Geollery: A Mixed Reality Social Media Platform

Ruofei Du
Department of Computer Science
University of Maryland, College Park
me@duruofei.com

David Li
Department of Computer Science
University of Maryland, College Park
dli7319@terpmail.umd.edu

Amitabh Varshney
Department of Computer Science
University of Maryland, College Park
varshney@cs.umd.edu

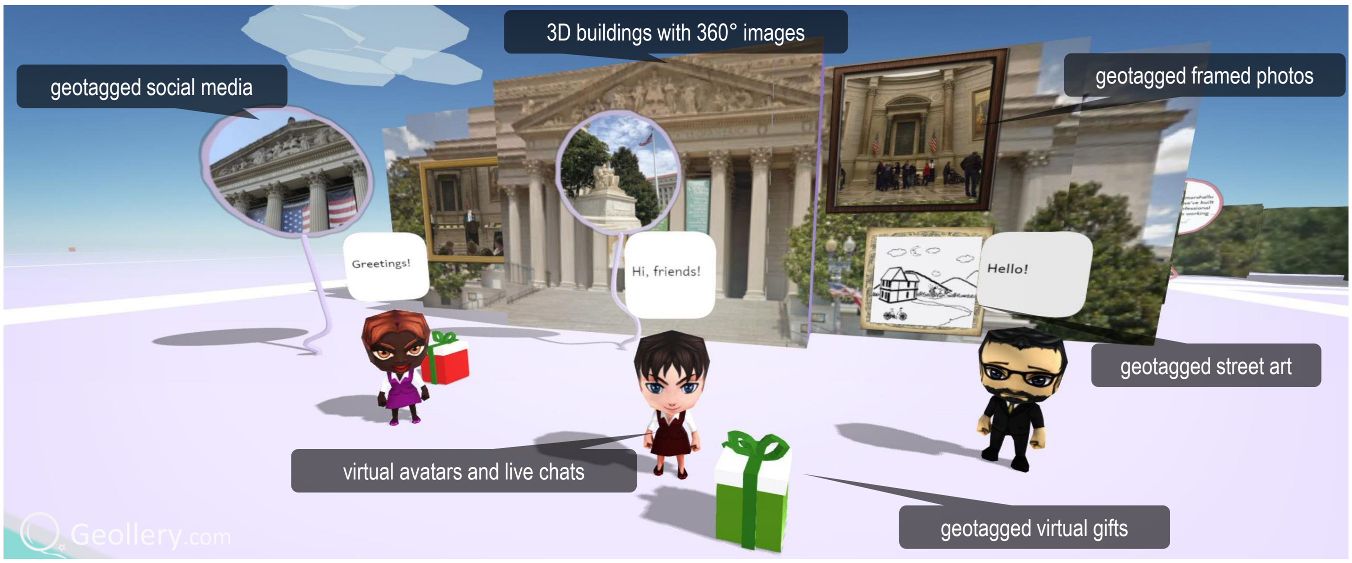

Figure 1: Geollery creates an interactive *mirrored world* where users are immersed with 3D buildings, live chats, and geotagged social media. The social media are visualized as balloons, billboards, framed photos, and gift boxes *in real time*.

ABSTRACT

We present Geollery, an interactive mixed reality social media platform for creating, sharing, and exploring geotagged information. Geollery introduces a real-time pipeline to progressively render an interactive mirrored world with three-dimensional (3D) buildings, internal user-generated content, and external geotagged social media. This mirrored world allows users to see, chat, and collaborate with remote participants with the same spatial context in an immersive virtual environment. We describe the system architecture of Geollery, its key interactive capabilities, and our design decisions. Finally, we conduct a user study with 20 participants to qualitatively compare Geollery with another social media system, Social Street View. Based on the participants' responses, we discuss the benefits and drawbacks of each system and derive key insights for designing an interactive mirrored world with geotagged social media. User feedback from our study reveals several use cases for Geollery including travel planning, virtual meetings, and family gathering.

CHI 2019, May 4–9, 2019, Glasgow, Scotland Uk

ACM ISBN 978-1-4503-5970-2/19/05…$15.00
https://doi.org/10.1145/3290605.3300915

CCS CONCEPTS

• **Human-centered computing** **Mixed / augmented reality**; *Virtual reality*; Web-based interaction.

KEYWORDS

virtual reality, augmented reality, social media, GIS, street view, visualization, 3D user interface, 3D reconstruction

ACM Reference Format:
Ruofei Du, David Li, and Amitabh Varshney. 2019. Geollery: A Mixed Reality Social Media Platform. In *CHI Conference on Human Factors in Computing Systems Proceedings (CHI 2019), May 4–9, 2019, Glasgow, Scotland Uk.* ACM, New York, NY, USA, 13 pages. https://doi.org/10.1145/3290605.3300915

1 INTRODUCTION

Social media plays a significant role in many people's daily lives covering a wide range of topics such as restaurant reviews, updates from friends, and local events. Despite the huge innovation in virtual and augmented reality, existing social media platforms typically use a linear narrative or a grid layout. While these traditional layouts are efficient for quickly browsing through social media posts, they lack the spatial context associated with social media posts. By incorporating spatial context, several use cases emerge where 3D social media platforms may outperform traditional social media platforms: business advertising, crowdsourced tourism, immersive storytelling, and learning about the culture [18].

Recently, several technologies and designs [7, 18, 30] (Figure 2) have emerged for visualizing social media in *mirrored worlds*[1] [21]. Nevertheless, designing an interactive social platform with immersive geographical environments remains a challenge due to the *real-time* constraints of rendering 3D buildings. In addition, the design space of visualizing and interacting with social media in mixed reality settings is not yet fully explored.

For example, *Social Street View* [18, 19] has made some initial contributions in blending immersive street views with geotagged social media. Nevertheless, interaction is limited to street-level panoramas. Consequently, users can not virtually *walk* on the streets but can only *teleport* among the panoramas. Bulbul and Dahyot [7] further reconstruct three cities with the street view data and visualize the popularity and sentiments with virtual spots lights. However, their system requires 113 - 457 minutes to reconstruct each city and lacks the interactivity with online users. Kukka *et al.* [30] have presented the conceptual design of visualizing street-level social media in a 3D virtual city, *VirtualOulu* [2]. However, such pre-designed 3D city models are not practical for deployment in larger areas.

We present *Geollery* (Figures 1 and 2D), an interactive mixed-reality social media platform in 3D which uses a mirrored world rendered in *real-time*. We introduce a progressive pipeline that streams and renders a mirrored world with 3D buildings and geotagged social media. We extend the design space in several dimensions: progressively streamed meshes and view-dependent textures, virtual representations of social media, aggregation approaches, and interactive capabilities.

To evaluate the benefits and drawbacks of our system, we conduct a user study with 20 participants for comparing Geollery with Social Street View. The quantitative evaluation and individual responses reveal the strengths and

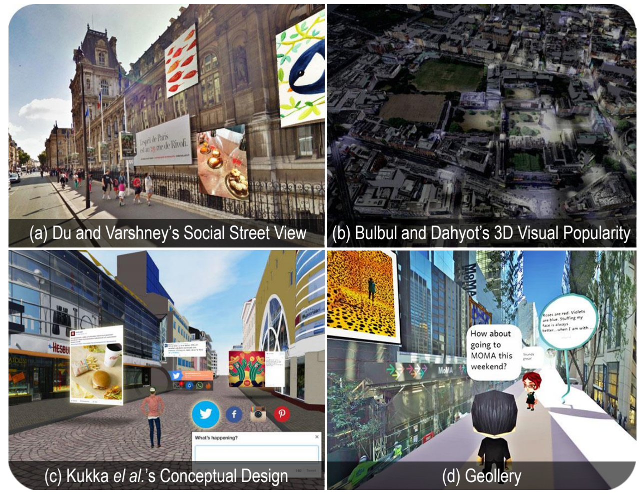

Figure 2: Comparison amongst mixed reality systems or designs for visualizing geotagged social media. (a) shows *Social Street View* [18], a real-time system which depicts social media as billboards via maximal Poisson-disk sampling [26], (b) shows Bulbul and Dahyot's offline system [7] which leverages virtual lighting to visualize popularity and sentiments of social media, (c) shows the conceptual design by Kukka *et al.* [30], which explores presentation manner, visibility, organization, and privacy during co-design activities, and (d) shows our results in *Geollery*, which fuses 3D textured buildings, geotagged social media, and virtual avatars in *real time.*

weaknesses of both systems. Our evaluation compares the different navigation methods used by each system and examines whether users would prefer walking or teleporting. Based on the responses in the semi-structured interviews, we further summarize the challenges and limitations of both systems, as well as their potential impact on future 3D social media systems. Finally, we improve *Geollery* based on the user feedback and deploy our system via Amazon Web Services (AWS). Please refer to https://geollery.com for the supplementary videos and live demos.

