

Automated Generation of Accurate 3D Building Interiors Models: Lessons Learned and Challenges

Aaron Katzman
Dignitas Technologies
Orlando, Florida
akatzman@dignitastechnologies.com

Joseph Moran, Jr.
Dignitas Technologies
Orlando, Florida
jomoran@dignitastechnologies.com

ABSTRACT

Modeling a building interior presents a complex problem to the Modeling and Simulation community. The traditional approach to modeling building interiors, by hand using CAD or other 3D modeling tools, is cumbersome and restrictive. This is especially true for multi-story buildings where there are numerous unique floor layouts or buildings with unusual geometry, as they require additional time and resources to model properly. Without an interior paired with the exterior model, there are limitations to the situations and scenarios where a building model can be used. Every data collection method presents unique advantages and disadvantages impacting its value to the automated processes employed to create an interior model.

Since future civilian emergency response and military engagements are increasingly likely to occur in dense urban environments, there is a need to generate accurate representations with known confidence of both the exterior and interior of a building. This need is not just exclusive to defense applications for mission rehearsal but also to emergency first responders and cultural resources. A variety of formats are available to support training, operational, and intelligence communities. Current data collection methods and automated data processes that generate 3D synthetic models are not mature enough to handle the complexities of dense urban environments on their own.

Our research suggests that a broad spectrum of data formats, techniques, and algorithms can be employed to create geo-specific interiors models, rather than to rely on a single automated collection and processing mechanism. In this paper, we discuss how a combination of commercial and open-source software, organized in a pipeline, incrementally improves and optimizes interior geometry modeling.

ABOUT THE AUTHORS

Joseph Moran Jr. (JJ) has over 15 years of experience in a variety of sectors; ranging from Geographic Information Systems, Commercial Game Development, and Modeling and Simulation. Through his career, Mr. Moran has managed and influenced the development of multiple products, including the Conform application and multiple AAA game titles. He is currently the technical lead for the Navy's Validate project which seeks to measure the level of compatibility for combined arms training across the Navy's collection of flight simulators. Mr. Moran is also providing expertise and guidance to the dense urban environments research being performed under the MUSE project.

Aaron Katzman has over 10 years of experience mainly within the simulation industry, especially from Virtual and Live fire trainers. In his career, Mr. Katzman has led and developed a wide range of projects within the defense industry from weapons and combat training to aerospace and synthetic networked environments; mainly specializing in 3d math, physics, and modeling. He is currently providing expertise and guidance, along with research, performed under the MUSE project involving dense urban environments.

Automated Generation of Accurate 3D Building Interiors Models: Lessons Learned and Challenges

Aaron Katzman
Dignitas Technologies
Orlando, Florida
akatzman@dignitastechnologies.com

Joseph Moran, Jr.
Dignitas Technologies
Orlando, Florida
jomorán@dignitastechnologies.com

INTRODUCTION

In the past, maps were sufficient for giving context to the end user regarding the place and space of objects. As technology has advanced, so has the concept of what a map can be used for and it has evolved from being a static 2D product to a **living** 3D model, including live crowd sourced data, embedded material attribution, and 3D models. Building a 3D model of a location has been made easier with laser scanning and photogrammetric techniques; however, they are usually missing a key element – the interior of those structures.

A review of open opportunities from Defense and Civilian Government Agencies on GovWin between 2018 to 2021 revealed a 67% increase of proposals involving building interior information over three years. Both commercial and non-commercial industries have attempted to solve this problem in a myriad of ways for an ever-growing list of applications - for example: ordinance effects, emergency management, or cultural heritage preservation. However, the solutions are often reliant on single pipelines with data formats that have distinct advantages and limitations. We do not suggest that one pipeline or file format is unequivocally better than the other. Rather, we propose that multiple pipeline approaches can be leveraged to achieve accurate, well-formed, geo-specific interior geometry.

In this paper we will evaluate broader industry needs of geo-specific 3D interior modeling, some of the past methods and research conducted for building interior model generation, and how pipelines and file formats can be consolidated for generating accurate 3D interior geometry for a wide range of applications and use-cases.

Defense Industry Needs

There are several industry viewpoints regarding the importance of geo-specific, modeled building interiors. Since the early 2000s, when the U.S. began conducting urban operations in Afghanistan and Iraq, geo-specific (and geo-typical) interiors modeling has remained a desirable operational capability - specifically for new concepts research, munitions effects research, and mission rehearsal (Mann & Eifert, 2006).

Most recently, the Department of Defense (DoD) has placed a renewed emphasis for geo-specific building interiors modeling for training purposes in the Synthetic Training Environment (STE) program, which is set to become the next generation in live, virtual, constructive, and gaming environments, as well as provide data for training and mission rehearsal exercises for the U.S. Army (U.S. Army, 2023). According to the STE's current Statements of Work for integrated providers, there is a capability need for generating specific, high fidelity 3D building models for live virtual training scenarios. This capability need is being explored by the STE's One World Terrain (OWT) (Tilton, 2021) and Training Management Tool datasets.

