

NEW POTREE SHADER CAPABILITIES FOR 3D VISUALIZATION OF BEHAVIORS NEAR COVID-19 RICH HEALTHCARE FACILITIES

Chris Carey¹, José Romero¹, Debra F. Laefer^{1,2}

¹Center for Urban Science + Progress, New York University, USA - (cpc9982, jrr9496, debra.laefer)@nyu.edu

²Dept. Civil and Urban Engineering, Tandon School of Engineering, New York University, USA

KEY WORDS: Three-dimensional epidemiology, Point cloud, WebGL, COVID-19, Healthcare facility, Geoscientific information

ABSTRACT:

While data on human behavior in COVID-19 rich environments have been captured and publicly released, spatial components of such data are recorded in two-dimensions. Thus, the complete roles of the built and natural environment cannot be readily ascertained. This paper introduces a mechanism for the three-dimensional (3D) visualization of egress behaviors of individuals leaving a COVID-19 exposed healthcare facility in Spring 2020 in New York City. Behavioral data were extracted and projected onto a 3D aerial laser scanning point cloud of the surrounding area rendered with Potree, a readily available open-source Web Graphics Library (WebGL) point cloud viewer. The outcomes were 3D heatmap visualizations of the built environment that indicated the event locations of individuals exhibiting specific characteristics (e.g., men vs. women; public transit users vs. private vehicle users). These visualizations enabled interactive navigation through the space accessible through any modern web browser supporting WebGL. Visualizing egress behavior in this manner may highlight patterns indicative of correlations between the environment, human behavior, and transmissible diseases. Findings using such tools have the potential to identify high-exposure areas and surfaces such as doors, railings, and other physical features. Providing flexible visualization capabilities with 3D spatial context can enable analysts to quickly advise and communicate vital information across a broad range of use cases. This paper presents such an application to extract the public health information necessary to form localized responses to reduce COVID-19 infection and transmission rates in urban areas.

1. INTRODUCTION

The effects of the built environment on human behavior conducive to COVID-19 transmission are not fully understood at the community level. To contribute to the state of understanding, this paper introduces a visualization mechanism using a three-dimensional (3D) point cloud on which to project egress behaviors of individuals leaving a COVID-19 exposed healthcare facility in Spring 2020 in New York City (NYC), during the initial peak outbreak of COVID-19. The goal was to demonstrate in an aggregated manner how individuals exhibiting specific characteristics (e.g. whether they touched their body or their cell phone) physically interacted with the surrounding environment and community and how these behaviors differed by subgroup (e.g. men vs. women).

Geographic Information System (GIS) applications have long been used to model and visualize disaster-related behaviors in two-dimensions [e.g. egress behaviors of Californians at risk from wildfires (Cova and Church, 1997)]. Infectious disease behavior has also been visualized in two-dimensions. However, the tendency of cases to cluster typically requires data points to be aggregated into density plots to enable sufficient visual interpretation. Atkinson and Unwin (2002) noted that visualizing such data as density plots in 3D space could further improve their interpretability in geographic contexts.

The increasing availability of 3D spatial data in the form of national aerial laser scans (U.S. Geological Survey, 2013) has created the opportunity for a new generation of visualization tools to aid understanding behavioral response to external agents (Castrillón et al., 2011). These 3D visualizations can enhance the understanding of the spatial context of the underlying data by improving spatial thinking. Spatial thinking is the ability to

visualize and solve problems spatially (Nielsen et al., 2011) and studies have shown the benefits of 3D visualization in improving these skills. Rautenbach et al. (2014) conducted a comparative experiment in South Africa with planners with low/medium levels of map literacy to evaluate their spatial skills in two-dimensional (2D) maps and 3D models of their country. Participants were asked to solve three tasks: cardinal direction, relative direction, and distance estimation. Most participants felt the assignments were easier to solve using the 2D maps. However, the results showed better performance when 3D representations were used. Three-dimensional models have the potential of improving users' geospatial thinking in specific tasks. For example, Carbonell-Carrera and Hess-Medler (2019) conducted two workshops with engineering students using visuospatial displays in 3D software environments to enhance their abilities in relief interpretation and landform comprehension. The Topographic Map Assessment (TMA) test, developed by Jacovina et al. (2014), was used to assess the effectiveness of these 3D software technologies on developing geospatial thinking. Evidence from the study showed gains of 10.7%-12.6% of TMA test scores between pre-workshop and post-workshop assessments.

