency linear acceleration and angular velocity for Visual-Inertial Odometry (VIO).
Magnetometer	SiteManager (Custom Script)	Measures Earth's magnetic field vector to establish initial True-North heading alignment.
GNSS / GPS Receiver	CesiumGlobeAnchor / SiteManager	Acquires real-world WGS84 geographic coordinates to anchor virtual models to the physical site.
Capacitive Touchscreen	Unity Input System / UI Canvas	Processes screen raycast taps on borehole cylinders and manages UI modal dialogs.
Network Interface (Wi-Fi/5G)	Cesium ion Client	Streams high-resolution Google Photorealistic 3D Tiles dynamically in Off-Site simulation mode.

The data flow governing spatial registration and tracking is illustrated in Figure 17. The architecture coordinates five functional tiers:

Sensors: Acquires raw optical, inertial, geomagnetic, and satellite signals from the physical site environment.

Tracking: Google ARCore executes visual-inertial odometry (VIO) to establish local relative movement, while the heading subsystem filters raw geomagnetic noise. AR Foundation standardizes these outputs into a local camera pose and real-time camera yaw.

Registration: Evaluates device satellite coordinates against an operational quality threshold and pairs the accepted horizontal position with a fixed, surveyed benchmark elevation. The geographic position is transformed from WGS84 ellipsoidal notation into Earth-Centered, Earth-Fixed (ECEF) vectors anchored to reference borehole S89 via Cesium.

Alignment: Reconciles the local camera coordinate frame with real-world orientation by subtracting the filtered compass heading from the live camera yaw, applying the resulting rotation directly to the tracking root.

Output and Offline Preparation: Renders the borehole masses and transparent cell boundaries in real-world space. To avoid runtime coordinating drift and performance overhead, the model's pivot offset relative to reference borehole S89 is pre-calculated and baked directly into the asset before compilation.

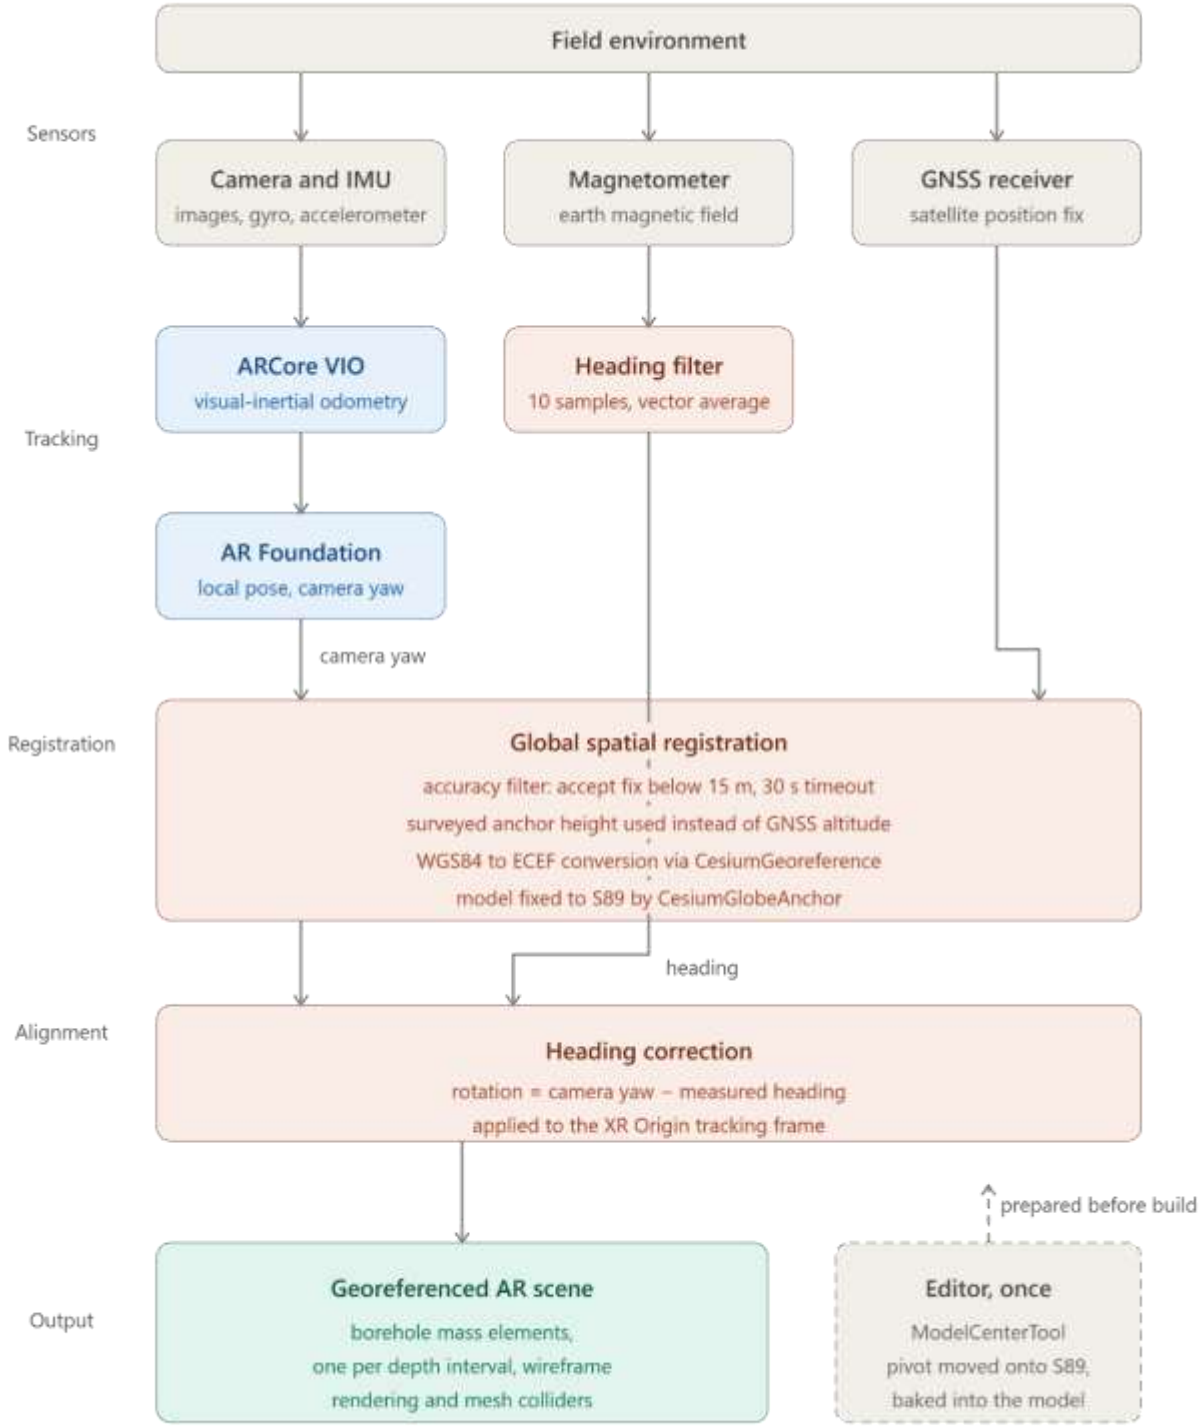

Figure 17 Data flow schema of the sensor-fusion and on-site positioning layer, showing heading reconciliation and geographic-to-local coordinate alignment.

The on-site tracking and spatial registration pipeline operate through three functional stages:

1. **Visual-Inertial Odometry:** The Google ARCore subsystem tracks local motion relative to the environment by matching optical features across camera frames and pairing them with high-frequency IMU data. AR Foundation translates this tracking into local camera translation (X, Y, Z) and rotation.

2. **Heading Alignment:** The raw magnetometer output is prone to transient magnetic disturbances on site. The system averages 10 consecutive magnetometer samples to compute a stable magnetic heading. To calculate the final rotational correction applied to the scene, this smoothed heading is subtracted from the real-time camera yaw ($\Delta\theta =$ Camera Yaw $-$ Heading), aligning virtual geometry with true field orientation.

3. **Geospatial Anchoring:** Geographic positioning queries the device's native satellite location service via Android. Position updates are evaluated against a validation gate: the reading must achieve an estimated accuracy of ≤ 15 mwithin a 30-second timeout. Once validated, real-world georeferencing is executed by CesiumGlobeAnchor using reference borehole S89 as the origin. Because the geometric center offset between the CAD model and the S89 borehole anchor was pre-calculated and baked into the model pivot within the editor (via ModelCenterTool), runtime operations require no secondary pivot shifts. Cesium directly transforms the validated WGS84 coordinates into Earth-Centered, Earth-Fixed (ECEF) Cartesian vectors, rendering the individual borehole depth interval meshes in their exact real-world positions.

Dual Operating Modes and Shared Data Core

The system architecture is bifurcated into two operational modes managed across dedicated scenes to prevent unnecessary computational overhead:

- **On-Site AR Mode:** Designed for field verification. It utilizes live camera feeds, local sensor fusion, and satellite georeferencing to overlay the subsurface Revit model and InfraWorks context directly onto the terrain. To ensure reliability in remote locations with poor connectivity, all geometries, CSV analytical tables, and PDF logs are hosted entirely on local device storage.

- **Off-Site Exploration Mode:** Designed for preliminary office planning and desktop inspection. Hardware tracking and camera feeds are deactivated, while Cesium ion streams global photorealistic 3D tiles over a network connection to provide comprehensive urban context.

Both modes interface with a unified data pipeline layer. The parsing routines that ingest the 715 interval analyses and D.Lgs 152/06 statutory limits, as well as the interactive UI popup panels, operate identically across both environments. A borehole selected in either mode invokes the exact same analytical logic and provides identical access to archival geotechnical documentation.

4.4 Implementation of the Layers

4.4.1 AR Visualization and Interaction

Augmented reality applications can support several interaction modalities, such as mid-air hand gestures, gaze-based selection, and touch interfaces. On an active construction site, mid-air gestures—common in head-mounted displays—are impractical, while gaze-dwelling is slow and tiring. A direct capacitive touch interface on a handheld smartphone was selected instead. Touching the screen is reliable, precise, and immediately familiar to field personnel, making it well-suited for site inspections.