First Responder Industry Needs

First responders have noted a need for accurate building interiors to promote heightened safety and improved efficiency in emergency response since the early 2000s. To rectify this need, the Building Tactical Information System (BIM) for Public Safety Officials project was proposed. The proposed system would aid first responders by providing live building data (e.g., fire detection system positions within a building) and the commercial building's interior layout (Holmberg, Davis, Treado, & Reed, 2006). This research was promising as it provided a set of standard BIM information to first responders, but it was ultimately not successful at being adopted.

During the Orlando Pulse Nightclub shooting in 2016, to save hostages and neutralize the shooter, first responders used an armored vehicle to breach a wall “...between the two bathrooms, but they had picked the wrong spot” (Goldman & Berman, 2016). This unfortunately removed the element of surprise and gave the shooter advanced warning of an impending breach on his position. This tactical error ultimately cost the lives of several hostages as the shooter began firing upon them immediately afterward.



Figure 1. FBI Agent investigating the scene of the Pulse Nightclub Shooting (Raedle, 2016)

From this instance, it is clear that U.S. first responders will gain a tactical advantage by having accurate geo-specific 3D interior models accessible at the point of need.

Cultural Resource Industry Needs

From a cultural resource standpoint, building interiors data collection and generation remains an important capability needed at different historic sites. There is a broad acknowledgement from many cultural resource sites that either man-made or natural disasters can and will eventually impact a historic site. The importance of retaining records of these historic interiors has been used for “...research, preservation, and future restoration...” efforts at several sites (Bent, Pfaff, Brooks, Radpour, & Delaney, 2022).

Most recently, 3D data that was collected by a research team prior to the fire that devastated the Notre Dame Cathedral in Paris is being reviewed. By leveraging the data and taking careful measurements, reconstruction of key structural features can be more precisely restored to their original form (Wiesmayer, 2020).

BUILDING INTERIORS LITERATURE REVIEW

Past Approaches for Modeling Geo-typical Building Interiors

In 2010, Merrell et al. proposed a method for generating floorplans using a combination of deep learning and grammars. The architecture combines a convolutional neural network (CNN) with a context-free grammar (CFG) to generate floorplan layouts. The CNN is used to extract features from input images, while the CFG is used to enforce constraints and generate coherent floorplan layouts that could be rendered in 3D. This research demonstrated that given a set of parameters and constraints, an end user could generate a geo-typical building and interior (Merrell, Schkufza, & Koltun, 2010). One of the limiting factors with this research was that it required almost no **real-world** exterior building data to be used as a reference to generate the interior.

GameSim’s Procedural Model Generation Service (PMGS) exemplifies procedural, geo-typical building generation. In 2014, GameSim initially developed this service to procedurally generate geo-typical exterior building models from GIS building footprints for simulation and gaming purposes (Marketwired, 2014). Later, PMGS’ capabilities were expanded to include geo-typical interiors based on a logic algorithm using exterior features such as window and door locations in Moore et al.’s research paper: “Exterior Attribute Extraction and Interior Layout Speculation of 3D Structures” (Moore, Reilly, & Pelham, 2018). The result was a generated exterior that closely matched the real-world and an interior that was visually impressive but not necessarily reflective of the building’s real-world interior configuration.

The Advanced Synthetic Natural Environments (ASNE) project researched and developed for the U.S. Army, worked towards building an automated pipeline to generate underground environments, more specifically subway stations in 2019. As part of the pipeline, the ASNE project leveraged openly available GIS data to identify the location of subway rails, stations, and associated attribution to derive the connectivity and layout of the network. From there, the pipeline used real-world engineering rules to create a set of diagrams, like architectural blueprints, which laid out all the critical components that make up a subway station. Once all the critical components were identified, the pipeline would generate 3D geometry that was integrated with the existing 3D terrain database. The ASNE project demonstrated that procedurally modeling interiors, especially underground environments, was possible. However, like PMGS, the focus was on geo-typical modeling.

Past Approaches for Modeling Geo-specific Building Interiors

In 2002, Pigora, Graniela, and Kraus published a paper at IITSEC titled "A System for the Rapid Capture of Virtual Environments" which presented a Virtual Environment Capture System (VECS) which was designed to quickly capture virtual environments. The system attempted to address the time-consuming process of manually modeling accurate virtual, 3D environments by automating and streamlining the data capture process. The authors discussed various components and techniques used in VECS, including the use of panoramic images, depth sensors, and advanced algorithms. This method was effective for modeling geo-specific interiors and relied on instancing geo-typical models (e.g., monitors and desk chairs) for interior furnishings (Pigora, Graniela, & Kraus, 2002). The main drawback was that it was time-consuming to collect interiors data with their system and required a great deal of manual sensor calibration.