With the increasing popularity of 3D models, understanding how certain design decisions of 3D visualizations affect the user experience and their map cognition abilities is essential. Elements such as the extent of the model (scale), scene complexity (number, structure, and density of objects), the field of view (perspective), familiarity, as well as training and instruction clarity, can affect the user experience and the achieved outcomes (Rautenbach et al., 2015). Research by Cubukcu (2011) suggests that improper configuration of a 3D model can cause anxiety, disorientation, frustration, and stress that would affect performance when using or interpreting the model. Introducing interactivity that enables users to control the configura-

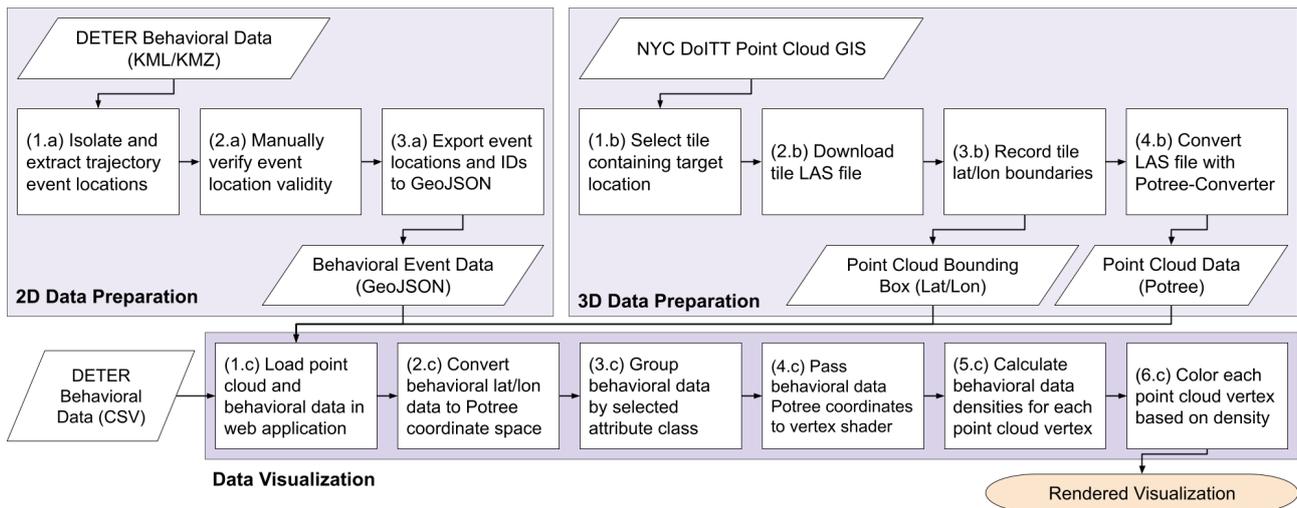


Figure 1. Process of data preparation and visualization.

tion of these 3D models could improve users' spatial cognition. Herman et al. (2018) noted that there are significant differences between interactive and static 3D models in effectiveness, efficiency, and subjective preferences. Additionally, the flexibility that interactivity creates contributes to improved spatial and contextual understanding for complex and combined tasks (Herman et al., 2018). Introducing such interactivity into a GIS-based, 3D geovisualization of human behavior in particular could enable researchers and policymakers to explore the relation of the built environment to complex spatial patterns among population subgroups in a medium that is flexible (Kwan, 2000) and beneficial towards understanding inherently 3D structures (St. John et al., 2001).

Recent advances in processing and rendering large quantities of 3D spatial data in web browsers have made the visualization of these data in web applications more feasible to produce. For example, Potree (Schütz, 2015) is an open-source Web Graphics Library (WebGL) point cloud renderer capable of rendering billions of points by storing and rendering data subsamples at varying resolutions based on the view frustum. For the visualization and analysis of epidemiological data, such web-based systems have the advantage of lower implementation costs and easier distribution compared to systems without web-based components (Carroll et al., 2014). In fact, numerous web-based tools for the spatio-temporal analysis of COVID-19 case data have been developed and are distributed publicly (Franch-Pardo et al., 2020). Because of Potree's ability to host point clouds from laser scanning, it was selected to demonstrate the applicability of GIS-based, 3D visualization towards understanding community-level COVID-19 transmission in urban environments. To date, such an approach has yet to be applied to point cloud data.