The AR visualization and interaction layer is the visible face of the application. It is where the user enters the augmented reality view, sees the soil model placed over the real world, and interacts with it by looking, moving, and tapping. This layer is responsible for starting the AR session, preparing the model for display, and turning the user's taps into borehole selections. The augmented reality functionality is built on AR Foundation, configured through an XR Origin in the scene. The XR Origin defines the origin of the AR world and carries the AR camera, whose position and rotation are driven by the device's tracking. The AR session manages the device's camera feed and its understanding of the surrounding space. Together these provide a live camera view that the model is drawn over, and the motion tracking that keeps the model fixed in place as the user moves. This setup is what allows the user to walk around the site and see the model hold its position in the real world. This layer is designed using three scripts that work together: Apply Wireframe, Add Colliders, and Tap Manager. Figure 18 shows how these scripts operate across two stages of preparing the model when the scene loads and responding to the user when a borehole is tapped.

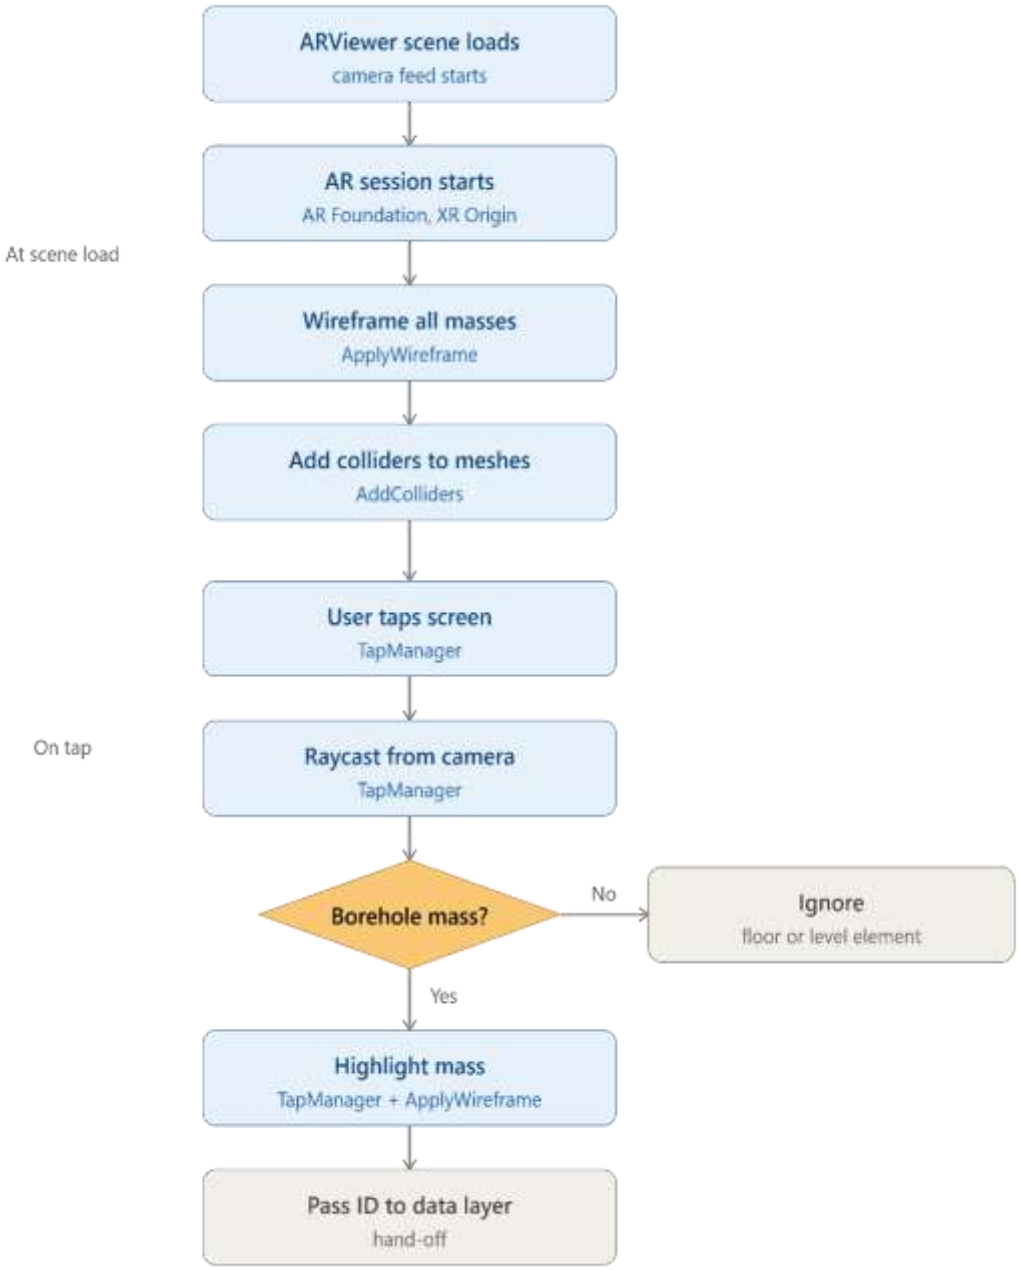

Figure 18 AR visualization and interaction. At scene load the model is prepared for display and made tappable; on tap, the selected borehole is identified, highlighted, and its identifier passed to the data layer.

The operational sequence divides naturally into two phases: preparing the scene when it loads and responding when the user taps an element.

Scene Initialization (At Scene Load)

When the ARViewer scene loads, the camera feed starts and Unity's AR Foundation initializes the AR Session and XR Origin. The device tracks the physical movement of the phone directly to the virtual camera pose, anchoring the digital model to the site environment.

Before any user interaction takes place, two automated preparation steps run:

- **Wireframe Material Application (ApplyWireframe):** If the subsurface soil model were rendered as solid blocks, the outer volumes would hide the inner borehole columns. ApplyWireframe solves this by assigning a clean, transparent wireframe shader to every mass element in a single pass at startup. This outlines the excavation cells while leaving the deeper stratigraphy visible.
- **Collider Generation (AddColliders):** The imported 3D geometry does not contain Unity colliders by default, meaning touch rays cannot detect it. Because adding colliders by hand across hundreds of elements is impractical and error-prone, AddColliders runs automatically on load. It steps through each mesh in the model and adds a matching MeshCollider, making every depth interval selectable.

Touch Interaction Pipeline (On Tap)

Once the scene is ready, the user interacts with the model by tapping borehole columns, which is handled by TapManager:

1. **Raycasting:** When the user touches the screen, TapManager fires a ray from the camera through that touch coordinates into the virtual scene.
2. **Object Filtering:** The ray might hit elements that are not boreholes, such as mass floors or project level planes. The script checks the name of the struck object; if it belongs to a floor or level, the event is ignored, ensuring that only genuine borehole segments respond.
3. **Visual Feedback:** When a valid borehole mass is struck, TapManager coordinates with ApplyWireframe to swap that element’s material from wireframe to a distinct highlight color. Only one segment is highlighted at a time—the script returns the previous selection back to wireframe before highlighting the new one.
4. **Data Layer Hand-Off:** Finally, TapManager reads the borehole identifier and depth interval from the selected object and passes them directly to the data pipeline layer, which opens the attribute panel with the matching measurements.

4.4.2 Data Pipeline

This layer handles the project's data, loads and prepares the data in the background, then retrieves and presents it when a borehole is selected. In the flow of the application it is the second layer, acting after a borehole has been chosen in the AR layer. It is built from two scripts: Data Manager, which loads and organizes the data, and Borehole Popup Manager, which displays it. Figure 19 shows how the data is prepared and how it is retrieved and presented once a borehole is selected.

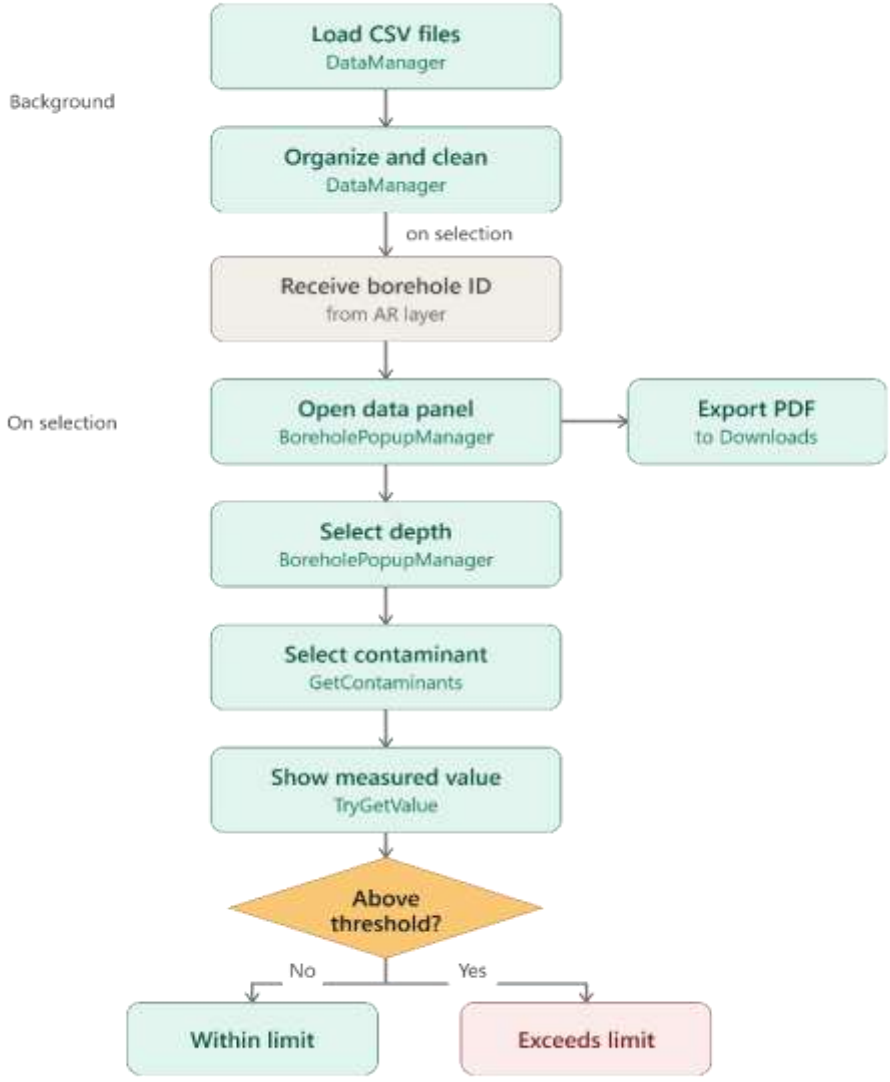

Figure 19 Execution workflow of the data pipeline layer, showing background CSV ingestion, UI attribute binding on selection, regulatory threshold evaluation, and PDF report export

The data is prepared ahead of time so that it is ready the moment it is needed. While the application starts up, Data Manager loads two files in the background. The first is a CSV of contaminant measurements, where each row records the values measured at a given depth in each borehole. The second is a CSV of regulatory thresholds, listing the legal limit for each contaminant. On Android, files bundled with the app cannot be read directly from disk, so they are loaded through a web request to the streaming assets path; this is necessary because the files are stored inside the compressed application package.