In "Reconstructing Building Interiors from Images" authored by Furukawa et al 2009, the researchers focused on a fully automated method for generating 3D models of building interiors using only 2D images taken from the interiors. This photogrammetric approach combines computer vision techniques with structure-from-motion and multi-view stereo algorithms to extract geometric and textured information from images. In the paper, it also noted specific challenges with collecting interior data such as lighting and *texture poor* surfaces such as uniformly painted walls. They applied their pipeline to four distinct locations, but the most relevant for building interiors were their Residential Home and the more complicated Gallery interior collections (Furukawa, Curless, Seitz, & Szeliski, 2009). These showed promising results for generating interior geometry. However, the geometry still contained erroneous data artifacts in the mesh and the result of the interior outputs did not contain discrete attribution for the interior layout such as walls, doors, and windows.

Miles et al. published a paper at IITSEC 2019 titled "Enhancing Situational Awareness Anywhere in the World with Geospatially Accurate Scene Simulation using Automated 'Real World' Content". The authors provide a brief background on the National Geospatial Intelligence Agency's Field Geospatial 3D (FG3D) program. This program is dedicated to advancing the development of technologies for generating virtual environments with realistic representations of real-world locations. One of the specific focuses of the FG3D program is the development of 3D building models. These building models aim to accurately capture the geometric details and visual appearances of various types of structures found in urban environments leveraging data from diverse data sources such as satellite imagery and light detection and ranging (LiDAR) data (Miles, Creel, Kenton, Abrams, & Wilder, 2019). Although the buildings were captured at a high fidelity, it was limited by interior data availability and what interior data was able to be captured from the exterior.

Conclusions for Past Approaches

After reviewing the work done on building interiors modeling within the last 20+ years, two conclusions can be drawn: 1.) Geo-typical modeling of building interiors has several approaches that can be easily leveraged to create generic 3D interiors. 2.) Geo-specific interior data is totally reliant on either existing data availability (such as blueprints) and/or what can be efficiently collected using LiDAR or photographic sensors. The data collected from these sensors are also constrained to available data storage and total processing time. An additional takeaway is that as it relates to the exterior of the building, the Modeling and Simulation community has successfully streamlined exterior building mesh generation by leveraging the latest point cloud and photogrammetric sensor pipelines.

CONTEMPORARY TECHNIQUES FOR INTERIOR GENERATION

There have been many open-source efforts to advance reading, manipulation, and enhancements within point cloud space. Current approaches benefit from several high-level algorithms and machine learning techniques that, usually with great overhead, feed into one another to clean-up, identify, and produce a final mesh. Current libraries typically used are CGAL, PCL, Yolo, and OpenCV. Each offers a set of helper functions that can be executed in different combinations to process point clouds and generate meshes.

In the next section we will discuss a new method, using a flexible high-level pipeline, which covers all critical steps of building interior processing and mesh generation. One of the major advantages of this pipeline is that all the steps are modularized so they can be reassembled in different combinations to adjust to the data processing needs. Also, new modules can be introduced into the pipeline with little or no negative impact to the overall system. Another advantage of modularization is that each step has well-known inputs and outputs that can be analyzed and saved. Analysis of the inputs and outputs allows us to determine whether the step is executing correctly and exporting valid data. Having the flexibility to save the outputs of each step will allow us to restart the processing at any step in the pipeline without the need to execute each and every service from the beginning. This speeds up experimentation and analysis, leading to a more mature product in less time.

PIPELINE STEPS & DEPENDENCIES

The diagram in Figure 2 is made up of microservices (i.e., modules) that can be configured to run in any order or can be duplicated in the pipeline for additional processing. The order of pipeline execution is from left to right in the diagram, but many of the steps could be executed concurrently or in isolation. Initial inputs into the pipeline are a building interior point cloud (or a mesh) and a point cloud (PC) for the building exterior. If a mesh is given as the interior input, the mesh will be converted to a PC before processing continues. Converting to PCs is necessary because they are easier to manipulate than static geometry.