Our main contributions in this paper are:

(1) conception, development, and deployment of *Geollery*, an online system that can depict geotagged social media, 3D buildings, and panoramas in an immersive 3D environment,
(2) further extending the design space of 3D social media platforms such as aggregation approaches, virtual representations of social media, co-presence with virtual avatars, and collaboration modes,
(3) conducting a user study with 20 participants to qualitatively compare two 3D social media platforms (*Geollery* and *Social Street View*) by discussing their benefits, limitations, and potential impacts to future 3D social media platforms.

[1]A mirrored world is defined as "a representation of the real world in digital form [which] attempts to map real-world structures in a geographically accurate way" [49].

2 RELATED WORK

Our work builds upon the rich literature of previous research on geospatial visualization of social media in 2D maps and 3D spaces.

Geospatial Visualization in 2D Maps

Visualizing information in a geospatial manner has been around for as long as there have been maps. The ability to map, understand, see patterns, and draw conclusions from information presented in a spatially significant way is potent and intuitive. An early example of visualizing geotagged social media can be seen in *TwitterStand* [54] and *NewsStand* [35, 52, 60], where Twitter posts and news information are analyzed, streamed, and distributed on a map of the world as different types of icons [33, 34, 36]. *Panoramio*[2] and *PhotoStand* [51] aggregate and visualize geotagged images from professional photographers or news articles. In this way, users can see what information is available, where it originates from, and the density of the information.

In addition to icons and images, previous seminal research also explores various ways to analyze and visualize geotagged information on 2D maps. For example, MacEachren *et al.* [38] present a seminal system for visualizing the heat maps of health reports on a map. Their further work, *SensePlace2* [39], presents a geospatial visualization of Twitter messages with user-defined queries, time filters, spatial filters, and heap maps of tweet frequencies. Chae *et al.* [10] present a social media analysis system with message plots on a map, topic filtering, and abnormality estimation charts. Recent research also focuses on gridded heat maps [58], multivariate kernel methods [40], movement patterns [12], Reeb graphs [40], sentiment modeling [23, 37, 55], and flow visualizations of spatio-temporal patterns [29]. Using domain-specific knowledge, previous research has analyzed geotagged social media to improve emergency responses [64, 69], assist disease control [24], understand the dynamics of neighborhoods [13] and cities [63, 65], and travel route planning [31].

The key differentiator of our work is the ability to offer a third-person or first-person walking experience in immersive virtual environments. We discuss challenges such as creating a digital city in real time, designing virtual representations of content in 3D, aggregation approaches, and exploring interaction capabilities.

Geospatial Visualization in 3D Spaces

Creating an immersive visualization of geotagged social media is a challenging task due to the lack of 3D data. For example, reconstructing a 3D mirrored world from images typically requires intensive computation for a few hours or days. Early seminal work [1, 50, 56, 57, 61, 66] focuses on offline, image-based modeling approaches to generate virtual 3D cities. In these systems, 3D models are generated from a large collection of unstructured photos via different structure-from-motion pipelines. Despite the impressive results, the data requirement of such systems may not be satisfied in every city and the slow processing speed limits real-time applications. We direct readers to a thorough survey [45] for urban reconstructions. As discussed in Section 1, recent research offers more practical solutions to integrate geotagged social media with street-level panoramas [18], pre-reconstructed cities in several minutes [7], and virtual city or terrain models [3, 6, 30]. However, generating 3D models for the city is not quite applicable to real-time applications. On one hand, the texturing [7] of 3D buildings suffers from artifacts on complex geometries. On the other hand, the pre-crafted digital cities used in [3, 7, 8, 27, 68] are usually unavailable in rural areas and require enormous amounts of collaborative work from crowd workers, artists, researchers, and city planners [2, 22, 48, 62]. Moreover, without a partitioning algorithm, the digital cities (over 100 MB as mentioned in [3]) may be a bottleneck for practical online deployment.

In contrast to the prior art, we circumvent the offline reconstruction or manufacture of a digital city by progressively streaming open 2D maps. With 2D polygons and labels, *Geollery* extrudes and textures geometries on demand in *real time* using nearby street view data, enabling visualization of geo-relevant social media with their spatial context, and allowing user interaction in a mirrored world.

As for human factors in 3D social media platforms, Kukka *et al.* [30] conduct a pioneering qualitative *anticipated* user experience study with 14 participants to explore the design space of geospatial visualization of social media in mirror worlds. Badri *et al.* [3] further evaluate a banner editor system for adding and visualizing social media banners in *VirtualOulu*, a virtual digital city. Nevertheless, the human factors have not yet been fully discussed for *experiencing* a *real-time* mixed-reality social media platform such as *Geollery* or *Social Street View* [18]. In this paper, we conduct a comparative study with 20 participants and derive key insights from the semi-structured interviews. Our qualitative evaluation further reveals the strengths and weaknesses of *Geollery* and *Social Street View*.

3 SYSTEM OVERVIEW

In this section, we present an overview of *Geollery*'s system architecture. *Geollery* consists of a data engine which streams 2D polygons and labels from OpenStreetMap[3] and social media data from our internal database and external sources

[2]Panoramio: https://en.wikipedia.org/wiki/Panoramio

[3]OpenStreetMap: https://openstreetmap.org, an open world map, .

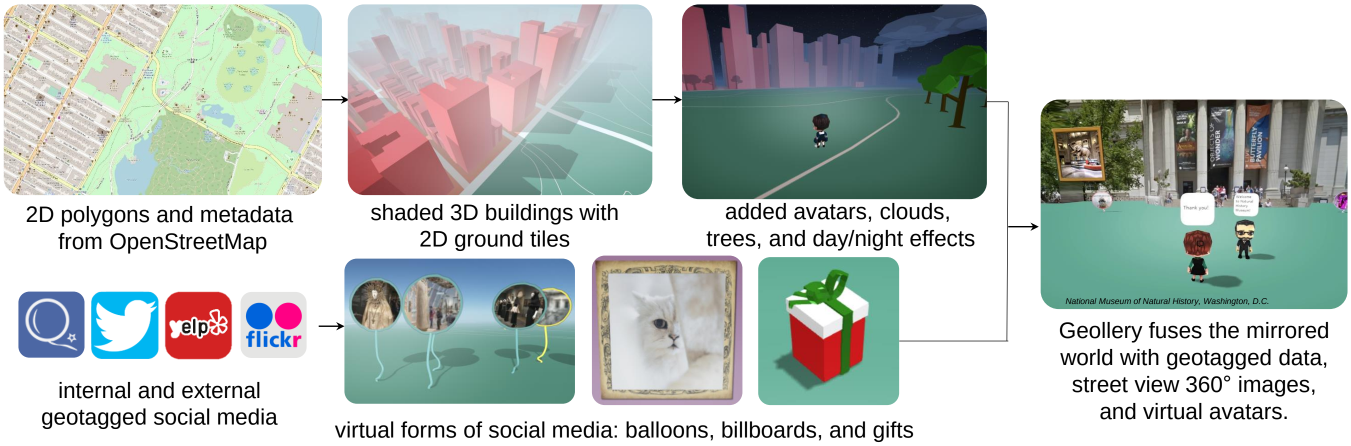

Figure 3: The workflow of *Geollery*. Based on users' geo-location requests, *Geollery* loads the nearby 2D map tiles, extrudes 3D geometries, and renders social media in real time. We take advantage of WebGL to enable users to access *Geollery* via modern browsers on a desktop, a mobile phone, or a head-mounted display.

such as *Twitter*[4], *Yelp*[5], and *Flickr*[6]. Deployment of *Geollery* requires only an SQL database and a web server powered by Apache and PHP. We take advantage of the *B+* tree to index the geotagged information for querying in real time. We build the rendering system upon `Three.js`[7], a cross-browser, GPU-accelerated JavaScript library.