As the measurement file is read, Data Manager organizes it into a nested structure so that any value can be found quickly later. The data is grouped first by borehole, then by depth, then by contaminant. This three-level organization mirrors how the user will navigate the data, first choose a borehole, then choose a depth, then select a contaminant. Building this structure once, ahead of time, means that later lookups are immediate, rather than searching the whole file each time.

Several details are handled during loading. The borehole identifier is extracted from the location code in each row, so that it matches the identifier read from the model when a borehole is tapped. Depths are stored as a combined key in the form "top base", representing the depth interval the sample was taken from. Empty cells and entries marked as not measured, such as "n.d.", are skipped, so that only real measurements are stored. Numbers are parsed using a fixed format, so that decimal values are read correctly regardless of the device's regional settings. The thresholds file is loaded in the same way, into a simpler list pairing each contaminant with its limit.

The layer then acts when a borehole is selected in the AR layer. Its identifier is passed to Borehole Popup Manager, which opens a panel for that borehole. The panel has two tabs, one for depth and one for contaminant. The user first selects a depth, and DataManager returns the contaminants measured at that depth. The user then selects a contaminant, and its measured value is shown.

At this point the value is compared against statutory limits established by Italian Legislative Decree No. 152/2006 (D.Lgs. 152/06) Borehole Popup Manager requests the value from Data Manager, then requests the corresponding threshold, and compares the two. If the measurement is above the limit, it is flagged as exceeding the threshold; if it is below, it is marked as within the limit. This comparison is the central purpose of the application, as it turns a raw number into a clear regulatory result that the user can act on in the field.

The panel also gives access to the full reports. The user can download the borehole's PDF report or the complete contaminant data table. The requested file is read from the application package

and copied to the device's Downloads folder, then registered with the system so that it appears immediately in the device's file browser and can be opened in any PDF viewer.

4.4.3 Geospatial Positioning

The fundamental challenge in outdoor augmented reality is the reconciliation of two incompatible spatial frameworks. Engineering models are authored in a local Cartesian system, a flat Euclidean space in which coordinates are expressed as offsets along orthogonal axes. The physical site, by contrast, lies on a curved Earth, described by geodetic coordinates: latitude and longitude referenced to an ellipsoid, and height referenced either to that ellipsoid or to the geoid. For a model to be displayed correctly over the physical world, a rigorous transformation must bridge the two.

The transformation required is a rigid-body one, since the model's scale is fixed and must not be altered. It therefore has six degrees of freedom: three of translation and three of rotation. Two of the rotational degrees are resolved directly by the device, which measures the direction of gravity and can establish its vertical axis without external reference. The remaining unknowns are the three translational components and the rotation about the vertical axis, the azimuth. Determining position and azimuth is thus the whole of the alignment problem.

A simplification applies at the scale of a single site. The investigated area extends roughly three hundred meters in each direction, over which the curvature of the ellipsoid is negligible relative to the accuracy attainable. The curved geodetic surface may therefore be approximated locally by a tangent plane, in which positions are expressed as east, north and up displacements from a chosen origin. This reduces the problem into relating one local Cartesian frame to another, rather than performing a curved-surface transformation continuously.

Evaluation of alignment strategies

Four approaches were considered.

Visual fiducial tracking places **physical markers** at surveyed positions on site, which the device recognises to establish its pose. Accuracy is high and independent of satellite reception, but the approach requires the site to be physically prepared in advance, and requires markers to remain in place, visible and undamaged. On an active brownfield site this cannot be assumed.

Multi-point registration determines the transformation by measuring several known points on site and fitting the model to them. Since the reference points are boreholes, whose coordinates are already surveyed, the measurement could be made with the device itself, without additional

instruments. This is theoretically the strongest of the approaches: it derives azimuth from geometry rather than from a magnetic measurement and provides redundancy that would reveal an inconsistent observation. It was not adopted because it requires the reference points to be measured at the site, and site access was not available during the development period. It is identified as the principal direction for further work.

Image-based localization determines position and azimuth by matching the camera view against a reference imagery database, resolving both unknowns simultaneously and without magnetic measurement. It was not adopted here because it depends on prior imagery coverage of the site, which cannot be assumed for an enclosed brownfield area, and because it imposes a substantially greater computational demand on the device.

Single-point geodetic anchoring with sensor fusion binds one predetermined point of the model to a known surveyed coordinate, establishing translation, and derives azimuth from the device's magnetic sensor. It requires no site preparation, no additional instruments, and no prior imagery.

Single-point anchoring with sensor fusion was adopted, as the only approach viable under the constraints of this work. Translation is established from a single surveyed coordinate combined with the device's own position estimate; the vertical axis follows from the measured direction of gravity; and the azimuth is obtained from the magnetic sensor.

The consequence of this choice is that position and azimuth are determined independently rather than jointly, and each carries the error characteristics of its own source. Magnetic azimuth is susceptible to local disturbance, which is common in-built environments, and therefore requires filtering to be usable. The accuracy achieved and the limitations that follow from this choice are examined in Chapters 5 and 6.

The geospatial positioning layer aligns the virtual subsurface model with its physical position on site, orchestrated by SiteManager.cs and the Cesium for Unity runtime. The runtime sequence, illustrated in Figure 20, separates offline geometric preparation from live on-site positioning.