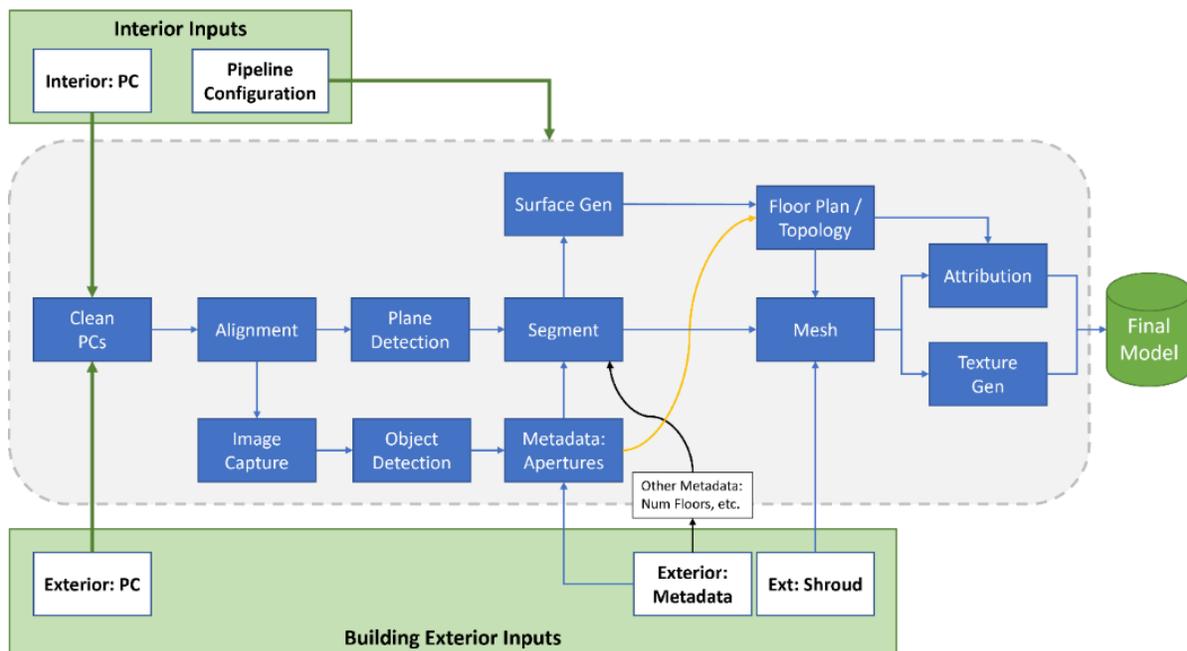


Figure 2. Microservice Pipeline for Interior Generation

The Building Exterior Inputs may include an exterior building PC that has geo-positioning and scaling information that would be beneficial for aligning the interior and exterior PCs. The Exterior Metadata, for example: the minimum and maximum thickness of a wall, may provide additional information at later steps to aid in segmentation of floors,

rooms, and hallways. The Exterior Shroud contains information about the terrain surrounding the building, the shell of the main structure, and other information that can be used to auto-align and fuse the interior with the exterior.



Figure 3. Exterior Shroud Example: Terrain surrounding the building (left), Main structure (center), Terrain with main structure combined (right)

Regardless of the specific details within the processing pipeline, there are some critical concepts that each pipeline must address (Figure 4). First, input data often has errors that must be cleaned (i.e., corrected) to ensure that all follow-up steps have the best data possible. Topology analysis must be executed on the cleaned data to determine the layout of the building interior and how it relates to the exterior

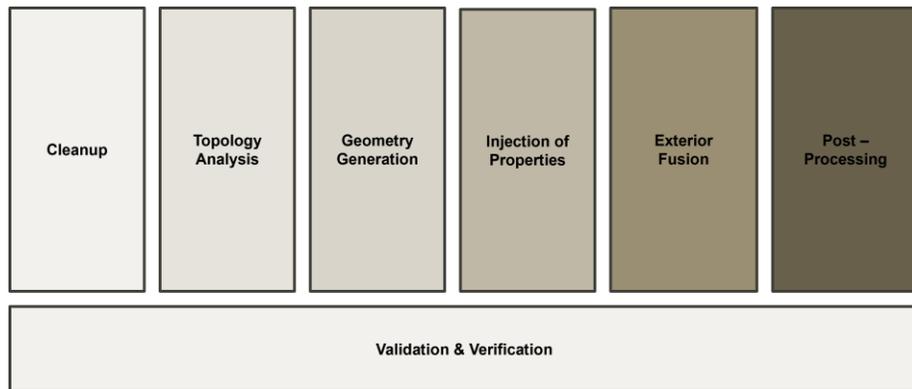


Figure 4. Pipeline Critical Concepts

structure. Seamless geometry must be generated with the correct dimensions and orientation. Outputs must be enhanced with metadata and properties to support the final use cases, including things like material composition of internal structures. Final output must be properly fused with the surrounding external environment to create a seamless world. Finally, post-processing may be necessary to meet very specific needs of the runtime simulation systems.

In the following sections we will describe the various pipeline microservices.

Clean PCs

Clean PCs ingests a raw point cloud and reduces common problems, such as noise, down-samples the PC to lower the point density, and merges (or subdivides) the PC to improve processing speed. This microservice will both start as well as run multiple times throughout the pipeline as any alterations or new outputs may produce unwanted artifacts or noise. Note that this microservice is on an as-needed basis as the input PCs could already be optimized.

The major output of Clean PCs is the calculations of normal that several other microservices may be dependent on. Additional outputs from this are mainly filtered or a reduced version of the original PC.



Figure 5. Raw point cloud (left); Cleaned, post-processed point cloud (right)

Alignment/Georegistration

Alignment takes an existing PC and correctly aligns it to an adjoining PC. This is also where the PC can be merged with exterior PC collects of a building. An additional benefit for PC Alignment is that Georegistration can be easily applied to the interior point cloud once it has been aligned/merged with a properly georegistered exterior building PC. Alignment is scaled to the same units (i.e. feet, meters, etc.) along with the same roll, yaw, pitch and z-level taking its cues from a known, ground-truth PC (ideally an exterior PC).