Geollery allows users to explore social media nearby or at a custom location. Users have a choice of either sharing their device's current location or entering a query into a search box. Unlike prior art which aims to reconstruct the entire city, *Geollery* leverages a progressive approach to partially build the mirrored world. We present the workflow of *Geollery* in Figure 3.

First, given latitude and longitude coordinates, our system queries 2D map tiles and renders the ground plane within a radius of about 50 meters[8]. The ground plane visualizes roads, parks, waters, and buildings with a user-selected color scheme. As users virtually walk on the street, *Geollery* streams additional data into the rendering system. Next, *Geollery* queries 2D map data from OpenStreetMap to gather information about buildings and terrains. 3D geometries are extruded from 2D polygons and then shaded with the appropriate lighting and shadows to form buildings. Trees are randomly generated in forests. Finally, the system renders a mirrored world within the user's field of view in real time, which contains 3D buildings, virtual avatars, trees, clouds, and different forms of social media, such as balloons, billboards, framed photos, and virtual gifts.

[4]Twitter: https://twitter.com, social networking service.
[5]Yelp: https://yelp.com, local city guide.
[6]Flickr: https://flickr.com, an image hosting service.
[7]Three.js: http://www.threejs.org.
[8]Users can change this parameter in the settings.

Our system acquires geotagged social media by querying our server for internal and external geotagged social media. The client uses POST requests to gather all social media within a specific radius of the user's location, encoded as a latitude-longitude pair. For internal queries, our PHP server retrieves social media from our MySQL database. We search only for social media the requesting user is permitted to view and return the list of social media to the client as a JSON encoded array of metadata including the location, author, image URL, and the text caption of each post. For external queries, our system uses public APIs documented by Twitter and Yelp to acquire the latest public social media near a geographic location. Our server parses the social media returned by each API and passes it to the client. We also use Flickr photo metadata from a dataset by [44] imported into in our MySQL database. For each form of social media, 3D models are generated on the client to represent the social media in the mirrored world.

For registered users, *Geollery* connects clients to our server via HTML5 WebSockets, allowing real-time communication and collaboration with other nearby participants. We explain our design and implementation details for social media and live interaction in the next section.

4 DESIGN SPACE

As listed in Table 1, we explore and compare several variables in the design space of 3D social media platforms between *Geollery* and *Social Street View* [18], including the choices of meshes and textures, availability, degrees of freedom (DoF) in motion, virtual avatar, and social media representations. We further discuss other dimensions of interest such as privacy concerns, real-world phenomena, and temporal filters.

Variable	Geollery	Social Street View
Mesh	Ground, 3D Buildings, trees, and clouds	Sphere
Textures	Geollery v1: No texture Geollery v2: With 360° street views	Textured by 360° street views
Availability	Almost always available	Only available for the locations with 360° street view data
Motion	6 DoF	3 DoF + Teleport
Virtual Avatar	Available	Not applicable
Collaboration	Available	Not applicable
Social Media Location Accuracy	Almost the exact location in the world	Estimated by distance and orientation
Virtual Representation	Billboards / Balloons / Framed photos / Doodles / Gifts	Billboards (v2: added balloons and gifts)
Aggregation	Based on spatial relationship	Based on direction and distance

Table 1: Comparison between *Geollery* and *Social Street View* along different variables. Note that the original version of *Social Street View* only uses billboards as a virtual form of social media while the latest version also uses balloons and virtual gifts.

Meshes and Textures

During the design process of selecting meshes and textures, we consider the tradeoff between the processing speed and the visual appearance. While prior art [7, 61, 66] includes various approaches to reconstruct textured 3D buildings in minutes or hours, we prefer a progressive approach to only reconstruct the nearby building geometries. This allows us to create buildings in real time as needed. We circumvent preconstructed models to allow *Geollery* to be used at any location where 2D building data are available.

Social Street View is another approach for real-time rendering of immersive street-level environments with geotagged social media. Nevertheless, it reconstructs textured spheres with depth maps and normal maps rather than 3D building blocks. Since the building geometries are not fully recovered, its degrees of freedom in motion are limited to pitch, roll, and yaw. Users have to *teleport* to the other locations by clicking on the streets or a 2D map.

To achieve six degrees of freedom in movement, we decide to progressively stream data from OpenStreetMap to build 3D meshes in real time. *Geollery* extrudes polygons of nearby buildings into 3D blocks according to metadata such as building heights (usually available in the dense urban areas) and building levels. Although this approach cannot reconstruct complex geometries such as the Effiel Tower or the London Eye, it provides the spatial context necessary for augmented reality scenarios (when the user holds a mobile device).

Figure 4: Chatting in *Geollery* with geotagged social media and virtual buildings provides users with spatial context.

In the first version of *Geollery*, we explore different color schemes for visualizing the mirrored world. Based on participants' feedback from the user study, we added images from *Google Street View* to *Geollery*, so that the closest street views are rendered with the building geometries in real time. We discuss the technical details in Section 6.

Interactive Capabilities

The real-time mirrored world enables new interactive capabilities in *Geollery*. Here, users can see nearby friends as virtual avatars, chat with friends, and paint street art collaboratively on the virtual building walls.

Avatars. First-time visitors to *Geollery* are asked to select a 3D avatar from a collection of 40 rigged models. These models are stored in glTF[9] format for efficient transmission and loading in the WebGL context. After selecting an avatar, users can use the keyboard or the panning gesture on a mobile device to virtually walk in the mirrored world.

Chat. As shown in Figure 4, when two participants virtually meet with each other, *Geollery* allows them to chat with each other in the form of text bubbles. Users can click on other avatars to send private chat messages or their own avatar to send public chat messages.

Collaborative Street Art. Inspired by street art, *Geollery* enables two or more users to share a single whiteboard, draw on it, and add pictures or text via WebSockets. The server updates drawings on nearby users' canvases after every edit enabling real-time collaboration.

Virtual Representations of Social Media

In classic 2D interfaces, social media are usually laid out linearly (*Twitter*, *Instagram*) or on a grid (*Pinterest*) within

[9] glTF: https://github.com/KhronosGroup/glTF, GL Transmission Format.