2. METHODOLOGY

2.1 Scope

This paper introduces an initial demonstration of a method for the web-based visualization of 2D spatio-temporal human behavior data within a 3D representation of the built environment, in the context of the COVID-19 pandemic. The data subset used in this paper consisted of records collected from a single healthcare facility in NYC from March 30-May 7, 2020.

2.2 Data Preparation

The web application developed herein renders 2D behavioral data and 3D point cloud data together in a single, interactive WebGL visualization with Potree. These data were filtered and formatted for combination and rendering. The processes for data preparation and visualization applied in this paper are presented in Figure 1.

2.2.1 Two-dimensional Data Preparation: Behavioral data came from the National Science Foundation funded project Developing Epidemiology Mechanisms in Three-Dimensions to Enhance Response (DETER), which produced a publicly-released dataset available as individual KML/KMZ files, an aggregated shape file, and an aggregated CSV file (Laefer et al., 2021). The entire dataset contained more than 5,100 records from more than 6,100 subjects observed at 19 NYC hospitals and urgent care centers from March 23-May 18, 2020. The behavioral KML/KMZ data included text annotations from data collectors, trajectory points of observed individuals, and event locations where observed subjects performed notable actions such as entering a store, meeting with other individuals, or exiting the observation area via public transport.

The CityMD Parkchester Urgent Care in Bronx, NY was selected to demonstrate the visualization capabilities in this paper, due to the quantity of available behavioral data collected in a format which lent itself to automatic classification and behavioral event location extraction. For each observation in this subdataset, a single polygon represented a trajectory line of the observed individual, while a separate polygon represented an approximate circle encompassing the location of a behavioral event. While an observation could contain multiple trajectory lines and multiple event locations, each was represented by an individual polygon. As such, the behavioral event polygons could be automatically distinguished from trajectory polygons by filtering out polygons whose bounding boxes exceeded a maximum width or height. A Python script was written to partially automate this classification process (1.a) through the open-source, GIS software QGIS (QGIS Development Team, 2021). This automated process was followed with manual, human verification to correct misclassifications (2.a). The bounding box centers of the verified event polygons and associated observation IDs were then exported to the GeoJSON format (3.a).

2.2.2 Three-dimensional Data Preparation: Point cloud data downloaded from NYC Open Data (NYC DoITT GIS Unit, 2017) were used for the 3D representation of the area surrounding Parkchester Urgent Care. The tile encompassing the site location was selected using the provided web interface (1.b) and downloaded in an LAS format (2.b). The displayed bounding latitude and longitude values were collected for later conversion into the coordinate space within the Potree visualization (3.b). Lastly, the LAS file was converted to the Potree format using the accompanying Potree-Converter tool (Potree Development Team, 2021) (4.b).

2.3 Data Visualization

To render the visualization, the web application first loaded the necessary datasets into memory and initialized Potree with the point cloud data (1.c). The behavioral event latitude and longitude values were then mapped to their corresponding coordinates in the Potree coordinate space using a linear transformation between the point cloud tile's bounding latitude and longitude values and the resulting minimum and maximum x- and y-values of the point cloud tile in the Potree coordinate space (2.c). Each 2D behavioral event location also needed a z-value to be rendered in 3D space. Every behavioral location was assigned a constant z-value which approximately matched ground level for all point cloud points in the observation space. Each z-value could be set according to the localized ground elevation determined by a digital elevation model to improve precision in areas of varying elevation. Behavioral event locations were then conditionally grouped according to their class within a particular attribute (e.g. day of week), as selected by the user (3.c). This was performed by using behavioral event IDs to combine the GeoJSON and CSV behavior data and group them according to their recorded attribute value within the CSV table.

The last step was to visualize the behavioral locations as density plots projected onto the 3D point cloud rendered using Potree. Behavioral locations were passed into the Potree vertex shader grouped by their corresponding attribute class (4.c). The primary vertex shader was modified to estimate each attribute class' behavioral event location density at each point cloud vertex (5.c). The number of behavioral event locations within a fixed search area surrounding each point cloud vertex was counted for each attribute class and weighted using a linear kernel function to reduce the influence of behavioral locations closer to the edge of the search area (Atkinson and Unwin, 2002). The resulting density values were then converted to vertex colors with intensities representing their corresponding event densities (6.c).