The output of Alignment/Georegistration is a single, composite PC that has been properly scaled and georegistered.

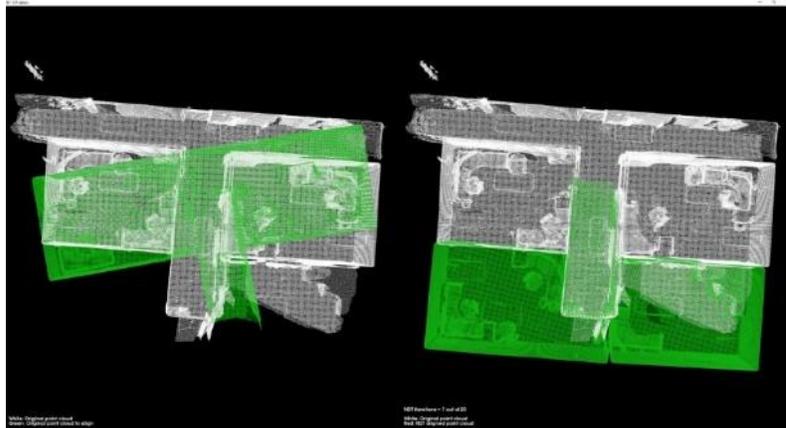


Figure 6. Incorrectly oriented PC shown in green (left); Post alignment of the two PCs (right)

Plane Detection

Plane Detection uses a clean PC to group a set of points into planes based on the normal orientation on those points. Once points are collected into planes, filtering can be executed to eliminate clusters of points that do not constitute large enough features for initial processing. These small features will be considered in later processing steps.

The output of Plane Detections is a calculated index associating each point with a known plane using the calculated normal. It also has the functionality to remove points that are not associated with or belong to any plane.

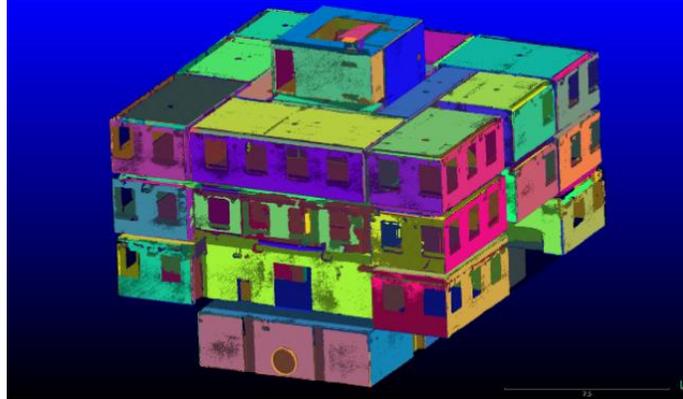


Figure 7. Planarization of a PC mesh. Note that each color represents a different plane

Segmentation

Segmentation is an important service that splits cleaned PCs into rooms, hallways, floors, etc. It does this by first sorting the generated planes into vertical and horizontal groups, then processing those groups to determine which of them make up floors/ceiling. These will be used to segment building stories which will, in turn, be broken down into rooms.

The output of Segmentation, using the precalculated plane index and normals, produces an index that associates each point with a known segment within the building i.e., ceiling, floor, wall, etc. Similarly to Plane Detection, this would allow for functionality to remove any points that are not associated with valid segments and potentially treat them as noise.

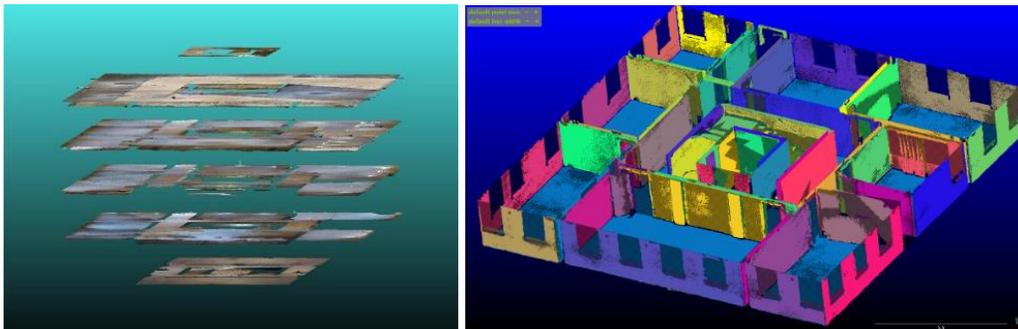


Figure 8. Floor and ceiling segmentation of the PC mesh (left); Building Floor segmented (right)

Surface Generation

Surface Generation uses the segmented rooms and planarized clusters of points to define the polygons that best identify with the surfaces.