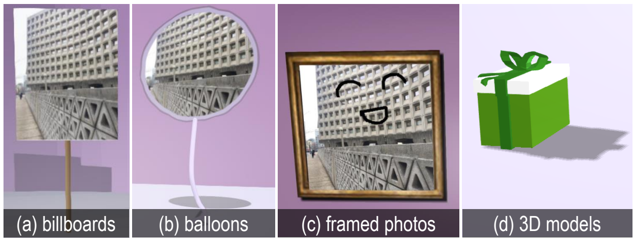

Figure 5: Four virtual representations of geotagged social media: (a) billboards, (b) balloons, (c) framed photography, and (d) 3D models such as gift boxes.

the screen space. Nevertheless, in a 3D space, the virtual forms of social media can have more diversity. In Geollery, we have designed the following four virtual representations of social media:

(a) **Billboards**. Billboards, newsstands, and posters are widely used in the physical world for displaying information. As shown in Figure 5(a), billboards show thumbnails of geotagged images or text. We implement four levels of detail for thumbnails: 64^2, 128^2, 256^2, and 512^2 pixels and progressively load higher resolution thumbnails as users approach different billboards. When users hover over a billboard, it reveals associated text captions, truncated to four lines. When users click on a billboard, a window appears with detail including the complete text caption, the number of likes, and any user comments.

(b) **Balloons**. To attract users' attention and sustain their interest, we design floating balloons in Figure 5(b) to showcase nearby social media. The border colors of balloons categorize their social media based on the text of each social media post.

(c) **Framed photos or street art**. These two representations are inspired by galleries and street art, respectively. *Geollery* allows the users to put on framed photos or a public whiteboard on building walls. Creators of street art can allow nearby users to collaborate in drawing doodles or writing text on the whiteboard. When creating whiteboards, users also have the option of selecting from multiple sizes and frame styles.

(d) **Virtual gifts**. To encourage users to engage with their friends, we design virtual gift boxes. Users can leave a gift box at any location in the world and their friends can open it up and get rewards in the form of a message or a picture. Gifts can also be secured via questions and answers.

Geollery allows users to create billboards, balloons, or gift boxes at their avatar's location by uploading photos or text messages. To create a framed photo or whiteboard, users simply click on or touch an empty section of virtual wall with the drawing mode enabled. *Geollery* hangs the frame outside the building by aligning the normal vectors of the wall and the frame [18].

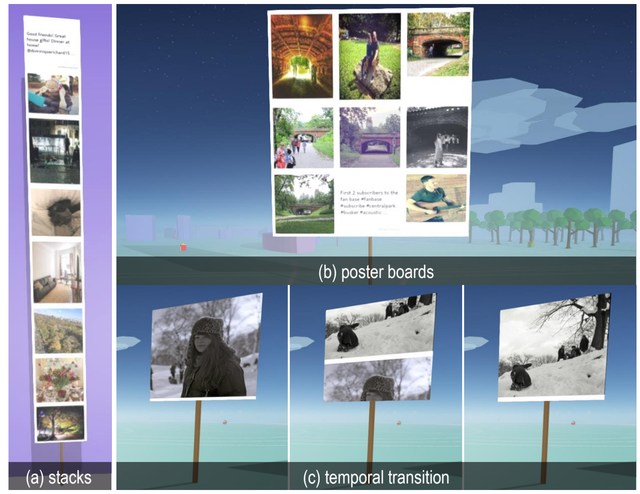

Figure 6: *Geollery* spatially aggregates social media into: (a) stacks, (b) poster boards, or (c) a single billboard or balloon with temporal transition.

Aggregation Approaches

One of the challenges of visualizing a large amount of social media in 3D spaces is visual clutter. When multiple social media are placed close together, their virtual representations may occlude each other. We propose three modes as shown in Figure 6 to resolve this issue:

(a) **Stacks**. This mode stacks older billboards upon the newer ones so that all co-located social media can be viewed at the same time.

(b) **Poster boards**. This mode is similar to the stacks but lays out the social media in a grid on a large poster board. Compared to stacks, posts are not placed as high when more than three are aggregated together.

(c) **Temporal transition**. This mode clusters nearby social media within a radius of approximately 12 meters into a single standard size billboard or balloon. The content displayed dynamically changes between aggregated social media every 10 seconds. This method greatly reduces the visual clutter while displaying the latest information to the user.

The advantage of stacks or poster boards is that multiple posts can be viewed at a glance, while the advantages of temporal transition are reducing the visual clutter and avoiding information overload. We provide all options and discuss the participants' feedback in Section 5.

Privacy

When designing *Geollery*, we take privacy concerns into consideration right at the beginning since location data may reveal details of people's lives [32]. We tackle privacy in two ways:

(1) **Social Media**. When creating social media, users can select among multiple privacy options including: only visible to themselves, only visible to friends, and visible to the public. Although we do not support tagging on photos for now, we note that future systems with the tagging feature should mitigate the multiparty privacy conflicts [59].
(2) **Avatars**. Users can set their avatar to be invisible to prevent exposing themselves to the public. Users can also customize their display name to remain anonymous in the mirrored world.

Real-world Phenomena

As suggested in [18, 30], real-world phenomena such as day and night transitions and changing seasons make virtual worlds more alive and realistic. In *Geollery*, we have designed a day/night transition system which adjusts the lighting and sky based on the local time of the user.

Filtering of Social Media

Applying topic models [5, 43] and temporal filters [24, 64] to social media has been researched intensively in recent years. In *Geollery* v2, we allow the users to filter the social media within the day, month, or year, or by keywords.

5 USER STUDY

To unveil the potential use cases and challenges for designing a 3D mixed-reality social media platform, we evaluated our prototype, *Geollery* v1, against another social media system, *Social Street View*, in a user study with semi-structured interviews. The key differences between the two systems are discussed in Section 4 and Table 1.

We recruited a total of 20 participants (10 females; age range: 21 - 30, with an average of 25.75 and standard deviation of 3.02) via campus email lists and flyers. Each participant was paid 10 dollars as compensation for their time and effort. None of the participants had been involved with this project before. The individual semi-structured interviews took place in a quiet room using two side-by-side workstations with 27-inch displays and NVIDIA GTX 1080 graphics cards. Participants interacted with the systems using keyboards and mice alongside the interviewer. The session for each participant lasted between 45 – 60 minutes and involved four stages: a background interview, an exploration of *Geollery* and *Social Street View*, a quantitative evaluation, and a discussion about the future of 3D social media platforms.

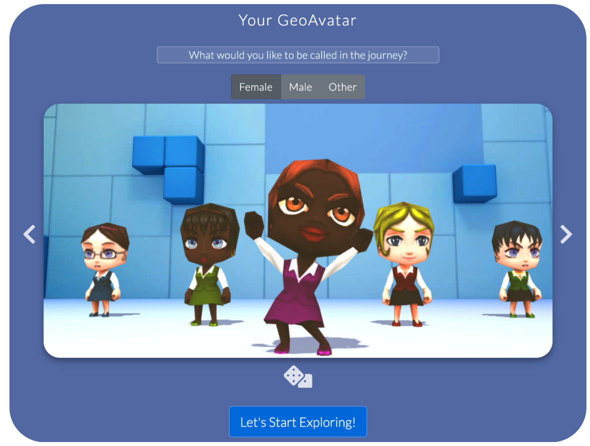

Figure 7: The welcome interface of Geollery.

Background Interview

In the first stage (5 minutes), the interviewer introduced *Geollery* and asked the participant about their prior experiences of social media. All of our participants reported social media usage of at least several times per week with few actively posting. Furthermore, 16 out of 20 responded with usage of several times per day. However, only 5 out of 20 actively posted social media frequently: "*I post news about sports and games every day.* (P7/M)"; "*I majorly use Instagram, I post from my own portfolio.* (P17/F)". The rest of our participants primarily use social media for viewing friends' updates and photos.

Exploration of *Geollery* and *Social Street View*

In the second stage (30-40 minutes), the interviewer instructed the participant to virtually visit four places using each of the target systems, *Geollery* and *Social Street View*. Participants were asked to explore the university campus where the study took place, the Manhattan District of New York, the National Gallery of Art in Washington D.C, and another location of the participant's choice. We counterbalanced the order of system conditions (*Geollery* or *Social Street View*), as well as the order of the three places using the Latin square design [28]. For the duration of the study, the interviewer observed the participants' behaviors and took notes about their comments and interaction.