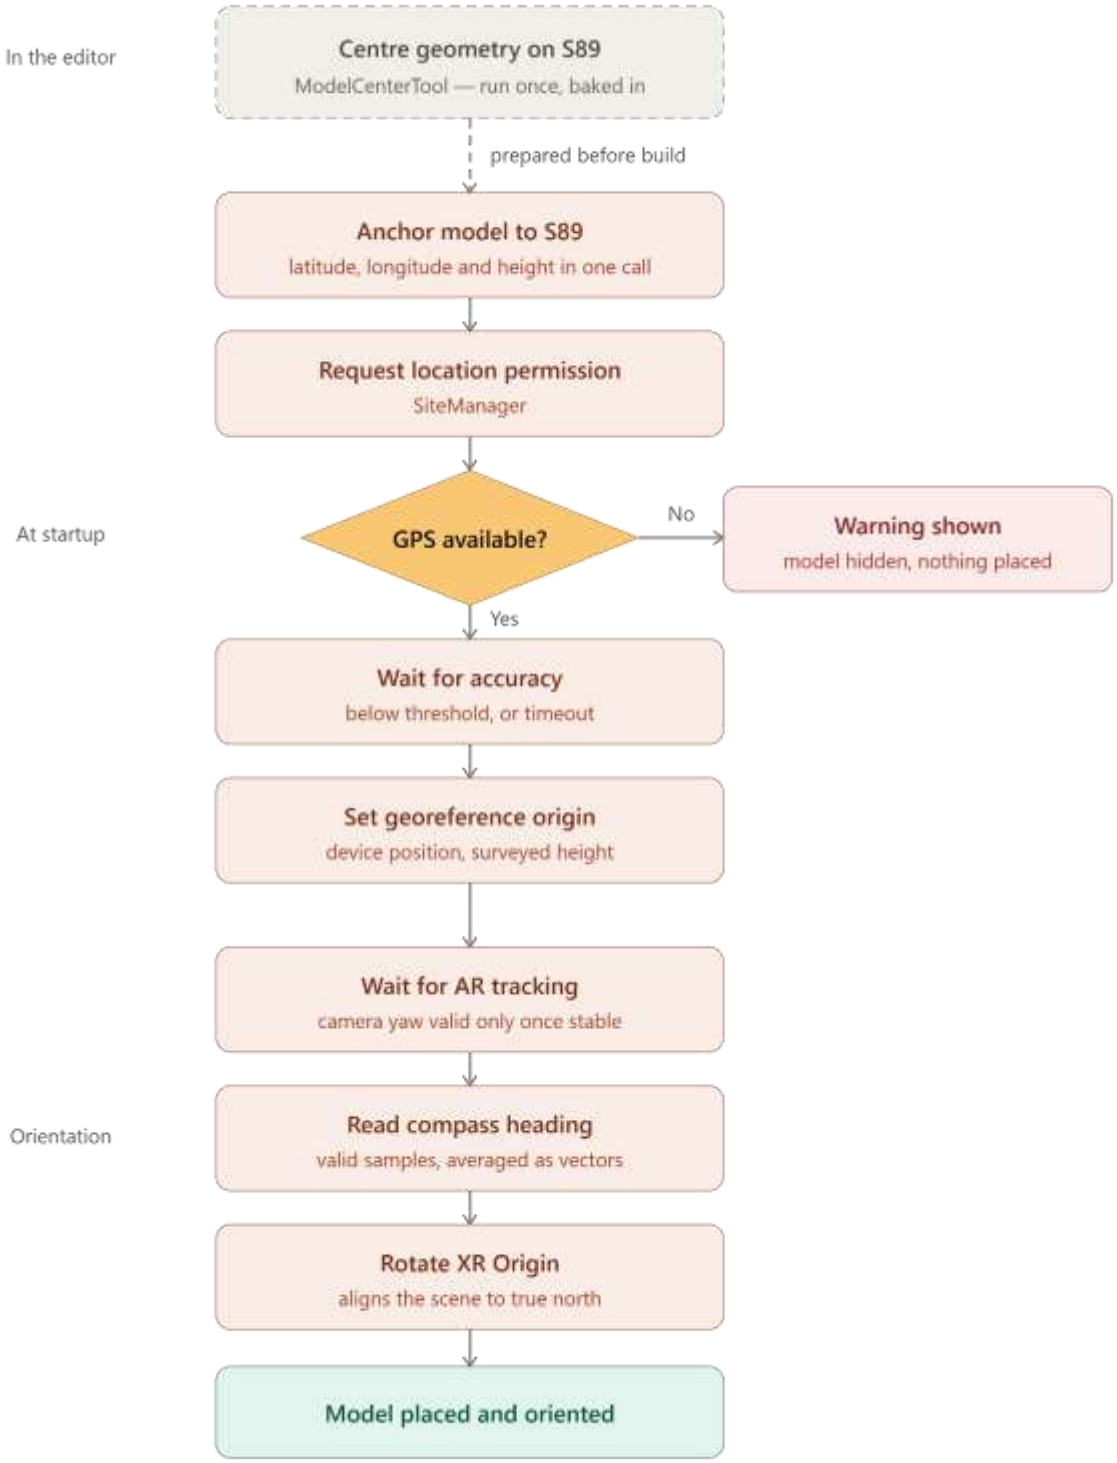

Figure 20 Execution workflow of the on-site geospatial positioning layer

1. Offline editor preparation

Before the application is built, the model's geometric origin must be aligned with a physical reference anchor. When exported from Autodesk Revit as FBX, the geometry retains Revit's internal origin, which lies several kilometres from the investigated area.

The correction is applied once in the Unity Editor, using ModelCenterTool.cs. The utility locates the reference borehole S89 within the model, measures its offset from the file origin, and applies the translation so that S89 coincides with the model pivot. The corrected model is then saved, and S89 serves as its permanent geometric origin.

Performing this once, rather than at runtime, is a deliberate decision. Cesium positions the model through its own globe anchor, and a runtime correction would act on the same object simultaneously. Geometric correction is permanent and belongs to the asset, while live positioning belongs to the geospatial runtime.

2. Startup initialization and satellite positioning validation

When the on-site scene loads, the application anchors the pre-prepared model to the surveyed coordinates of borehole S89. This process relies on a few critical steps to ensure the placement is both stable and accurate:

- **Single-call coordinate assignment:** The latitude (45.025369° N), longitude (7.664259° E), and surveyed ground elevation (234.00 m) are passed to the CesiumGlobeAnchor all at once. Assigning them one by one forces the system to recalculate after each step, which can momentarily strand the anchor at an invalid coordinate and cause a persistent placement error.
- **Location permissions and availability:** Before attempting to place the model, SiteManager.cs verifies Android location permissions. If GPS is disabled or access is denied, the app displays a warning and completely hides the model. This is a deliberate safety measure: showing unanchored subsurface geometry is worse than showing nothing at all, as field personnel might rely on an inaccurate placement.
- **Accuracy filtering gate:** Mobile devices usually show a rough, network-based location before securing a true satellite lock. To clear out this initial drift, SiteManager.cs waits until the reported horizontal accuracy improves to at least 15 meters. If it hasn't reached this threshold after 30 seconds, the system times out and simply proceeds with the best available fix to avoid stalling the application.
- **Georeference origin initialization:** Once a reliable fix is acquired, the CesiumGeoreference origin is centered on those horizontal coordinates. However, it ignores the device's reported altitude. Satellite altitude is measured against a

mathematical ellipsoid rather than actual sea level, meaning it can easily be off by tens of meters. Instead, the app pairs the user's horizontal position with the fixed, surveyed ground elevation. Because this exact same height is applied to both the user's origin and the model's anchor, any discrepancy between the surveyed datum and the ellipsoid cancels out naturally, ensuring the model sits exactly where it should relative to the user.

3. Sensor fusion and heading alignment

Anchoring the model to S89 establishes geographic position but leaves orientation unresolved. When an AR session initializes, the device's inertial sensors establish the vertical axis from gravity, but the system has no knowledge of geographic azimuth; its forward direction defaults to wherever the device was pointing when the session began.

To align the scene with true north, SiteManager.cs executes an orientation routine once ARCore reports a stable tracking state, since camera orientation is not meaningful before then.

Validity filtering. Readings reported as invalid by the Android sensor subsystem are discarded rather than using them to get averaged. Including them yields a heading of zero and therefore no correction at all.

Vector-averaged directional sampling. Because compass angles wrap from 359° to 0°, a standard arithmetic mean across the boundary — of 359° and 1°, for instance — returns 180°, pointing south. The system buffers N = 10 valid readings α_i, decomposes each into unit vector components, and computes the directional mean (Mardia & Jupp, 2000):

$$\bar{X} = \frac{1}{N}\sum_{i=1}^{N}\cos(\alpha_i),\ \bar{Y} = \frac{1}{N}\sum_{i=1}^{N}\sin(\alpha_i)$$

$$\bar{\alpha} = \operatorname{atan2}(\bar{Y}, \bar{X})$$

XR Origin reorientation. The angular offset is obtained by subtracting the filtered heading from the camera yaw reported by AR tracking:

$$\Delta\theta = \theta_{\text{cam}} - \bar{\alpha}$$

Applying Δθ about the vertical axis of the XR Origin rotates the tracking reference frame, bringing the model into alignment with true north. The direction of this correction was established empirically rather than derived, since the sign depends on whether the rotation is applied to the scene content or to the tracking frame. A manual adjustment of one degree in either direction is also provided, allowing residual error from local magnetic disturbance to be corrected on site.

4.4.4 Off-Site Exploration Mode

The augmented reality mode described so far can only be used at the site itself. This is by design, since the purpose of the application is to deliver subsurface data in the field. It does, however, leave two practical gaps.

The first is access. Reviewing the model, preparing for a site visit, presenting the work to colleagues, or demonstrating the application all require being physically present at the site, which is rarely convenient and sometimes not possible at all.

The second is scale. The site measures roughly 270 by 330 meters. A user standing inside the model in augmented reality sees only their immediate surroundings and cannot take in the site as a whole or understand how the boreholes are distributed across it.

A second mode was therefore developed to address both. In this mode the application does not use device camera, GPS, or augmented reality tracking. Instead, the soil model is displayed within a three-dimensional representation of the real site, streamed as photorealistic terrain tiles from a geospatial tile service. The user navigates freely with touch controls, orbiting, panning and zooming around the model rather than walking through it.

The two modes are implemented as separate scenes, selected from the opening menu. This separation is deliberate: the two have almost nothing in common at the technical level, since one depends on a live camera feed and motion tracking while the other does not. Keeping them apart means that a change to one cannot break the other, and that the augmented reality scene carries no code for functionality it never uses.

What the two modes do share is the georeferencing. Both anchor the soil model to the same reference borehole using the same geospatial library, so the model occupies its true geographic position in both. The difference is only in how the surrounding world is represented, and how the user moves through it.

One deliberate deviation from true position is made. The soil model lies entirely below ground, while the streamed terrain represents the ground surface, so the terrain would completely conceal the model. The model was aligned with the cesium tiles. Also, the horizontal position remains georeferenced to the true site coordinates. This mode is consequently intended for orientation and overview rather than for judging depth relationships, which are represented correctly in the on-site mode.

Because the terrain tiles are streamed rather than stored in the application, this mode requires an internet connection.

4.4.5 Build and deployment

With the scenes, scripts, and data files fully set up, the project was packaged as an Android application and installed on the test device. All the field results and observations presented in Chapter 5 were gathered directly from these final builds, which were evaluated both in the Unity Editor and on the physical phone. To safely manage the development process, the entire project was kept under version control from the start in GitHub. This provided a reliable safety net, making it easy to roll back to a stable, working state whenever experimental changes or new features were unsuccessful.

Chapter 5 Results

5.1 Overview of the Implemented Application

The work resulted in a functioning mobile application, built with Unity and deployed as an Android package. It runs on a standard consumer smartphone, a Samsung Galaxy A53, and requires no specialist hardware beyond the augmented reality support already present in current Android devices. This means the application can be used by professionals already on site, with the device they are carrying.

When the application is opened, the user is presented with a menu scene offering two modes. The On-Site mode uses the device camera and satellite positioning to display the soil model at its true location, so that the user can walk across the site and see the subsurface model in place beneath them. The Off-Site mode displays the same model within a photorealistic three-dimensional representation of the site, navigated with touch controls, and can be used anywhere.

Both modes present the same model and give access to the same information. The soil model is rendered as a set of three-dimensional masses, one for each borehole location, drawn as a wireframe so that the internal structure remains visible rather than being hidden behind solid surfaces. Selecting a borehole highlights it and opens a panel showing the measurements recorded there. The user chooses a depth, then a contaminant, and the measured concentration is displayed together with a comparison against the corresponding regulatory limit, marked clearly as either exceeding or within the threshold. The field documentation, including the individual borehole report and the complete contaminant table can be downloaded as PDF files within the application and opened in a standard document viewer.

The dataset underlying the application comprises laboratory measurements from 138 boreholes, recorded across 715 depth intervals and covering 57 contaminant parameters, together with the regulatory limits defined under D.Lgs 152/06. Figure 21 summarises how much of this data the application can reach.