The output of Surface Generation can be used to export to different formats and types of 2D or 3D models of the surface. In Figure 9, this Surface Generation output was exported to FLT to display the surface and planes as simple polygons in 3D form.

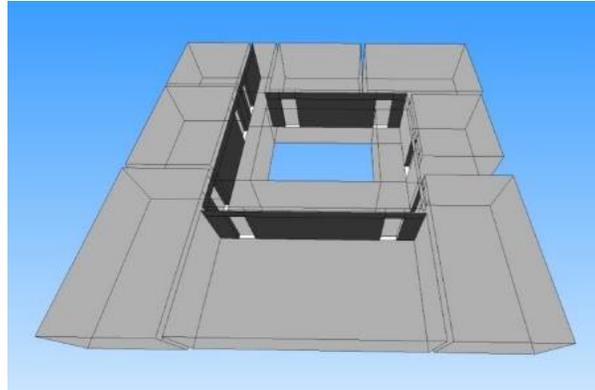


Figure 9. Surface planes derived from PCs

Floor Plan /Topology

Floor Plan/Topology creates a generic representation of the building interior composition. This includes room volumes, aperture definitions, and the relationship between those rooms and apertures. In essence, this defines a structure of interrelated components within the interior.

The output of Floor Plan/Topology is data that associates how different rooms and structures coalesce to make a floor and define the different apertures between these intersections.

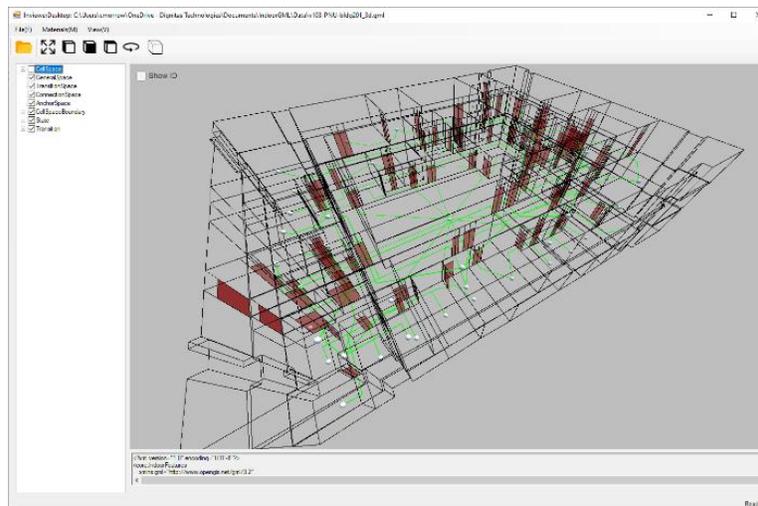


Figure 10. Floorplan of a building with doors and windows identified in red

Mesh Generation

Mesh Generation will use cleaned PCs and other metadata to create detailed meshes for the building interior. Unlike the geometry generated in the Surface Generation step, the Mesh Generation will not ignore “clutter” content, such as chairs, desks, etc. Mesh generation resolution may also be adjusted to capture very high detail content, such as light fixtures, door handles, fire suppression systems, etc.

The output of Mesh Generation is created using the calculated segments, planes, and normal to produce a high-fidelity 3D model that includes more than just simple geometry. For example, Figure 11 shows a high resolution collect of a bench which includes the individual bench leg geometry.



Figure 11: Sample Interior Mesh

Texture Generation

Texture Generation takes the interior surfaces and wraps them with a texture that is sampled from the PC points which are adjacent to any given surface. It requires a cleaned input PC and that all points face in the correct direction.

The output of Texture Generation can be one or several texture files to apply to a 3D Model plane. For example, in Figure 12, the textures that have been embedded using images from the PC, which are then applied to plane surfaces thus giving a realistic representation.

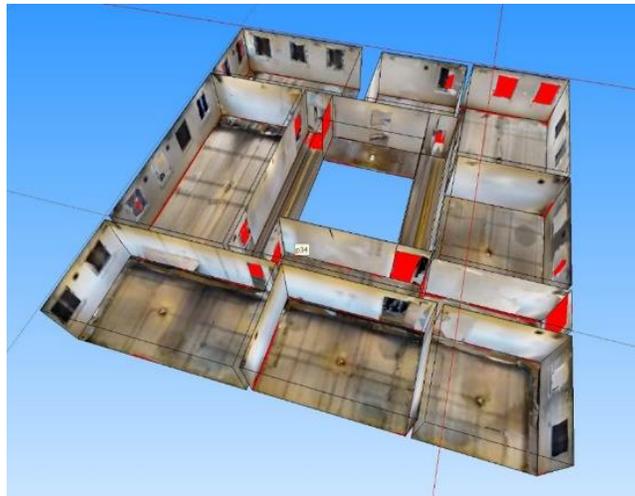


Figure 12. PC derived texture applied to polygon planes

Attribution

Attribution injects metadata into the final output product. Attribution can come from several different sources including a mesh, BIM data, and data derived from different microservices. It is not necessarily limited to common material attribution mapping for an interior surface such as concrete or drywall.