First, the participant was asked to choose a nickname and an avatar from welcome interface in Figure 7. Meanwhile, the interviewer logged in to the same system on the other workstation so the participant could virtually interact with the interviewer.

Next, we asked if the participant was aware of their location in each virtual setting. In *Social Street View*, all participants quickly figured out their virtual locations. In *Geollery*, participants who noticed the minimap would immediately know where they were, but four out of 20 users became confused. For example, P5/F asked: *"Am I in a gallery?"*, and P16/M responded: *"I believe I am in a museum."*

After allowing the participants to freely explore each interface for 3 minutes, we interviewed them about their first impressions. In *Geollery*, many participants were amazed by walking in the mirrored world and the progressive loading of the geometries: *"I think it's a very good start, it's very good experience to walk around."* (P6/F); *"I like that the buildings are forming while I am walking."* (P16/M); *"I really like the fact that it's scaled, so I don't have to walk 15 minutes from one place to the other."* (P17/F).

In *Social Street View*, many participants appreciated the texturing of the 360° views: *"I think the texturing actually helps me."* (P17, F); *"It's like you don't have to be there."* (P11, F).

However, several participants found navigating *Social Street View* frustrating as they could not freely *walk* around, only *teleport* by clicking the mouse: *"So how do I walk here?"* (P5/F) [The interviewer instructed her how to teleport.] *"Oh, I see, it zooms in when I scroll. It's like Google Street View."*

We further asked their preferences of different virtual representations, different aggregation methods, and whether they preferred the system to read out the social media contents on demand. In regards to billboards versus balloons, 14 out of 20 participants preferred balloons: *"Balloons are informal and billboards can have notices. Balloons may be better for social media."* (P13/F). The other 6 participants preferred billboards: *"I like billboards. First thing, balloons keep moving, it's a little distracting. Billboards look like you are announcing something. It's more neat"* (P17/F). In addition, 75% of participants preferred the temporal transition approach to aggregate nearby social media into one billboard or balloons and 80% users preferred audio on demand.

In the end, we encouraged the participants to input any desired location and compare *Geollery* with *Social Street View*. Most participants chose their homes while a few participants input locations where *only Geollery* is available. For example, P12/M typed the Statue of Liberty in New York City, where only *Geollery* was able to present the geotagged social media with its spatial context, the Liberty Island.

Quantitative Evaluation

After exploring the two interfaces for 30 minutes, we asked the participants to comparatively and quantitatively rate the two systems along 9 attributes in an AttrakDiff-based antonym word pair questionnaire inspired by [30]. The average ratings are visualized as a radar chart in Figure 8. From a Welch's paired t-test, we found a significant effect for interactivity ($t(20) = 3.04, p < 0.01$, Cohen's $d = 0.83$) and creativity ($t(20) = 2.10, p < 0.05$, Cohen's $d = 0.66$) with *Geollery* outperforming *Social Street View*. In addition, 14 out of 20 found *Geollery* more or equally immersive compared to *Social Street View* and 16 out of 20 found *Geollery* more or equally entertaining compared to *Social Street View*.

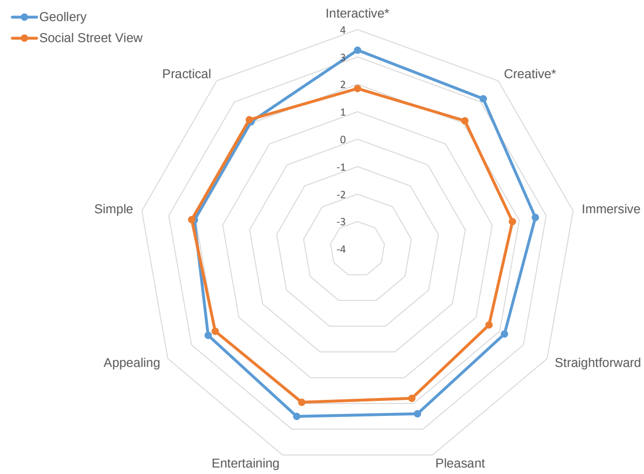

Figure 8: The radar chart visualizes the quantitative evaluation between *Geollery* and *Social Street View* along 9 dimensions with 20 participants. With a Welch's paired t-test, there is a significant effect ($p < 0.05$) that *Geollery* is rated more interactive and more creative than *Social Street View*.

We then asked the participant which appealed more to them. Overall, more participants (13 out of 20) preferred *Geollery* to *Social Street View* due to its interactivity: *"I prefer Geollery in terms of moving around, and because you have the options to draw on walls and interact with people."* (P17/F); *"I like Geollery because I have free roaming there, and it's kind of cool that I can walk over the world."* (P11/F).

Several participants pointed out that *Geollery* is more similar to a massively multiplayer online (MMO) game: *"That one* (*Geollery*) *I was in a game. This one* (*Social Street View*) *feels depressing, nothing exciting. (P18/M) "I think it's more like a game, it's more fun to interact with the virtual world."* (P19/M) *"Having more people makes the place feel more interesting and immersive."* (P10/F).

In this study, the participants did not try *Geollery* v2 with textured buildings. Some participants preferred *Social Street View* due to the immersive panoramas: *"[In Geollery,] the buildings don't like the buildings in the real world, but Social Street View allows me to explore my environment."* (P13/F); *"I like Social Street View better. There, I understand the environments better."*.

The Future of 3D Social Media Platforms

At the end of the user study, we asked the participant to discuss their demands as users of future 3D social media platforms and the features they would add if they were designers or product managers. We interviewed the participants with the following three questions:

1. Suppose that we have a polished 3D social media platform like Geollery or Social Street View, how much time would you like to spend on it?

For this question, we categorize our participants into three classes: supporters, followers, protesters. Supporters (75%, 15 out of 20) are generally more optimistic about the future of 3D social media platforms. They envision *Geollery* or *Social Street View* being used for daily exploration or trip planning. Here are some responses: *"I would like to use it every day when I go to work, or travel during weekends. [...] I may spend about 8 hours per week using it."* (P4/M); *"If it's not distracting like Facebook and Instagram, I would use it every day on a couple of things."* (P17/F); *"I love travelling, [so] I would like to use it [Social Street View] to preview my destinations before my trips."* (P3/M) .

The followers (4 out of 20) typically preferred to switch to 3D social media platforms once their friends joined. For example, here are some followers' responses: *"I am a follower on most social media sites. I would only join a 3D social media platform once my friends are there."* (P4/M); *"If my friends are all on this, I can see myself spend a couple of hours every week. We can have a meet-up point at one place. My friends could go to my home and post social media."* (P12/M).

As for protesters (1 out of 20), P2/F responded: *"I don't think I will use this. I prefer to use Yelp to see comments [of nearby restaurants]."*

2. Can you imagine your use cases for Geollery and Social Street View? What would you like to use 3D social media platforms for?

Many participants (17 out of 20) mentioned food and travel planning as their majority use cases: *"I would like to use it for the food in different restaurants. I am always hesitating of different restaurants. It will be very easy to see all restaurants with street views. In Yelp, I can only see one restaurant."* (P13/F); *"[I will use it for] exploring new places. If I am going on vacation somewhere, I could immerse myself into the location. If there are avatars around that area, I could ask questions."* (P17/F).

Family gathering and virtual parties are also potential use cases according to the participants' responses: *"I think it* (*Geollery*) *will be useful for families. I just taught my grandpa how to use Facetime last week and it would great if I could teleport to their house and meet with them, then we could chat and share photos with our avatars."* (P2/F); *"...for communicating with my families, maybe, and distant friends, [so] they can see New York. And, getting to know more people, connecting with people based on similar interests."* (P19/M); *"We can use it (Geollery) on parties [...] like hide some gifts around the house and ask people to find."* (P4/M).