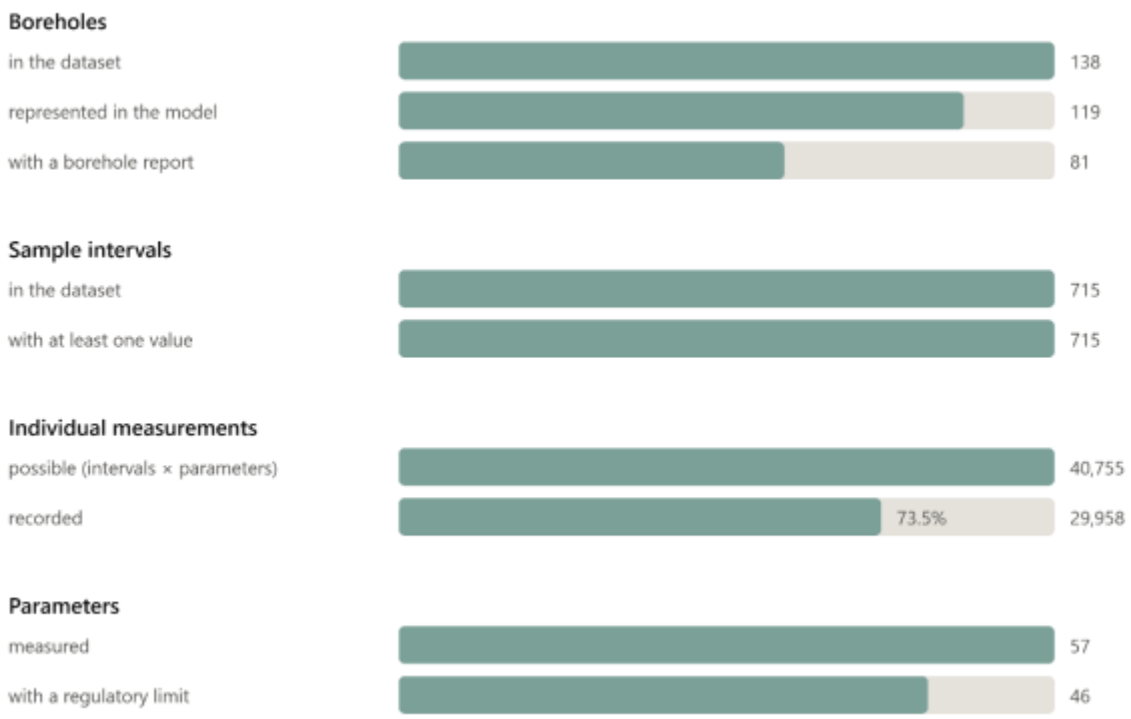

Figure 21 Coverage of the dataset. Each bar is scaled against the total for its category. All boreholes and all sample intervals contain measured data; what varies is the completeness of the analysis within each interval.

No borehole or depth interval is empty; each one contains at least one measured value. Completeness within each interval varies, since not all 57 parameters were analysed on every sample, and the application lists only the parameters measured for the interval selected.

Of the 57 parameters, 46 have a regulatory limit under D.Lgs 152/06 and 11 do not. For the first group the application shows the measured value and states whether it exceeds the limit; for the second it shows the value alone. The 11 without a limit are parameters added by later investigation campaigns, for which no limit is defined in the reference table.

Of the 138 boreholes, 119 are represented as geometry in the BIM model and can therefore be selected in the application. The remaining nineteen come from earlier or separate investigation campaigns and were not included in the model as it was inherited. Their measurements are present in the dataset but cannot be reached, because selection depends on there being a mass element to tap. This is a limitation of the inherited model rather than of the application, which retrieves data by borehole identifier and would present these boreholes as soon as the corresponding geometry were added.

Individual borehole reports, documenting the stratigraphy and field observations at each location, are available for the 81 boreholes of the main survey series and can be exported when that borehole is selected. Where a selected borehole has no report, the export option is not

offered. The complete contaminant table is available as a single document covering the whole site and can be exported at any time.

The implemented application consists of three scenes. The opening scene, MenuScene, loads automatically when the application starts and presents the two modes. Because it contains no model, no data files and no augmented reality session, it appears immediately; the cost of loading the model and preparing the data is deferred to whichever mode is selected. Figure 22 shows the menu as it appears on the device.

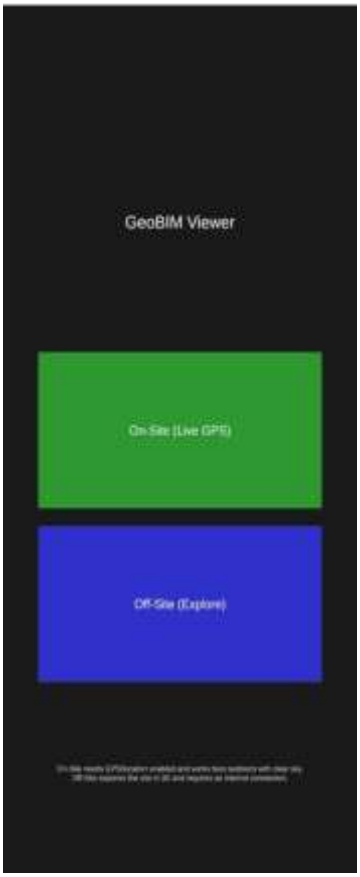

Figure 22 The opening menu, from which the user selects one of the two operating modes.

Selecting a mode loads one of the two remaining scenes, ARViewer for on-site use or ExploreScene for off-site use, each replacing the menu entirely. The two are independent scenes

rather than one scene behaving differently and share only the components that load and present the borehole data.

The sections that follow examine each layer of the application in turn, reporting what was verified, what was observed during testing, and what remained unresolved.

5.2 AR Visualization and Interaction

When the on-site scene was loaded, the augmented reality session started, and the soil model appeared over the live camera view.

The wireframe rendering applied by ApplyWireframe proved well suited to this type of model. Drawing the soil masses as outlines rather than solid volumes kept the internal structure visible, so that boreholes behind the outermost masses could still be seen and selected. Had the model been rendered as solid geometry, only the outer surface would have been reachable, which would have defeated the purpose of visualising the subsurface.

Collider generation by AddColliders worked without manual intervention. Every mass in the model became selectable despite the number of elements involved, confirming that generating colliders in code was the appropriate choice at this scale. Assigning them by hand would have been slow and prone to omissions, and any element missing would have appeared visible but unresponsive to the user.

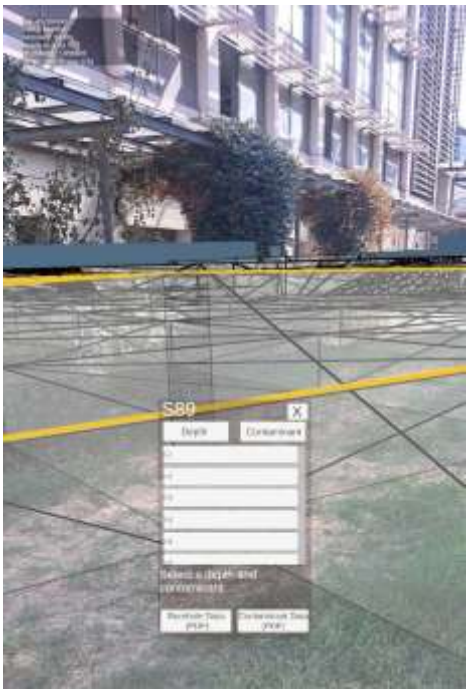

Figure 23 On-Site mode using test coordinates

Selection behaved as intended. Tapping a mass selected the intended borehole, and the selected mass was switched to the highlight material so that the current selection was unambiguous. Returning the previously selected mass to wireframe before highlighting a new one ensured only one borehole ever appeared selected at a time. As shown in figure 23, The soil model is drawn as a wireframe over the live camera view, so that boreholes behind the outermost masses remain visible; the selected borehole becomes highlighted.

Performance was measured by recording the frame rate over a session in the on-site mode with the complete model in view. The frame rate averaged approximately 20 frames per second, falling to around 10 at times. This is below the 30 frames per second generally taken as the threshold for comfortable augmented reality use, and the display was visibly less fluid than a native camera application.

5.3 Data Pipeline

The data pipeline worked correctly on the test device. The data was prepared in the background while the application was starting, so the user did not have to wait for it. The DataManager script successfully loaded both the measurement and threshold files and reported the number of boreholes and contaminant parameters found. This confirmed that the files were located and parsed correctly. Since the loading happens while the augmented reality session is being initialized, it did not introduce any noticeable delay.

On Android, the files had to be loaded through a web request instead of being accessed directly from the device's file system. This is because files included in the application are stored inside a compressed package and cannot be accessed using normal file operations. Therefore, this was a platform requirement rather than a design choice. Both data files were successfully loaded using this approach.

The nested structure used to organize the data also worked well. The measurements were grouped first by borehole, then by depth, and finally by contaminant when the files were loaded. This meant that the application could retrieve the required values directly instead of searching through the entire file each time. As a result, the information was displayed without any noticeable delay. Preparing this structure once at startup, rather than performing a new search for every request, was important for keeping the panel responsive.

The data cleaning performed during the loading process also worked as expected. Empty entries and values marked as not measured were removed, meaning that the panel only displayed actual measurements instead of empty or irrelevant rows. The numerical values were also parsed using a fixed format. This ensured that decimal numbers were interpreted consistently, regardless of

the regional settings of the device, avoiding a potential source of errors that could otherwise go unnoticed.

The data-handling layer is shared by both application modes (Off-site and On-site) without modification. The same scripts are used to load the measurements and thresholds, and the same panel is used to display them in both the on-site and off-site scenes. Regardless of the mode, selecting a borehole follows the same process and produces the same type of information. This includes the comparison of measurements with their regulatory limits and the generation of the corresponding reports.

The selection and display process also worked as intended. Selecting a borehole opened its information panel, selecting a depth interval displayed the contaminants measured at that depth, and selecting a contaminant showed its corresponding value. The threshold comparison provides the main purpose of this part of the application: each measurement is compared with its regulatory limit and clearly identified as either exceeding or remaining within the allowed threshold. This turns the stored measurement into an immediate and understandable regulatory result.