2.3.1 Vertex Shader Algorithm: The vertex shader assigns the color of every point cloud vertex. The application was configured by default to color each vertex in grayscale with a linear relationship between the vertex brightness and elevation. In order to change the color of point cloud vertices to correspond to the density of nearby behavioral events, the operation described in Algorithm 1 was appended to the primary Potree vertex shader.

The Potree vertex shader ran this operation for every point cloud vertex. This resulted in an $O(V \times S)$ time operation where V is the number of vertices and S is the number of event locations. Therefore, computationally intensive calculations were limited. As such, in Algorithm 1, to determine if an event location was within the kernel search radius of a vertex, the dot product of a , the vector between them, was compared to the squared search

radius. The exact magnitude of a was only calculated if this condition was met, as square root operations were more computationally expensive to perform than dot product operations.

The color for maximum density was set as $c_{\max} = rgb(1.0, 0.5, 0.0)$. This resulted in a bright orange color when $\rho = \rho_{\max}$ which distinguished areas of high density apart from monochromatic vertices representing areas with zero density.

Algorithm 1 Univariate Vertex Color

Input:

- B (behavioral event positions)
- c_0 (initial vertex color)
- c_{\max} (max density color)
- r (kernel search radius)
- v (vertex position)
- ρ_{\max} (max density value)

Output:

- c (new vertex color)

```

1: function GETDENSITY( $B$ )
2:    $\rho \leftarrow 0$ 
3:   for each  $b \in B$  do
4:      $a \leftarrow b - v$ 
5:     if  $\|a\|^2 < r^2$  then
6:        $\rho \leftarrow \rho + (r - \|a\|)$ 
7:     end if
8:   end for
9:   return  $\rho$ 
10: end function
11:  $\rho \leftarrow \text{GETDENSITY}(B)$ 
12:  $c \leftarrow c_0 + \frac{\rho}{\rho_{\max}}(c_{\max} - c_0)$ 
13: return  $c$ 

```

2.3.2 Bivariate Color Shading: To estimate densities of behaviors of more than one attribute class (e.g. male and female), two separate density estimations were maintained – one for each attribute class: ρ_1 and ρ_2 . This required a separate maximum density color for each attribute class: $c_{1\max}$ and $c_{2\max}$. As each point cloud vertex could only represent a single color, these colors were combined in a manner representative of the proportionality of their corresponding densities. As such, a third color was used to represent the maximum total density consisting of equal densities of both classes: c_{midMax} . A piecewise linear transformation was performed such that:

$$c_{\max} = \begin{cases} c_{1\max}, & \text{if } \rho_1 \geq \rho_{\max} \text{ and } \rho_2 = 0 \\ c_{2\max}, & \text{if } \rho_2 \geq \rho_{\max} \text{ and } \rho_1 = 0 \\ c_{\text{midMax}}, & \text{if } \rho_1 = \rho_2 \geq \frac{1}{2}\rho_{\max} \end{cases} \quad (1)$$

The colors representing their corresponding maximum density values were selected to maximize visual differentiation between attribute classes, such that the majority class of areas of high density could be easily identified (Figure 2). Algorithm 2 describes the vertex shader color assignment process when densities of two attribute classes were visualized.



Figure 2. Bivariate color scale for c_{\max} based on ρ_1 and ρ_2 .

Algorithm 2 Bivariate Vertex Color

Input:

- B_1, B_2 (behavioral event positions by class)
- c_0 (initial vertex color)
- $c_{1\max}, c_{2\max}$ (max density color by class)
- c_{midMax} (max density color when $\rho_1 = \rho_2$)
- r (kernel search radius)
- v (vertex position)
- ρ_{\max} (max density value)

Output:

- c (new vertex color)

```

1:  $\rho_1 \leftarrow \text{GETDENSITY}(B_1)$ 
2:  $\rho_2 \leftarrow \text{GETDENSITY}(B_2)$ 
3:  $k \leftarrow \frac{\rho_{\max}}{\min(\rho_1 + \rho_2, \rho_{\max})}$ 
4:  $c_{\text{mid}} \leftarrow c_0 + k(c_{\text{midMax}} - c_0)$ 
5: if  $\rho_1 \geq \rho_2$  then
6:    $c_1 \leftarrow c_0 + k(c_{1\max} - c_0)$ 
7:    $c \leftarrow c_1 + \frac{\rho_2}{\rho_1}(c_{\text{mid}} - c_1)$ 
8: else
9:    $c_2 \leftarrow c_0 + k(c_{2\max} - c_0)$ 
10:   $c \leftarrow c_2 + \frac{\rho_1}{\rho_2}(c_{\text{mid}} - c_2)$ 
11: end if
12: return  $c$ 