3. If you were a designer or product manager for Geollery or Social Street View, what features would you like to add to the systems?

Many participants mentioned texturing the buildings on *Geollery* v1: *"A mapping of the texture, high-resolution texture, will be great."* (P12/M); *"if there is a way to unify the interaction between them, there will be more realistic buildings [and] you could have more roof structures. Terrains will be interesting to add on.* (P18/M).

Participants suggested more data be integrated into 3D social media platforms: *"If I'm shopping around in a mall, if I could see deals and coupons, and live comments..."* (P7/M); *"I would like to add traffic and parking information."* (P6/F).

Many participants also suggested adding a better avatar system, more 3D objects, and additional interactive capabilities in *Geollery*: *"[I would like] the flexibility to build your own avatar. Customizing avatar will be one useful feature."* (P18/M); *"I would like to see kitties and puppies running around, and birds flying in the air."* (P13/F); *"I could also add a bike, add a vehicle, a motorcycle in Geollery, this will add some fun."* (P17/F).

6 DISCUSSION

In this section, we summarize the key insights we learned from the user study, as well as the further improvements we have made since the user study.

Insights from User Study

From the user study with 20 participants, we summarize our findings and insights as follows:

(1) Data sources of social media play key roles in developing a 3D social media platform. Since many users do not post to social media frequently, obtaining high-quality data from external sources or seed users to generate high-quality content is of great significance.
(2) Interactivity and panoramic textures have different levels of importance for different groups of users. Users with better geospatial awareness may appreciate more on interactivity in *Geollery* while others may appreciate more on the panoramic texturing.
(3) Customization of avatars, diversity [9], and accessibility [41] are important for developing future 3D social media platforms. All users should be able to represent themselves and share the virtual mirrored world equally.

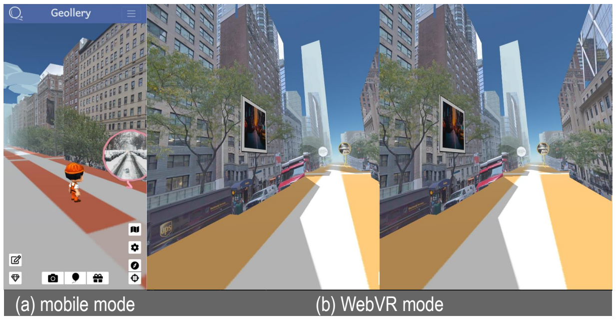

Figure 9: We further combine the best features of *Geollery* and *Social Street View* by texturing the buildings in *Geollery* v2: (a) shows a screenshot on an Android phone, where users can track their current location and orientation while exploring Geollery; (b) shows a screenshot in the VR mode, where users are immersed in a first-person experience and are able to walk around via an XBOX controller.

Limitations of our User Study

We recruited participants for our user study via email lists and on-campus flyers. Therefore, most of our participants were undergraduate and graduate students at our campus. As the ages of our participants spanned 20 - 31, the results of our study may not generalize to other populations such as older adults who may prefer realism over walking around. As our user study was conducted in a constrained environment with only the user and the interviewer in the system, it is unclear how the users may perceive our system with more users. Some users may feel less comfortable in crowded virtual environments. Furthermore, as users of our system were only able to engage with it for 30 minutes, their self-assessment of whether and how they would regularly use a mixed-reality social media system such as Geollery is not fully conclusive.

Combining *Geollery* and *Social Street View*

Thanks to the participants' feedback, we have developed *Geollery* v2 which combines progressive geometries with street views to create textured buildings. We achieve this by projecting street views from Google Street View onto the building geometries in *Geollery*. As users walk around in *Geollery*, we continuously update the closest street view and use alpha blending to transition to textures obtained from new street views. This approach works for many urban areas in *real-time*. Nevertheless, as shown in Figure 9, this algorithm may project trees or the sky onto the geometries due to the approximation of the digital city. Accurate real-time creation of textured buildings from street view images remains an open challenge even in the state-of-the-art reconstruction systems [3, 61, 66]. Future development may take advantage of deep neural networks to semantically segment the sky [4, 11, 25] or in-paint the pedestrians [20, 46].

7 CONCLUSION AND FUTURE WORK

In this paper, we present *Geollery*, an interactive social media platform in mixed reality. We introduce our system architecture, design choices, and implementation details. We conduct a user study with semi-structured interviews to examine the challenges and limitations of the interfaces, as well as the types of decisions these could influence and their potential impact. The quantitative results indicate that *Geollery* is more interactive and creative than *Social Street View*. The user responses reveal several key use cases including searching for food, travel planning, and family gathering. Taking the participants' feedback into account, we combine *Geollery* and *Social Street View* by texturing buildings using street views.

There are several future directions for improving *Geollery*. First, we plan to fuse multiple street views onto the building geometries in real-time to achieve better photo-realistic rendering. Second, we aim to integrate additional useful information into the 3D world such as geotagged sales, services, and job listings. Adding mental health and sentiments [42, 67] extracted from social media and live surveillance videos [15, 53] could prove useful for social good. Third, we intend to use techniques from previous research [14] to improve the filtering mechanism, encouraging supportive comments and reducing negative emotions in *Geollery*.

We imagine *Geollery* existing as a standalone social media platform for those looking to explore new areas or looking to share their experiences. Currently, the majority of social media in *Geollery* is from external sources. However, as the *Geollery* community grows, we expect internal media posted to *Geollery* quickly becoming the predominant source. As we obtain more users in *Geollery*, we envision 3D social media playing a significant role in the realm of virtual and augmented reality. Users may eventually move from traditional text and 2D media into 3D. For example, VideoFields [15], Holoportation [47], and Montage4D [16, 17] systems take multiview videos and convert them into 3D visualization or stylized holograms in real time. Such techniques may eventually change the way we consume and create data, as well as the way we socialize with other people.

ACKNOWLEDGEMENT

We would also like to thank *Sai Yuan, Akanksha Shrivastava, Xiaoxu Meng, Shuo Li, Eric Krokos, Xuetong Sun*, all user study participants, and the anonymous reviewers for the insightful feedback on the Geollery system and the manuscript.

This work has been supported in part by the NSF Grants 1823321, 1564212, 1429404, and the State of Maryland's MPower initiative. Any opinions, findings, conclusions, or recommendations expressed in this article are those of the authors and do not necessarily reflect the views of the research sponsors.

REFERENCES

[1] Sameer Agarwal, Noah Snavely, Ian Simon, Steven M. Sietz, and Rick Szeliski. 2009. Building Rome in a Day. In *2009 IEEE 12th International Conference on Computer Vision (CVPR)*. IEEE, New York, NY, USA, 72–79. https://doi.org/10.1145/2001269.2001293

[2] Toni Alatalo, Timo Koskela, Matti Pouke, Paula Alavesa, and Timo Ojala. 2016. VirtualOulu: Collaborative, Immersive and Extensible 3D City Model on the Web. In *Proceedings of the 21st International Conference on Web3D Technology (Web3D)*. ACM, ACM, New York, NY, USA, 95–103. https://doi.org/10.1145/2945292.2945305

[3] Mahmoud Badri, Minna Pakanen, Paula Alavesa, Hannu Kukka, and Timo Ojala. 2017. Design, Development, and Usability Evaluation of a System for Adding and Editing Social Media Banners in the Immersive Street-Level 3D Virtual City. In *2017 9th International Conference on Virtual Worlds and Games for Serious Applications (VS-Games)*. ACM, New York, NY, USA, 102–108. https://doi.org/10.1109/VS-GAMES.2017.8056577

[4] Vijay Badrinarayanan, Ankur Handa, and Roberto Cipolla. 2015. SegNet: a Deep Convolutional Encoder-Decoder Architecture for Robust Semantic Pixel-Wise Labelling. *ArXiv Preprint ArXiv:1505.07293* (2015).