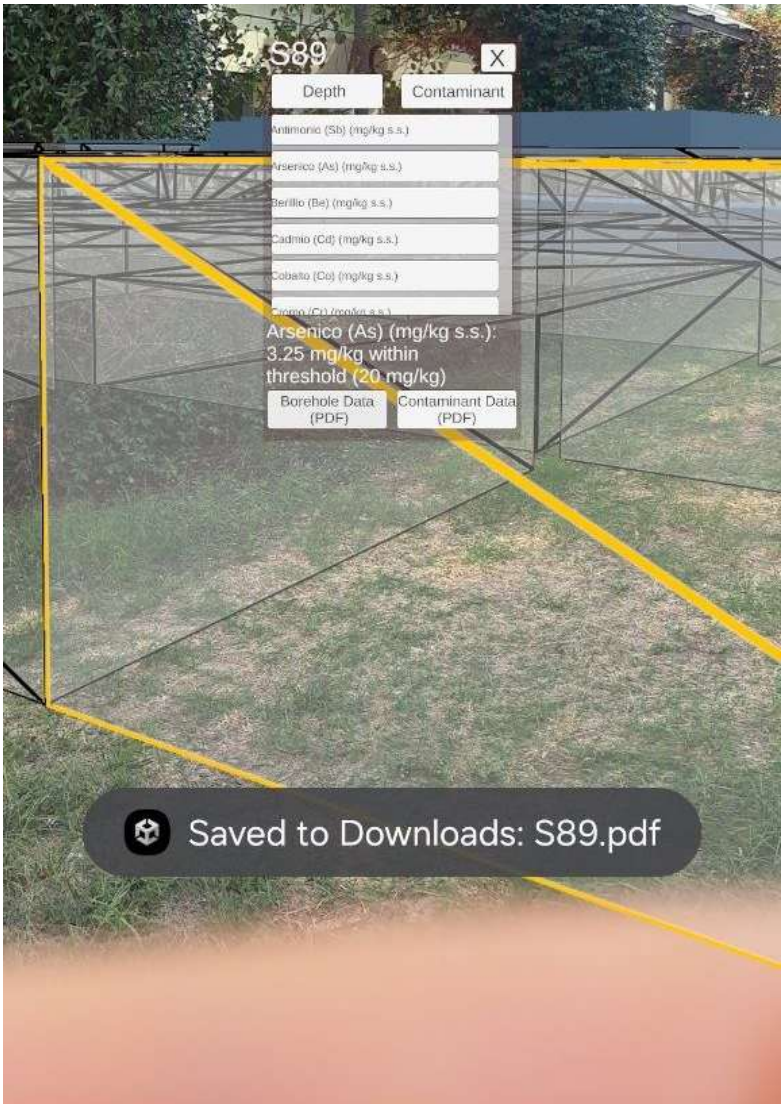

Figure 24 On-Site mode using test coordinates, Borehole S89 selected, with the data panel open and the report exported.

The export function also worked successfully. Both individual borehole reports and the complete contaminant table could be saved to the device's Downloads folder and opened using a standard document viewer. This provides access to the complete field documentation directly from the application. As shown in Figure 24, when a Borehole is selected in On-Site mode. After a depth interval and a contaminant are chosen, the measured concentration will be displayed together with its comparison against the regulatory limit. The notification confirms that the corresponding borehole report has been exported to the device.

5.4 Geospatial Positioning

The geospatial positioning layer was tested outdoors on the device, at a test location with the reference coordinates temporarily set to a point in an accessible open area.

To prepare the geometry for the application, the ModelCenterTool script applied one-time correction by removing the export offset permanently, after that the geospatial library positioned the model without conflict. This confirmed that separating the fixed geometry correction from the live positioning was the correct approach.

Without the opportunity to test on-site, the app was built using a test location as an anchor, allowing for rapid outdoor testing of the satellite positioning pipeline. Once anchored, the model appeared around the user at real scale. As shown in Figure 25, the app utilizes two anchors; the test anchor can be easily updated in the editor without modifying the SiteManager script.

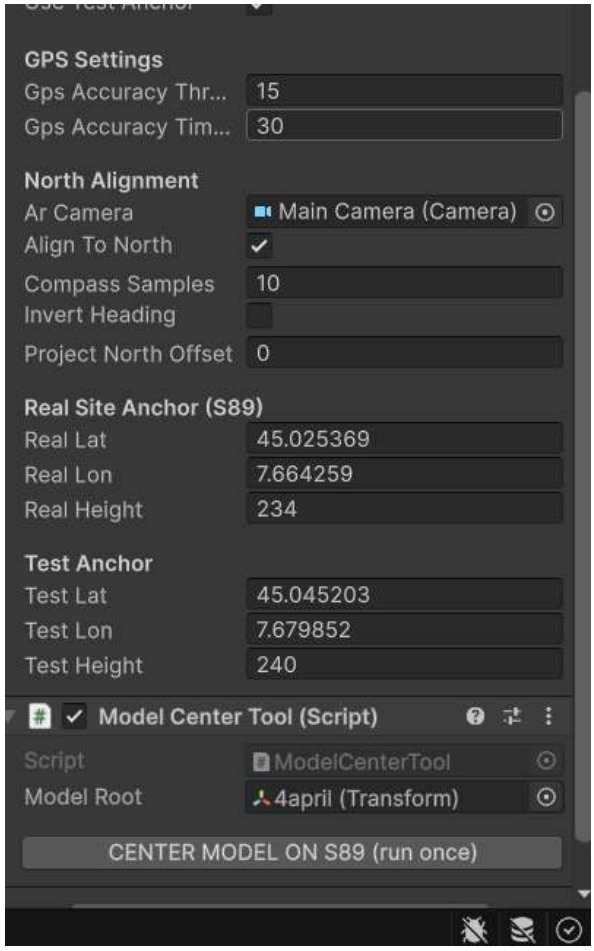

Figure 25 Unity environment, Site manager component, Showing Site anchor and Test Anchor

The top three lines in the panel on top left side of the screen in the on-site mode show the device's current position and its reported horizontal accuracy, updated continuously and they are derived from the GPSDisplay.sc. The accuracy is the receiver's own estimate of the radius within which the true position is expected to lie, derived from satellite visibility and geometry. It is a confidence estimate rather than a measured error. On the other hand, the bottom two lines are produced by SiteManager. These two are status displays that record the position accepted at startup and the accuracy reported at that moment. These values are fixed for the session, because the model is anchored once and the augmented reality tracking maintains its position thereafter. Continuous re-anchoring would cause the model to shift with every fluctuation in the satellite fix. The acceptance rule is defined by two parameters. The application waits until the **reported accuracy** falls below 15 m and waits **at most 30 seconds** for this to occur, after this time the best available fix is used.

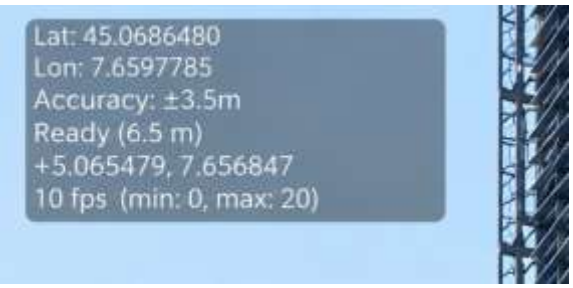

Figure 26 On-Site mode at the test location. The status display shows the accepted position and its reported accuracy; the live readout above shows the current device position

With the anchor set to the test coordinates, the model appeared around the user at real scale and remained fixed as the user moved. Walking approximately thirty steps in two directions, the model held its geographic position rather than travelling with the device, confirming that the placement is genuinely anchored to the world.

The accuracy of that placement is bound by the satellite fix. In the session shown in Figure 23 the position was accepted at a reported horizontal accuracy of 6.5 m, and the model was placed accordingly. The live readout displayed in the same figure shows the device subsequently reporting 3.0 m.

Positioning accuracy varied considerably with conditions, and this was the most significant finding of the testing. In tests in more open conditions the reported accuracy settled between one and a half and three meters, and the model appeared essentially at the user's position. The accuracy threshold described in Section 4.7 is the direct response to this variability: it prevents the application from silently accepting a poor position, at the cost of a short wait at startup.

It should be noted that the reported accuracy is the device's own estimate and is not independently verified. In one test the position was accepted at a reported accuracy of three

meters, yet the accepted coordinates lay approximately eighteen meters from the surveyed anchor. The reported figure is therefore an indication of the device's confidence rather than a guarantee of correctness. The reported accuracy is the device's own estimate and is not independently verified. A repeatability trial quantifying the relationship between reported accuracy and observed variation in position is presented in Section 5.6.

The decision to use the surveyed elevation rather than the satellite altitude reading was also justified in testing, as the vertical placement was consistent between sessions rather than varying with the altitude estimate.

Orientation. Before alignment was introduced, launching the application while facing different directions placed the model at different bearings while its distance from the user remained consistent: turning ninety degrees and restarting rotated the model by the same amount. This identified the tracking system's arbitrary initial heading, rather than the model itself, as the source of the misalignment.

With magnetometer-based alignment applied, the model's orientation no longer depended on the direction the device was facing at launch, and repeated launches from different orientations placed the model consistently. Reaching a working implementation required three refinements: applying the rotation only after the tracking system reports a stable state, discarding readings the operating system flags as invalid, and averaging the remaining readings as directions rather than as numbers. The direction of the correction was established empirically rather than derived.

5.5 Off-Site Exploration Mode

The off-site mode worked as intended during testing. When it was selected from the main menu, the scene loaded the streamed photorealistic terrain and placed the soil model in its correct geographic position. This allowed the site and its surroundings to be explored freely from different angles without the user needing to be physically present at the location.

The decision to keep the two modes in separate scenes also proved useful in practice. Since the augmented reality scene contains no code related to the off-site exploration mode, and vice versa, changes made to the positioning behavior of the on-site scene during testing did not affect the off-site mode. The off-site scene continued to work normally throughout the development process. At the same time, the components shared between the two modes, particularly the data loading system and the borehole information panel, behaved in the same way in both scenes without requiring any changes.