```

3. RESULTS

3.1 Visualization and User Interface

The resulting heatmap visualization enabled users to identify the most frequently exposed locations and surfaces based on color intensity. This method is comparable to traditional 2D heatmap visualizations, but with the advantage of encompassing structures in 3D (Figure 3). While it was possible to set the upper bound of ρ to ρ_{\max} for univariate color shading, allowing ρ to exceed ρ_{\max} produced the added benefit of a maximum density color of yellow [i.e. $rgb(1.0, 1.0, 0.0)$], when $\rho = 2\rho_{\max}$. This particular choice created a wider color range from gray to orange to yellow, thereby improving visual interpretability (Figure 4). User interface (UI) slider controls allowed the user to set the kernel density search radius r and the maximum density value ρ_{\max} which corresponded to the maximum color c_{\max} , due to the dependency of the observed maximum density value on the kernel density search radius and data quantity visualized.

Column name	Description
1st_Touch_Obj Coded	First object or body part touched coded into preset option
Day_Type	Day of the week (Mon-Sun)
Final_Destination _Coded	Final destination coded into one of several preset options
Gender	Majority of the gender in the recorded group
Re_Medical	Whether the subject used a mechanized means to exit area
Time_Type	Whether record occurred during common work hours
Touch_Binary	Whether the subject touched objects or parts of their body

Table 1. Attributes available for data segmentation.

Field of view could be freely translated and rotated using mouse controls directly on the visualization. Behavioral locations

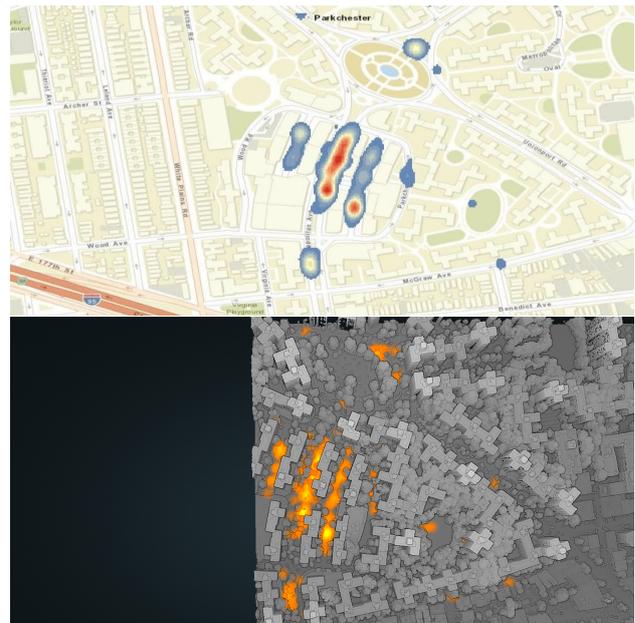


Figure 3. Cumulative egress behavioral data visualized as heatmaps in ArcGIS (top) and in Potree (bottom).

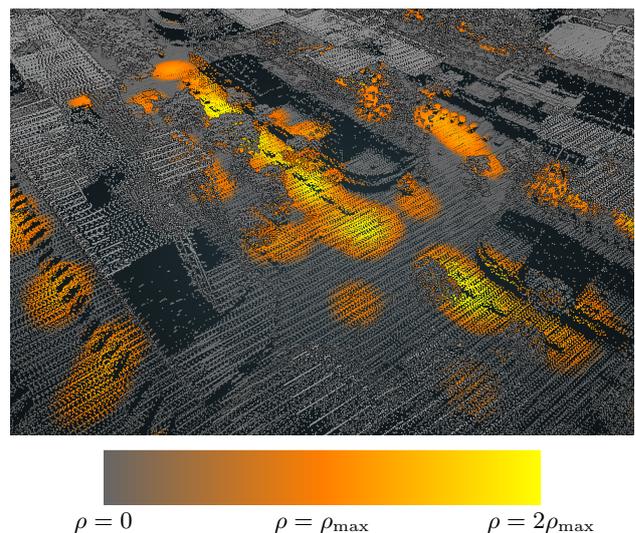


Figure 4. (Top) Egress behavioral data heatmap, with points shaded by behavioral event location density. (Bottom) Resulting color scale for vertex color c based on ρ when $c_{\max} = rgb(1.0, 0.5, 0.0)$ and ρ was unbounded.