[5] David M Blei, Andrew Y Ng, and Michael I Jordan. 2003. Latent Dirichlet Allocation. *Journal of Machine Learning Research* 3, Jan (2003), 993–1022.

[6] Jan Brejcha, Michal Lukác, Zhili Chen, Stephen DiVerdi, and Martin Cadík. 2018. Immersive Trip Reports. In *Proceedings of the 31st Annual ACM Symposium on User Interface Software and Technology (UIST '18)*. ACM, New York, NY, USA, 389–401. https://doi.org/10.1145/3242587.3242653

[7] Abdullah Bulbul and Rozenn Dahyot. 2017. Social Media Based 3D Visual Popularity. *Computers & Graphics* 63 (2017), 28–36. https://doi.org/10.1016/j.cag.2017.01.005

[8] O'Sullivan Carol and Cathy Ennis. 2011. Metropolis: Multisensory Simulation of a Populated City. In *2011 Third International Conference on Games and Virtual Worlds for Serious Applications*. IEEE, IEEE, New York, NY, USA, 1–7. https://doi.org/10.1109/VS

[9] Matthew Carrasco and Andruid Kerne. 2018. Queer Visibility: Supporting LGBT+ Selective Visibility on Social Media. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI)*. ACM, ACM, New York, NY, USA, 250. https://doi.org/10.1145/3173574.3173824

[10] Junghoon Chae, Dennis Thom, Harald Bosch, Yun Jang, Ross Maciejewski, David S Ebert, and Thomas Ertl. 2012. Spatiotemporal Social Media Analytics for Abnormal Event Detection and Examination Using Seasonal-Trend Decomposition. In *2012 IEEE Conference on Visual Analytics Science and Technology (VAST)*. IEEE, IEEE, New York, NY, USA, 143–152. https://doi.org/10.1109/VAST.2012.6400557

[11] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. 2018. Deeplab: Semantic Image Segmentation With Deep Convolutional Nets, Atrous Convolution, and Fully Connected Crfs. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 40, 4 (2018), 834–848. https://doi.org/10.1109/TPAMI.2017.2699184

[12] Siming Chen, Xiaoru Yuan, Zhenhuang Wang, Cong Guo, Jie Liang, Zuchao Wang, Xiaolong Luke Zhang, and Jiawan Zhang. 2016. Interactive Visual Discovering of Movement Patterns From Sparsely Sampled Geo-Tagged Social Media Data. *IEEE Transactions on Visualization and Computer Graphics* 22, 1 (2016), 270–279. https://doi.org/10.1109/TVCG.2015.2467619

[13] Justin Cranshaw, Raz Schwartz, Jason I. Hong, and Norman M. Sadeh. 2. The Livehoods Project: Utilizing Social Media to Understand the Dynamics of a City. In *ICWSM*, John G. Breslin, Nicole B. Ellison, James G. Shanahan, and Zeynep Tufekci (Eds.). The AAAI Press, New York, NY, USA, 8. http://www.aaai.org/ocs/index.php/ICWSM/ICWSM12/paper/view/4682

[14] Robert Deloatch, Brian P Bailey, Alex Kirlik, and Craig Zilles. 2017. I Need Your Encouragement!: Requesting Supportive Comments on Social Media Reduces Test Anxiety. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI)*. ACM, ACM, New York, NY, USA, 736–747. https://doi.org/10.1145/3025453.3025709

[15] Ruofei Du, Sujal Bista, and Amitabh Varshney. 2016. Video Fields: Fusing Multiple Surveillance Videos into a Dynamic Virtual Environment. In *Proceedings of the 21st International Conference on Web3D Technology (Web3D)*. ACM, Anaheim, California, 165–172. https://doi.org/10.1145/2945292.2945299

[16] Ruofei Du, Ming Chuang, Wayne Chang, Hugues Hoppe, and Amitabh Varshney. 2018. Montage4D: Interactive Seamless Fusion of Multiview Video Textures . In *Proceedings of ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games (I3D) (I3D)*. ACM, Montreal, Quebec, Canada, 124–133. https://doi.org/10.1145/3190834.3190843

[17] Ruofei Du, Ming Chuang, Wayne Chang, Hugues Hoppe, and Amitabh Varshney. 2018. Montage4D: Real-time Seamless Fusion and Stylization of Multiview Video Textures. *Journal of Computer Graphics Techniques* 7, 4 (10 December 2018), 1–34. https://doi.org/10.1145/3190834.3190843

[18] Ruofei Du and Amitabh Varshney. 2016. Social Street View: Blending Immersive Street Views With Geo-Tagged Social Media. In *Proceedings of the 21st International Conference on Web3D Technology (Web3D)*. ACM, New York, NY, USA, 77–85. https://doi.org/10.1145/2945292.2945294

[19] Ruofei Du and Amitabh Varshney. 2018. Systems, Devices, and Methods for Generating a Social Street View. US Patent App. 15/559,955.

[20] Arturo Flores and Serge Belongie. 2010. Removing Pedestrians From Google Street View Images. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition-Workshops*. IEEE, IEEE, New York, NY, USA, 53–58. https://doi.org/10.1109/CVPRW.2010.5543255

[21] David Gelernter. 1993. *Mirror Worlds: Or: The Day Software Puts the Universe in a Shoebox... How It Will Happen and What It Will Mean*. Oxford University Press, New York, NY, USA.

[22] Marcus Goetz and Alexander Zipf. 2013. The Evolution of Geo-Crowdsourcing: Bringing Volunteered Geographic Information to the Third Dimension. In *Crowdsourcing Geographic Knowledge: Volunteered Geographic Information (VGI) in Theory and Practice*. Springer, New York, NY, USA, 139–159. https://doi.org/10.1007/978-94-007-4587-_9

[23] Ming C Hao, Christian Rohrdantz, Halldor Janetzko, Daniel A Keim, Umeshwar Dayal, Lars Erik Haug, Meichun Hsu, and Florian Stoffel. 2013. Visual Sentiment Analysis of Customer Feedback Streams Using Geo-Temporal Term Associations. *Information Visualization* 12, 3-4 (2013), 273–290. https://doi.org/10.1177/1473871613481691

[24] Myung-Hwa Hwang, Shaowen Wang, Guofeng Cao, Anand Padmanabhan, and Zhenhua Zhang. 2013. Spatiotemporal Transformation of Social Media Geostreams: a Case Study of Twitter for Flu Risk Analysis. In *Proceedings of the 4th ACM SIGSPATIAL International Workshop on GeoStreaming*. ACM, ACM, New York, NY, USA, 12–21. https://doi.org/10.1145/2534303.2534310

[25] Forrest N Iandola, Song Han, Matthew W Moskewicz, Khalid Ashraf, William J Dally, and Kurt Keutzer. 2016. Squeezenet: Alexnet-Level Accuracy With 50x Fewer Parameters And< 0.5 Mb Model Size. *CoRR* abs/1602.07360 (2016). http://arxiv.org/abs/1602.07360

[26] Cheuk Yiu Ip, M. Adil Yalçin, David Luebke, and Amitabh Varshney. 2013. PixelPie: Maximal Poisson-Disk Sampling With Rasterization. In *Proceedings of the 5th High-Performance Graphics Conference (HPG '13)*. ACM, New York, NY, USA, 17–26. https://doi.org/10.1145/2492045.2492047

[27] Toru Ishida. 2002. Digital City Kyoto. *Commun. ACM* 45, 7 (2002), 76–81. https://doi.org/10.1007/1140754_8

[28] A Donald Keedwell and József Dénes. 2015. *Latin Squares and Their Applications* (second edition ed.). North-Holland, New York, NY, USA.