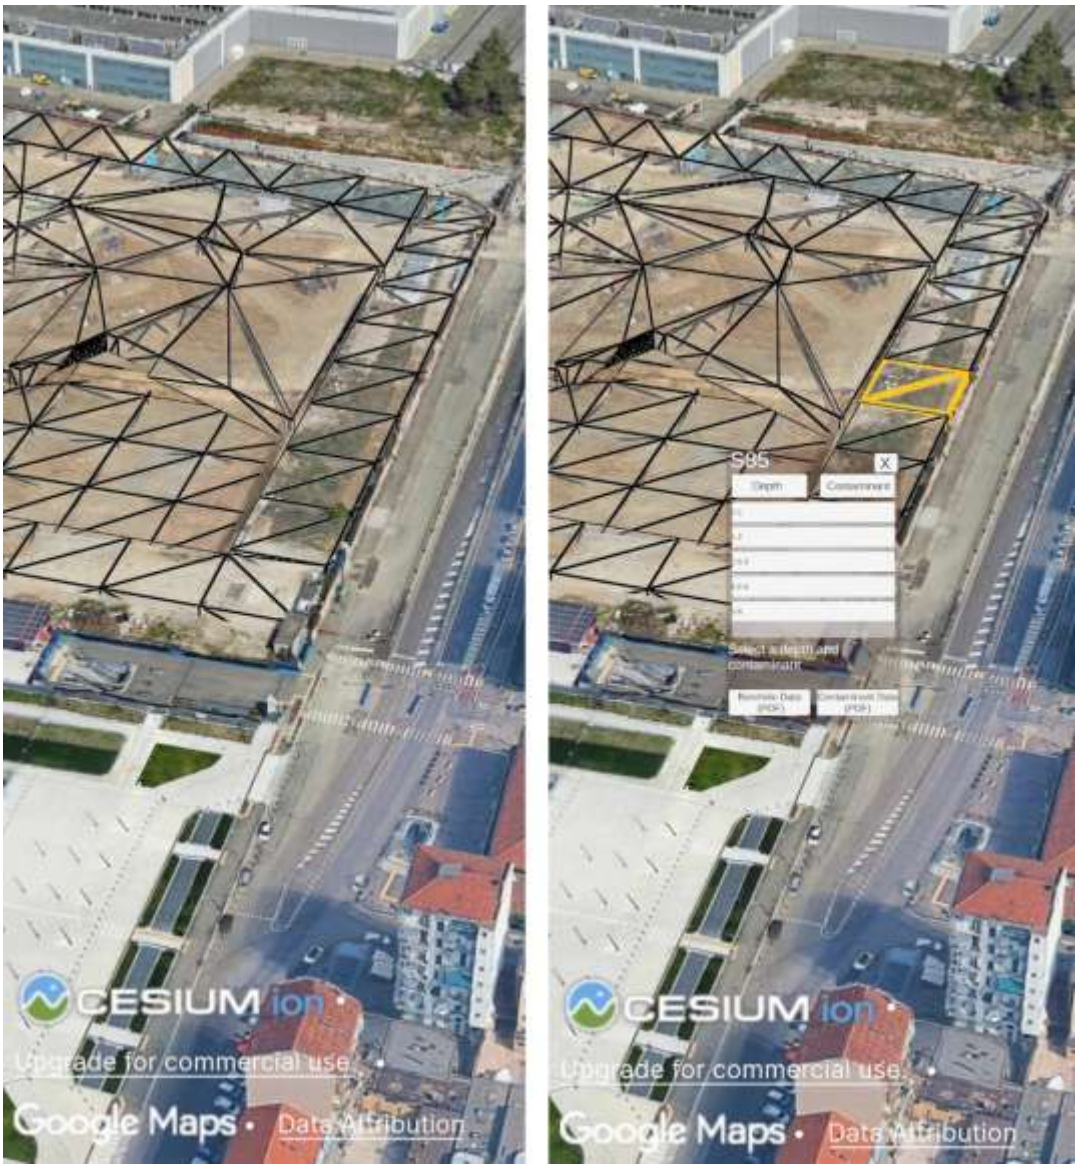

Figure 27 Two Views of the offsite mode, model visible in wireframe and the borehole selected in highlight the borehole panel functions in this mode

As shown in figure 27, the selection behavior is shared by both modes. The model, wireframe rendering, colliders, and highlighting are also present in the off-site scene, allowing boreholes to be selected in essentially the same way. Since the data-handling layer does not depend on how a borehole was selected, it can be reused in both application modes. The two modes have different visualization and interaction methods, but they can rely on the same underlying data-handling process.

The main difference between the two modes is how user input is handled. In the off-site mode, the same touch input is also used to orbit and move the camera. For this reason, a selection is only triggered when the user releases the finger, and the finger has moved very little. This prevents normal camera movements from being interpreted as borehole selections. The two modes therefore use separate input scripts, but the underlying logic for identifying boreholes, filtering out other model elements, and applying the highlight is shared between them. The components specifically related to augmented reality, including the AR session and camera tracking, are only used in the on-site scene.

Chapter 6 Discussion

This chapter returns to the research questions set out in Chapter 1, considers each result against what has been established in earlier work, and reflects on what the application means for practice on a construction site.

6.1 Integration of geometry and records

The application resolves each selected borehole to its laboratory record and presents the corresponding measurement against its regulatory limit. Of the 138 boreholes in the dataset, 119 are represented as geometry and are therefore reachable; of the 57 measured parameters, 46 have a regulatory limit and can be assessed automatically. Both shortfalls originate in the material inherited rather than in the application, which resolves records by identifier and would handle the remainder as soon as the corresponding geometry or limits were supplied.

This is worth setting against the earlier studies in the research chain. The first produced a geological model in specialist software; the second rebuilt it in a BIM environment where it could carry attributes and be shared. In both, the data and the geometry remained within one authoring environment, where correspondence between them is maintained by the software itself. Moving that correspondence onto a mobile device requires it to be re-established explicitly, through identifier matching across separate files, and it is at that boundary that silent loss occurs. The lesson is not specific to this application: any pipeline delivering BIM-linked data to a field device crosses the same boundary and needs to verify coverage rather than assume it.

6.2 Alignment accuracy

In open outdoor conditions with a clear sky at the designated test location, the reported satellite positioning accuracy settled between approximately one and a half and three meters. The virtual model was successfully rendered at the user's position, and this placement was maintained while the user moved. However, in less favorable conditions, the device took substantially longer to reach a usable fix, and in some cases did not reach the accuracy threshold within the waiting period at all. The accuracy threshold added in response helps but does not entirely solve the problem. It filters based on the device's own estimate of its accuracy, which is a statement of confidence rather than a true measurement. While the application can refuse positions the

device admits are poor, it cannot detect a position the device is confidently wrong about, as it has no independent reference against which to check.

Orientation presented a separate structural challenge. Augmented reality tracking measures gravity and therefore knows which way is up but has no knowledge of north. Its sense of forward direction comes from wherever the phone happened to be pointing when the application was launched. To resolve this, the model is anchored using a single reference point, which fixes the position but leaves the heading to be obtained separately from the device's magnetometer. This correction ensures the alignment is good enough to point the model roughly the right way, but it cannot match the surveyed geometry precisely. Deriving the heading geometrically from multiple positions measured on site would give an error of roughly one degree over the distances involved here, whereas relying on the compass yields an error of ten to twenty degrees. The manual rotation adjustment of one degree provided in the interface is a practical response to this discrepancy rather than a mathematical solution.

The single-point anchoring approach also leaves no redundancy: an error in that one measurement shifts the entire model, and nothing in the system can detect it. A multiple-anchor method would resolve this, but it requires measuring reference points directly at the target site, which was not possible during the development period.

Although testing was restricted to a proxy location, these observed figures are highly consistent with the published behavior of smartphone satellite positioning. Comparable accuracy of a few meters is typically reported under favorable conditions, while considerable degradation is observed in built-up environments (Merry & Bettinger 2019). This identical conclusion has been reached in parallel research.

Consequently, this limitation is recognized as an inherent characteristic of the hardware approach rather than a flaw in this specific software implementation. It is concluded that no refinement of the software would substantially improve this accuracy.

Therefore, the practical capabilities of this application must be stated plainly. Because boreholes at the target site are situated fifteen to twenty meters apart, an accuracy of a few meters allows the correct general area, and usually the correct borehole, to be reliably identified. However, permitting a precise point to be marked on the ground is not possible. Ultimately, the application is positioned as an instrument for orientation and data interpretation, rather than a tool for precise site setting-out. This boundary is acknowledged as a permanent property of the sensing approach rather than a temporary software shortcoming.

6.3 Delivery of consumer hardware

The complete model and dataset ran on a mid-range smartphone, with no additional equipment and no modification to the device. In the on-site mode the application operates entirely offline: the model, the measurements and the borehole reports are bundled within it, so even document export functions without a network. This matters for the intended use, since connectivity on a construction site is frequently poor. On the other hand, the off-site mode requires an internet connection because it streams terrain in increasing levels of detail rather than storing it. Consequently, it cannot serve as a fallback in the field. This streaming also impacts visual delivery on the test device; zooming out or changing viewpoints quickly left the terrain as flat, untextured blocks for a few seconds while surrounding buildings filled in, though the pre-bundled soil model itself was never affected.

The primary cost of using consumer hardware is rendering performance. The recorded frame rate fell to between 10 and 20 frames per second with the full model in view. This is below the 30 frames per second generally regarded as the threshold for comfortable augmented reality use, making the display visibly less fluid than a native camera application. The cause is the high number of draw calls generated by the rendered elements: specifically, one mass per borehole and depth interval, each carrying its own wireframe material and an individual mesh collider.

Reporting this measured cost, rather than a qualitative impression, seems the more useful contribution to this study. It establishes exactly what a mid-range consumer device of this class can carry and identifies where specific optimizations—such as combining geometry or rendering only the elements immediately near the user—would be required to carry more.

6.4 Implications for site management

The wider significance of this application concerns *when* and *where* subsurface information is consulted, rather than simply how it is displayed. At present, an investigation report is examined during planning and rarely afterwards. Decisions taken during excavation—about what to separate, what to stockpile apart, and where additional precautions apply—are typically made from 2D drawings and from a recollection of what the report contained. The information exists, but the link between a measured tabular value and a physical location must be reconstructed mentally by whoever is reading it.

By using Augmented Reality to connect the GeoBIM model and its tabular data directly to the real world, this work makes a significant methodological contribution: it changes the moment of consultation. Furthermore, it demonstrates that anchoring a model in physical reality is a mechanism for fulfilling the Level of Information Need (LOIN) for field personnel. Traditionally,

LOIN is increased by adding more geometric detail or semantic attributes within the software. However, for professionals working directly on a remediation site, the most critical missing information is the immediate physical context. By overlaying the digital data onto the site, the physical environment itself becomes part of the model's information structure. Reality acts as the ultimate spatial reference frame, delivering a level of actionable information that no desktop environment can provide.

This targeted fulfillment of LOIN has direct consequences for site operations. Under D.Lgs 152/06, the handling of excavated material depends on the concentrations measured within it. That determination is most useful while the material remains separable in the ground, rather than after it has been moved and mixed. Likewise, knowing that a specific area exceeds a limit for a particular contaminant has immediate consequences for the protective equipment and working methods required by the people standing over it, decisions that must be managed during active work, not just during planning.