could be grouped by certain attributes (Table 1) through UI controls, thereby allowing users to compare data segments (Figure 5). While only two attribute classes could be compared at a time, these attribute classes were user-selectable, enabling different comparisons of data segmented by multi-class classification attributes (Figure 6). Note that this only visualized the data subset containing records belonging to the two selected classes.

3.2 Performance

The visualization was then accessible through any modern web browser supporting WebGL. The application load times were observed for two machines to evaluate general performance. Each machine hosted and rendered the application in

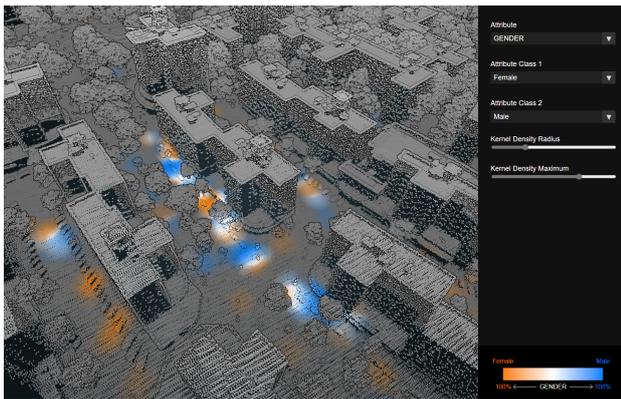


Figure 5. Distribution of female (orange) and male (blue) subject event locations with color scale.

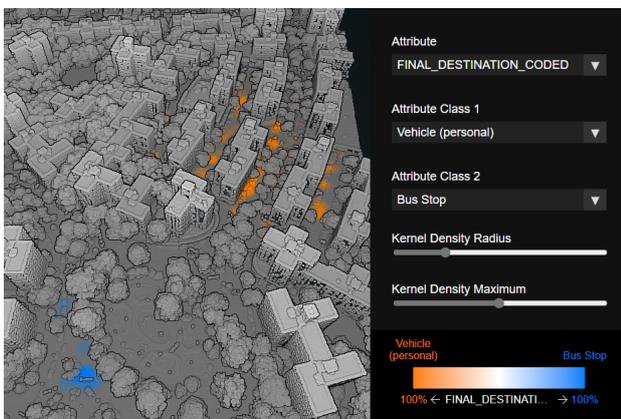


Figure 6. Event location distribution by a multi-class attribute: transportation choice destinations of subjects by private vehicle (orange) or bus stop (blue), with user control panel in detail.

Microsoft Edge 93.0.961.44 at a 1440x720 window size during testing. Performance metrics were captured using the available in-browser performance profiling tool. Given the full data subset of 223 behavioral event locations and 9,634,714 point cloud tuples, the Microsoft Surface Pro 6 with an Intel i7-8650U 1.90GHz Central Processing Unit (CPU) and Intel UHD Graphics 620 Graphical Processing Unit (GPU) took 4752 ms for scripting and 7752 ms in total. In comparison, the custom PC with an Intel i7-7700K 4.2 GHz CPU and Nvidia GeForce GTX 1080 Ti GPU took 3329 ms for scripting and 6126 ms in total. While performing interactive navigation actions with rendered behavioral data, noticeable input lag was encountered on the Surface Pro 6. This was not observed on the custom PC, thus, indicating the need to improve algorithm performance for less performant machines.

4. DISCUSSION

The web application presented in this paper provided 3D spatial context to data typically visualized and analyzed in traditional 2D GIS applications. By utilizing 3D perspective to visualize the built environment, color hue was freed as a channel to encode behavioral data without occluding physical features. The interactivity of the application and the ability to visualize and filter data based on specific characteristics augmented the context provided by this third dimension, thereby enabling quicker and more localized analyses of observed patterns. This ability

to quickly understand the spatial contexts of specific locations without being there in person is important in forming public health policies remotely when physical presence is limited.