The output of Attribution is clearly defined metadata applied to a mesh or PC that is used for the benefit for 2D capture or object detection such as parameters defining doorways, windows, or other objects. Alternatively, the data can be used for structural associations at a point level.

2D Image Capture

2D Image Capture allows us to take snapshots of the building's interior for 2D object-based detection which can help produce a floor plan. It requires the PC first be clean so that snapshots around the interior can be passed on as subsets for the Object Detection service.

The output of 2D Image Capture is an index for per point level data that associates a point to known objects within the image.

Object Detection

Object Detection identifies the locations of building features (e.g., windows, doors) by leveraging a trained machine learning model. This model outputs metadata that defines the interior features and can analyze the geometry to detect other structural elements such as walls, ceilings, and other interior objects like desks and chairs. Having these elements identified also allows for exclusion or inclusion of those objects in the final interior space.

The output of Object Detection, similarly to 2D Image Capture, is an index that associates the point to a known object within the image. The key difference being this is done from 3D information, allowing for easy visualization of the different objects. For example, in Figure 13 each valid object is represented in a different color.

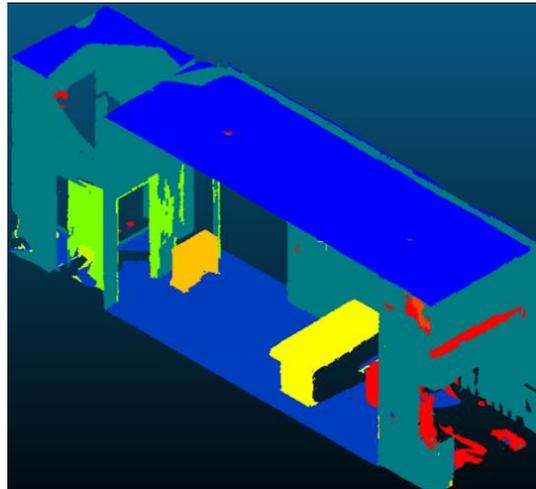


Figure 13. Machine learning object detection showing walls (teal), door apertures (green), window apertures (red), desk (yellow), and a shelf (orange)

Metadata: Apertures

The Metadata: Apertures takes in any defined apertures from Object Detection (e.g. windows or door openings) and combines it with any exterior metadata that would be relevant.

The output of Metadata creates a master set of aperture data structures using a combination of exterior and interior information. This combined dataset is used in a variety of ways to help generate the overall interior geometry. For example, the location of apertures is used to define the internal topology of the building, including the relationships between rooms and floors. The apertures are also used to determine the volumetric boundary of a room.

CONCLUSIONS AND RECOMMENDATIONS

Through the process of creating a texture mesh using collected interior building data, we discovered that multiple processes such as downsampling, noise reduction, segmentation, and object detection would be needed to accurately portray a high-resolution input mesh with embedded attribution needed for modeling and simulation purposes. By dividing each process into individual services and multiple iterations, we were able to start to link together processes in a configurable order to create the basis of this pipeline.

There are several industries, including military modeling and simulation, that require geo-specific interiors data to address needs. While there have been numerous advancements in the last 20+ years regarding digital representations of the interiors, there is more work that will need to be completed. The issue at hand is determining when the post-processed interiors data is *good enough* for different industry use cases. As an example, having realistic photogrammetric or point cloud meshes may be desired as a reference source for reproducing historic sites where

measurements and the object's texture/color are the primary focus. Likewise, tactical mission rehearsals need the mesh to be configurable to achieve goals, including allowing doors and windows to be opened and shut and identifying what materials the walls are made of (i.e., drywall, brick).

At present, there is no universal building interior data model or pipeline that can address every need. The pipeline we proposed highlights that users should incorporate steps that are required for the end use application, as well as consider if they are necessary to accomplish stated objectives. Steps such as "Planarizing" are extremely useful for making the mesh more compact and efficient for run-time environments, but it comes at a tradeoff by not capturing every minor detail from the data source.

Future research should focus on attribution of building features within a mesh or model as well as open standards on what constitutes the components of an interior of a building.

ACKNOWLEDGEMENTS

We wish to acknowledge the significant contributions by our fellow Dignitas Team Members (Freddie Santiago, Matt Rolfe), our customers from the U.S. Army STTC Center (Clay Burford, Patricia Montoya, Gage Jenners), and our IITSEC bird dog Einav Kiperman. Without their contributions and patience this paper would not have been possible.