The economic argument for this approach is straightforward. Site investigations are highly expensive, yet their output is used intensively during design and then largely set aside during construction. The earlier studies in this research chain made that data structured and shareable; this thesis makes it consultable exactly where the ground is. The investigation gains a longer, more active working life at almost no additional cost, since both the data and the model already exist.

Two qualifications belong to this. First, as established, the application supports orientation and data interpretation, not precise setting-out. Second, it does not replace the traditional investigation report as the formal legal record but rather acts as an interpretive lens through which that record can be utilized during active excavation.

6.5 Limitations of the Study

The technical and hardware constraints encountered during development, such as satellite tracking accuracy, data completeness, and rendering performance, have been addressed alongside the specific research questions they affect. The remaining limitations of this work relate to the methodological scope and the validation of the application itself.

Scope of testing and validation Because access to the target site was not available during the development period, all positioning and orientation testing were carried out at a proxy location. Testing in this manner successfully verified the software pipeline, confirming that the model is placed geographically and holds its position as the user moves. However, it cannot establish the accuracy achievable under the specific environmental conditions of the actual site, which has its

own sky visibility, surrounding structures, and magnetic environment. Consequently, the definitive validation step—standing at a known physical borehole to confirm that the virtual model aligns with reality—has not yet been performed.

Lack of end-user evaluation While the application successfully demonstrates the technical feasibility of delivering BIM-linked records to a mobile device, it has not been evaluated with the site professionals who would use it. The economic and operational arguments for consulting subsurface data in the field are straightforward, but the tool's practical usefulness, usability, and true integration into daily construction site management remain to be established through formal user testing.

Platform and hardware variability Finally, development and testing were restricted to a single mid-range Android device. Although the underlying frameworks support cross-platform deployment, no iOS build was generated, and no testing was conducted across a wider range of Android hardware. Because satellite positioning capabilities, sensor accuracy, and rendering performance are known to vary significantly between mobile devices, the performance baselines established in this study cannot be universally generalized without broader hardware testing.

Chapter 7 Conclusion

The goal of this study performed on the site of “Citta della salute” in Turin was to create a working mobile application that brings geotechnical subsurface data into the field. The platform used for developing this Augmented Reality application was Unity. The BIM model of the soil and a set of laboratory measurements were taken as inputs and, the application places that model at its real location on a phone screen with its data reachable.

The application operates in two modes. On site, it displays the soil model over the live camera view at real coordinates, anchored to a surveyed borehole and aligned to true north, so that a user can walk across the site and see the subsurface beneath them. In the explore mode, the application shows the same model within streamed photorealistic terrain, navigated freely, which allows the site to be examined and presented without being there. In both modes a borehole can be selected to retrieve the contaminant measurements recorded at each depth, compared automatically against the regulatory limits of D.Lgs 152/06, with the associated field documentation available for download. Testing confirmed that the placement is genuinely geographic.

The main contribution of this study is providing data access for the engineers. This work is the third in a chain, the first study produced the geological model, the second rebuilt it as a data-bearing BIM model. This work closes the gap between the outcome of those studies and the site they represent especially without specialist equipment.

The outcomes of this research can be the starting point of several further studies. The clearest direction follows from the limitation of position and heading that shaped this work. position come from satellite positioning and the heading comes from a magnetometer, and they are both approximaed.one possible study path could be anchoring the model to two or more positions measured on site, rather than to a single reference point, allowing the heading to be derived geometrically rather than magnetically. It reduces the error by roughly an order of magnitude across the distances involved at this site. This path needs no new hardware, only site access. The second study path is to replace sensor-based positioning altogether with visual positioning, which determines where the device is and which way it faces by matching the camera view against a reference database of imagery. Google's Visual Positioning Service, accessible through the same augmented reality framework used here, reports accuracy of around five meters in position and five degrees in rotation.

In the end it is worth mentioning three broader questions. The application was tested on a single device, so nothing is known about how the performance will vary across hardware. The second question is, no evaluation with intended users was carried out. in other words, the professionals

who must use this application in the end may have a more realistic vision on what how useful this augmented reality practically is. Finally, this application implemented a model prepared by hand in an earlier study. For the approach to be practical beyond a single site, the path from geotechnical investigation to field-ready model would need to be repeatable and largely automated.

Appendix A:

Project Repository The mobile augmented reality application described in this thesis was developed using [e.g., Unity and AR Foundation]. All associated scripts, 3D assets, configuration properties, and datasets are openly available for review. The complete project repository can be accessed via GitHub at: [Insert your GitHub URL here].

References

Amin, K., Mills, G. and Wilson, D. (2023) 'Key functions in BIM-based AR platforms', *Automation in Construction*, 150, 104816.

Aschieri, D.L.D., Khosravi, M. and Carlotto Cano, V. (2026) 'Bridging above and below: a multi-criteria evaluation of GeoBIM modelling workflows for contamination assessment—the PSRI case study (Italy)', *Journal of Environmental Management*, 401, 128959.

Borkowski, A.S. (2023) 'A literature review of BIM definitions: narrow and broad views', *Technologies*, 11(6), 176.

Hsieh, C.-C., Chen, H.-M. and Wang, S.-K. (2023) 'On-site visual construction management system based on the integration of SLAM-based AR and BIM on a handheld device', *KSCE Journal of Civil Engineering*, 27(11), pp. 4688–4707.

Htet, H.T., Jung, H.-Y., Lee, Y.-J. and Park, M.-W. (2026) 'Real-time AR-based information retrieval for dynamic objects on construction sites', *KSCE Journal of Civil Engineering*, 30, 100572.

International Organization for Standardization (2018) *ISO 19650-1:2018 Organization and digitization of information about buildings and civil engineering works, including building information modelling (BIM) Information management using building information modelling Part 1: Concepts and principles*. Geneva: ISO. Available at: https://www.iso.org/standard/68078.html (Accessed: 10 August 2026).

Kassou, F., Bouayach, N. and Rguig, M. (2025) 'GeoBIM for infrastructure geotechnics: parametric and generative design', *Transportation Infrastructure Geotechnology*, 12, 290.

Khorrami Shad, H., Feng, Z., Paes, D., Yiu, T.W. and Lovreglio, R. (2025) 'Augmented reality applications for the excavation industry: locating and protecting underground utilities', *Journal of Pipeline Systems Engineering and Practice*, 16(1), 04024072.

Kim, H.-S. and Chung, T.-K. (2026) 'Integrating 3D geotechnical characterization and Geo-BIM for geohazard assessment in soft ground', *Developments in the Built Environment*, 25, 100863.

Li, M., Feng, X., Han, Y. and Liu, X. (2023) 'Mobile augmented reality-based visualization framework for lifecycle O&M support of urban underground pipe networks', *Tunnelling and Underground Space Technology*, 136, 105069.

Raitviir, C.-R. and Lill, I. (2024) 'Conceptual framework of information flow synchronization throughout the building lifecycle', *Buildings*, 14(7), 2207.

Rajadurai, R. and Vilventhan, A. (2025) 'Integration of building information modeling, geographic information system, and augmented reality for visualization and management of multiple underground utilities', *Journal of Pipeline Systems Engineering and Practice*, 16(1), 04024058.

Safaan, M.H., Metawie, M. and Marzouk, M. (2025) 'Subsurface utility detection and augmented reality visualization using GPR and deep learning', *Automation in Construction*, 176, 106299.

Sanseverino, A. and La Placa, S. (2026) 'Open-oriented algorithmic approach for BIM modelling of complex historic water infrastructure', *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, XLVIII-2/W12-2026, pp. 431–438.

Schiavi, B., Havard, V., Beddiar, K. and Baudry, D. (2022) 'BIM data flow architecture with AR/VR technologies: use cases in architecture, engineering and construction', *Automation in Construction*, 134, 104054.

Shekargoftar, A., Taghaddos, H., Azodi, A., Nekouvaght Tak, A. and Ghorab, K. (2022) 'An integrated framework for operation and maintenance of gas utility pipeline using BIM, GIS, and AR', *Journal of Performance of Constructed Facilities*, 36(3), 04022023.

Sivakumar, A., Melina J V, T. and Suresh, E.S.M. (2025) 'BIM-GIS integration for efficient management of urban underground stormwater networks', *Proceedings of the Institution of Civil Engineers – Smart Infrastructure and Construction*, 178(3), pp. 186–199.

Skabek, K., Rola, D. and Zamarski, W. (2026) 'Combining indoor and outdoor positioning for navigation in AR environments', *Computer Assisted Methods in Engineering and Science*, 33(1), pp. 7–25.

Brata, K.C., Funabiki, N., Panduman, Y.Y.F. & Fajrianti, E.D. (2024). An Enhancement of Outdoor Location-Based Augmented Reality Anchor Precision through VSLAM and Google Street View. *Sensors*, 24(4), 1161. DOI: 10.3390/s24041161

Fonsati, A., Cosentini, R.M., Tundo, C. and Osello, A. (2023) 'From geotechnical data to GeoBIM models: testing strategies for an ex-industrial site in Turin', *Buildings*, 13(9), 2343.

Mardia, K.V. & Jupp, P.E. (2000). *Directional Statistics*. Chichester: John Wiley & Sons.

Merry, K. and Bettinger, P. (2019) 'Smartphone GPS accuracy study in an urban spatial environment', *PLOS ONE*, 14(7), p. e0219890.

Koka, E., Bozheku, E., Venditti, C., Picaro, R., Bashmili, J. and Veshi, D. (2025) 'Legal BIM in Italy: An Administrative Case-Law Approach', *Bratislava Law Review*, 9(1), pp. 113–134.

Pan, N.-H. and Isnaeni, N.N. (2024) 'Integration of Augmented Reality and Building Information Modeling for Enhanced Construction Inspection—A Case Study', *Buildings*, 14(3), p. 612.

Mužík, F. and Běloch, L. (2026) 'Using a Visual Positioning System for a Geolocated Visualization of an Archaeological Site in Augmented Reality', *ISPRS International Journal of Geo-Information*, 15, p. 219.

Barrile, V., La Foresta, F., Calcagno, S. and Genovese, E. (2024) 'Innovative System for BIM/GIS Integration in the Context of Urban Sustainability', *Applied Sciences*, 14(19), 8704.