While the 3D spatial context provided through the visualization method presented in this paper can enhance the understanding of 2D behavioral data, analysis is limited by the granularity level of point cloud data used. Smaller objects of interest that don't appear identifiable in the point cloud data, such as door handles, railings, and crosswalk signal buttons, would not appear in this visualization and would require additional data to identify. Given visual geotagged imagery data of the surrounding environment, this application could link vertex locations to corresponding imagery, though such analysis would be limited to 2D.

Point cloud data collected through higher-resolution terrestrial laser scanning methods could also enable these object identification tasks. While this would require storing, transferring, and rendering a greater quantity of data, object identification tasks within the built environment are important in many other contexts, such as investigating connections between the built environment and criminal behavior or monitoring changes between the natural environment and animal behavior. As currently built, these tasks would be human-driven. However, advances have been made in applying deep convolutional neural networks to urban environment classification of point clouds produced by both aerial (Zhao et al., 2018) and terrestrial (Guo and Feng, 2020) laser scanning. These deep learning methods could provide automated assistance in establishing such correlations between human behavior and identified objects, structures, and other environmental characteristics.

Visualizing 2D behavior data in higher resolution 3D point cloud data will require further techniques to improve efficiency beyond those demonstrated in this paper, such as spatial indexing of the behavioral data. These techniques, in addition to user control of performance-based tuning parameters, such as maximum points rendered, would also improve performance on less powerful machines. While point cloud rendering and interactive navigation were not substantially hindered on such machines during testing, interactivity was drastically improved when tested on a machine with a performant GPU. Further testing should measure and monitor input lag directly in these situations.

Despite expected differences in performance across machine types, visualization and analysis capabilities were still enabled in an open-source, easy-to-use manner that united publicly accessible 2D, GIS-based behavioral data with publicly available 3D point cloud data to identify locations of interest. While this paper presented a visualization of historical behavioral data, the construction of the application through web-based components provides the potential to ingest and visualize real-time streaming data, similar to the prevalent online COVID-19 metric dashboards released publicly by various governments and organizations during the COVID-19 pandemic. Such 3D visualization could accelerate the response by healthcare facilities and surrounding environs to communicate risk areas in the built environment during epidemiological events. Timely risk area identification can advise public health as to the need for targeted policies to reduce infections in urban areas. Understanding data from the COVID-19 pandemic through such visualization methods can also better inform researchers and policymakers as to the value of specific types of data collection for future natural disasters.

5. CONCLUSIONS & FUTURE WORK

This paper demonstrates a method for visualizing 2D behavioral data in point cloud data by modifying Potree vertex shaders to render 3D density plots in an interactive web application. Applying this visualization method to urban point cloud data and egress behaviors of COVID-19 rich healthcare facilities in NYC provides the additional spatial context that may highlight patterns potentially indicative of correlations between the environment, human behavior, and COVID-19 transmission. This can further inform localized public health responses and communicate important health and safety information based on spatio-temporal analyses conducted on COVID-19 transmission risk.

Further application development could expand these analysis capabilities. Annotating the visualization with additional geographic context through text labels, feature coloring, and feature overlays could enhance analysis efficiency for researchers. Date range selection and comparison capabilities would enable quicker spatio-temporal comparisons. Indeed, a great opportunity lies in comparing multiple observation sites and time periods to identify differences and changes in behavioral patterns as a result of localized or policy-driven conditions. The capability of Potree to render point cloud datasets of hundreds of billions of points could enable city-scale visualization. As such, work will continue to expand the spatio-temporal coverage of visualized human behavioral data from the COVID-19 pandemic and potentially other transmissible diseases to demonstrate further applications of public interest within urban environments.

ACKNOWLEDGEMENTS

This work was supported by the National Science Foundation (award #2027293), the Data Science and Software Services (DS3) which is funded by the Moore and Sloane foundations through the NYU Moore Sloane Data Science Environment, and Bluefield GIS.