REFERENCES

- Bent, G., Pfaff, D., Brooks, M., Radpour, R., & Delaney, J. (2022). A practical workflow for the 3D reconstruction of complex historic sites and their decorative interiors: Florence As It Was and the church of Orsanmichele. *Heritage Science, 10*. Retrieved from <https://heritagesciencejournal.springeropen.com/articles/10.1186/s40494-022-00750-1>
- Furukawa, Y., Curless, B., Seitz, S., & Szeliski, R. (2009). Reconstructing Building Interiors from Images. *2009 IEEE 12th International Conference on Computer Vision* (pp. 80-87). Kyoto: IEEE. Retrieved from <https://grail.cs.washington.edu/wp-content/uploads/2016/09/furukawa2009rbi.pdf>
- Goldman, A., & Berman, M. (2016, August 1). 'They took too damn long': Inside the police response to the Orlando shooting. *The Washington Post*. Retrieved from https://www.washingtonpost.com/world/national-security/they-took-too-damn-long-inside-the-police-response-to-the-orlando-shooting/2016/08/01/67a66130-5447-11e6-88eb-7dda4e2f2aec_story.html
- Holmberg, D., Davis, W., Treado, S., & Reed, K. (2006). *Building Tactical Information System for Public Safety Officials: Intelligent Building Response (iBR)*. National Institute of Standard and Technology . Gaithersburg: U.S DEPARTMENT OF COMMERCE. Retrieved from <https://www.govinfo.gov/content/pkg/GOVPUB-C13-b2c9928b4ef83598bb8cf0f8722ce1a3/pdf/GOVPUB-C13-b2c9928b4ef83598bb8cf0f8722ce1a3.pdf>
- Mann, J., & Eifert, L. (2006). Easing the Pain of Urban Modeling. *2006 Paper No. 2811* (pp. 1-10). Orlando: Interservice/Industry Training, Simulation, and Education Conference (IITSEC) 2006. Retrieved from <http://www.iitsecdocs.com/download/2006/8V1BAQYTGBGCBH6T>
- Marketwired. (2014, April 29). U.S. Army Research Laboratory Awards GameSim With Procedural Model Generation Service Contract. *Marketwired*. Retrieved from https://finance.yahoo.com/news/u-army-research-laboratory-awards-132011744.html?guccounter=1&guce_referrer=aHR0cHM6Ly93d3cuZ29vZ2xiLmNvbS8&guce_referrer_sig=AQAAALP3hHXi6JcIVv53LW9F23Sbsa283hQLx6CUPal7zgL_tWA4c0OwlWHTcf10TmHIOIIAK2qKbcu4QIIF48unz6u9nEcX11
- Merrell, P., Schkufza, E., & Koltun, V. (2010). Computer-Generated Residential Building Layouts. *ACM Transactions on Graphics, 29*(6), 1-12. Retrieved from <https://paulmerrell.org/wp-content/uploads/2021/06/floorplan-final.pdf>
- Miles, B., Creel, T., Kenton, A., Abrams, M., & Wilder, K. (2019). Enhancing Situational Awareness Anywhere in the World with Geospatially Accurate Scene Simulation using Automated 'Real World' Content. *2019 Paper No. 19112* (pp. 1-11). Orlando: Interservice/Industry Training, Simulation, and Education

- Conference (IITSEC). Retrieved from http://amz.xcdsystem.com/44ECEE4F-033C-295C-BAE73278B7F9CA1D_abstract_File4313/PaperUpload_19112_0612125333.pdf
- Moore, R., Reilly, M., & Pelham, T. (2018). Exterior Attribute Extraction and Interior Layout Speculation of 3D Structures. *2018 Paper No. 18243* (pp. 1-13). Orlando: Interservice/Industry Training, Simulation, and Education Conference (IITSEC) 2018. Retrieved from http://www.iitsecdocs.com/download/2018/2018_18243
- Pigora, M. A., Graniela, B., & Kraus, M. (2002). A SYSTEM FOR THE RAPID CAPTURE OF VIRTUAL ENVIRONMENTS. *2002 Paper* (pp. 1-7). Orlando: Interservice/Industry Training, Simulation, and Education Conference (IITSEC) 2002. Retrieved from <http://www.iitsecdocs.com/download/2002/WRX8JQG4DJ6N5K73>
- Raedle, J. (2016, June 13). *Orlando Gunman Discussed ISIS With Police, Called Boston Bombers His 'Homeboys'*. Retrieved from *Intelligencer*: <https://nymag.com/intelligencer/2016/06/orlando-nightclub-shooting-leaves-20-dead.html>
- Tilton, B. (2021). One World Terrain. *Tech Showcase for FT Meade* (pp. 6, 17). Orlando: Maxar. Retrieved from https://www.militaryexpos.com/wp-content/uploads/2021/09/U_0900_Tilton_OWT.pdf
- U.S. Army. (2023). *SYNTHETIC TRAINING ENVIRONMENT (STE)*. Retrieved from PEO STRI: <https://www.peostri.army.mil/synthetic-training-environment-ste/>
- Wiesmayer, P. (2020, January 14). Reconstructing Notre Dame with the help of 3D data from Germany. *Innovation Origins*. Retrieved from <https://innovationorigins.com/en/reconstructing-notre-dame-with-the-help-of-3d-data-from-germany/>