REFERENCES

- Atkinson, P., Unwin, D., 2002. Density and local attribute estimation of an infectious disease using MapInfo. *Computers & geosciences*, 28(9), 1095–1105.
- Carbonell-Carrera, C., Hess-Medler, S., 2019. 3D landform modeling to enhance geospatial thinking. *ISPRS International Journal of Geo-Information*, 8(2), 65.
- Carroll, L. N., Au, A. P., Detwiler, L. T., Fu, T.-c., Painter, I. S., Abernethy, N. F., 2014. Visualization and analytics tools for infectious disease epidemiology: a systematic review. *Journal of Biomedical Informatics*, 51, 287–298.
- Castrillón, M., Jorge, P. A., López, I. J., Macías, A., Martín, D., Nebot, R. J., Sabbagh, I., Quintana, F. M., Sánchez, J., Sánchez, A. J. et al., 2011. Forecasting and visualization of wildfires in a 3D geographical information system. *Computers & geosciences*, 37(3), 390–396.
- Cova, T. J., Church, R. L., 1997. Modelling community evacuation vulnerability using GIS. *International Journal of Geographical Information Science*, 11(8), 763–784.
- Cubukcu, E., 2011. Does the level of visual detail in virtual environments affect the user's spatial knowledge? *Environment and Planning B: Planning and Design*, 38(4), 741–752.
- Franch-Pardo, I., Napoletano, B. M., Rosete-Verges, F., Billa, L., 2020. Spatial analysis and GIS in the study of COVID-19. A review. *Science of The Total Environment*, 739, 140033.
- Guo, Z., Feng, C.-C., 2020. Using multi-scale and hierarchical deep convolutional features for 3D semantic classification of TLS point clouds. *International Journal of Geographical Information Science*, 34(4), 661–680.
- Herman, L., Juřík, V., Stachoň, Z., Vrbík, D., Rusznák, J., Řezník, T., 2018. Evaluation of user performance in interactive and static 3D maps. *ISPRS International Journal of Geo-Information*, 7(11), 415.
- Jacovina, M., Ormand, C., Shipley, T., Weisberg, S., 2014. Topographic Map Assessment. silc.northwestern.edu/topographic-map-assessment-tma/ (4 August 2021).
- Kwan, M.-P., 2000. Interactive geovisualization of activity-travel patterns using three-dimensional geographical information systems: a methodological exploration with a large data set. *Transportation Research Part C: Emerging Technologies*, 8(1-6), 185–203.
- Laefer, D. F., Kirchner, T., Cheong, D., Khan, A., Qiu, W., Tai, N., Truong, T., Virk, M. et al., 2021. Data Resource Profile: Egress Behavior from Select NYC COVID-19 Exposed Health Facilities March-May 2020. *arXiv preprint arXiv:2101.10079*.
- Nielsen, C. P., Oberle, A., Sugumaran, R., 2011. Implementing a high school level geospatial technologies and spatial thinking course. *Journal of Geography*, 110(2), 60–69.
- NYC DoITT GIS Unit, 2017. Topobathymetric LiDAR Data (2017) Dataset. data.cityofnewyork.us/City-Government/Topobathymetric-LiDAR-Data-2017-/7sc8-jtbz (1 April 2021).
- Potree Development Team, 2021. Potree Software, Version 1.8.0. potree.org (1 April 2021).
- QGIS Development Team, 2021. QGIS Software, Version 3.16.4-Hannover. potree.org (1 March 2021).
- Rautenbach, V., Coetzee, S. M., Çöltekin, A., 2014. Towards evaluating the map literacy of planners in 2d maps and 3d models in south africa. *AfricaGEO 2014 Conference Proceedings*. doi.org/10.5167/uzh-105682.
- Rautenbach, V., Çöltekin, A., Coetzee, S., 2015. Exploring the impact of visual complexity levels in 3D city models on the accuracy of individuals' orientation and cognitive maps. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 2, 499–506.
- Schütz, M., 2015. Potree: Rendering large point clouds in web browsers. PhD thesis, Wien.
- St. John, M., Cowen, M. B., Smallman, H. S., Oonk, H. M., 2001. The use of 2D and 3D displays for shape-understanding versus relative-position tasks. *Human Factors*, 43(1), 79–98.
- U.S. Geological Survey, 2013. 3D Elevation Program 1-Meter Resolution Digital Elevation Model. usgs.gov/core-science-systems/ngp/3dep/data-tools (4 June 2021).
- Zhao, R., Pang, M., Wang, J., 2018. Classifying airborne LiDAR point clouds via deep features learned by a multi-scale convolutional neural network. *International journal of geographical information science*, 32(5), 960–979.