[29] Seokyeon Kim, Seongmin Jeong, Insoo Woo, Yun Jang, Ross Maciejewski, and David S Ebert. 2018. Data Flow Analysis and Visualization for Spatiotemporal Statistical Data Without Trajectory Information. *IEEE Transactions on Visualization and Computer Graphics* 24, 3 (March 2018), 1287–1300. https://doi.org/10.1109/TVCG.2017.2666146

[30] Hannu Kukka, Minna Pakanen, Mahmoud Badri, and Timo Ojala. 2017. Immersive Street-Level Social Media in the 3D Virtual City: Anticipated User Experience and Conceptual Development. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW)*. ACM, ACM, New York, NY, USA, 2422–2435. https://doi.org/10.1145/2998181.2998341

[31] Takeshi Kurashima, Tomoharu Iwata, Go Irie, and Ko Fujimura. 2010. Travel Route Recommendation Using Geotags in Photo Sharing Sites. In *Proceedings of the 19th ACM International Conference on Information and Knowledge Management*. ACM, ACM, New York, NY, USA, 579–588. https://doi.org/10.1145/1871437.1871513

[32] Ilaria Liccardi, Alfie Abdul-Rahman, and Min Chen. 2016. I Know Where You Live: Inferring Details of People's Lives by Visualizing Publicly Shared Location Data. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI)*. ACM, ACM, New York, NY, USA, 1–12. https://doi.org/10.1145/3137616.3137617

[33] Michael D. Lieberman and Hanan Samet. 2011. Multifaceted Toponym Recognition for Streaming News. In *Proceedings of the 34th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '11)*. ACM, New York, NY, USA, 843–852. https://doi.org/10.1145/2009916.2010029

[34] Michael D. Lieberman and Hanan Samet. 2012. Adaptive Context Features for Toponym Resolution in Streaming News. In *Proceedings of the 35th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '12)*. ACM, New York, NY, USA, 731–740. https://doi.org/10.1145/2348283.2348381

[35] Michael D. Lieberman and Hanan Samet. 2012. Supporting Rapid Processing and Interactive Map-Based Exploration of Streaming News. In *Proceedings of the 20th International Conference on Advances in Geographic Information Systems (SIGSPATIAL '12)*. ACM, New York, NY, USA, 179–188. https://doi.org/10.1145/2424321.2424345

[36] Michael D. Lieberman, Hanan Samet, and Jagan Sankaranayananan. 2010. Geotagging: Using Proximity, Sibling, and Prominence Clues to Understand Comma Groups. In *Proceedings of the 6th Workshop on Geographic Information Retrieval (GIR '10)*. ACM, New York, NY, USA, Article 6, 8 pages. https://doi.org/10.1145/1722080.1722088

[37] Yafeng Lu, Xia Hu, Feng Wang, Shamanth Kumar, Huan Liu, and Ross Maciejewski. 2015. Visualizing Social Media Sentiment in Disaster Scenarios. In *Proceedings of the 24th International Conference on World Wide Web*. ACM, ACM, New York, NY, USA, 1211–1215. https://doi.org/10.1145/2740908.2741720

[38] Alan M. MacEachren, Francis P. Boscoe, Daniel Haug, and Linda Pickle. 1998. Geographic Visualization: Designing Manipulable Maps for Exploring Temporally Varying Georeferenced Statistics. In *Proceedings of the 1998 IEEE Symposium on Information Visualization (INFOVIS '98)*. IEEE, IEEE Computer Society, Washington, DC, USA, 87–. https://doi.org/10.1109/INFVIS.1998.729563

[39] Alan M MacEachren, Anuj Jaiswal, Anthony C Robinson, Scott Pezanowski, Alexander Savelyev, Prasenjit Mitra, Xiao Zhang, and Justine Blanford. 2011. Senseplace2: Geotwitter Analytics Support for Situational Awareness. In *2011 IEEE Conference on Visual Analytics Science and Technology (VAST)*. IEEE, IEEE, New York, NY, USA, 181–190. https://doi.org/10.1109/VAST.2011.6102456

[40] Ross Maciejewski, Stephen Rudolph, Ryan Hafen, Ahmad Abusalah, Mohamed Yakout, Mourad Ouzzani, William S Cleveland, Shaun J Grannis, and David S Ebert. 2010. A Visual Analytics Approach to Understanding Spatiotemporal Hotspots. *IEEE Transactions on Visualization and Computer Graphics* 16, 2 (2010), 205–220. https://doi.org/10.1145/2820783.2820817

[41] Haley MacLeod, Cynthia L Bennett, Meredith Ringel Morris, and Edward Cutrell. 2017. Understanding Blind People's Experiences With Computer-Generated Captions of Social Media Images. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. ACM, ACM, New York, NY, USA, 5988–5999. https://doi.org/10.1145/3025453.3025814

[42] Lydia Manikonda and Munmun De Choudhury. 2017. Modeling and Understanding Visual Attributes of Mental Health Disclosures in Social Media. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI)*. ACM, ACM, New York, NY, USA, 170–181. https://doi.org/10.1145/3025453.3025932

[43] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient Estimation of Word Representations in Vector Space. *ArXiv Preprint ArXiv:1301.3781* (2013). https://arxiv.org/pdf/1301.3781.

[44] Hatem Mousselly-Sergieh, Daniel Watzinger, Bastian Huber, Mario Döller, Elöd Egyed-Zsigmond, and Harald Kosch. 2014. World-Wide Scale Geotagged Image Dataset for Automatic Image Annotation and Reverse Geotagging. In *Proceedings of the 5th ACM Multimedia Systems Conference (MMSys '14)*. ACM, New York, NY, USA, 47–52. https://doi.org/10.1145/2557642.2563673

[45] Przemyslaw Musialski, Peter Wonka, Daniel G Aliaga, Michael Wimmer, Luc Van Gool, and Werner Purgathofer. 2013. A Survey of Urban Reconstruction. *Computer Graphics Forum* 32, 6 (2013), 146–177. https://doi.org/10.1111/cgf.12077

[46] Angelo Nodari, Marco Vanetti, and Ignazio Gallo. 2012. Digital Privacy: Replacing Pedestrians From Google Street View Images. In *2012 21st International Conference on Pattern Recognition (ICPR)*. IEEE, IEEE, New York, NY, USA, 2889–2893. https://doi.org/10.1109/CVPRW.2010.5543255

[47] Sergio Orts-Escolano, Christoph Rhemann, Sean Fanello, Wayne Chang, Adarsh Kowdle, Yury Degtyarev, David Kim, Philip L Davidson, Sameh Khamis, Mingsong Dou, Vladimir Tankovich, Charles Loop, Philip A.Chou, Sarah Mennicken, Julien Valentin, Vivek Pradeep, Shenlong Wang, Sing Bing Kang, Pushmeet Kohli, Yuliya Lutchyn, Cem Keskin, and Shahram Izadi. 2016. Holoportation: Virtual 3D Teleportation in Real-Time. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology (UIST)*. ACM, Tokyo, Japan, 741–754. https://doi.org/10.1145/2984511.2984517

[48] Erik Palmquist and Jonathan Shaw. 2008. Collaborative City Modeling. *Architecture in Computro* (2008), 249–256. https://doi.org/10.1007/978-3-642-15690-_9

[49] Wade Roush. 2007. Second Earth. *Technology Review* (2007), 38.

[50] Bc Russell and R Martin-Brualla. 2013. 3D Wikipedia: Using Online Text to Automatically Label and Navigate Reconstructed Geometry. *ACM Transactions on Graphics (TOG)* 32, 6 (2013), 1–10. https://doi.org/10.1145/2508363.2508425

[51] Hanan Samet, Marco D Adelfio, Brendan C Fruin, Michael D Lieberman, and Jagan Sankaranarayanan. 2013. PhotoStand: a Map Query Interface for a Database of News Photos. *Proceedings of the VLDB Endowment* 6, 12 (2013), 1350–1353.

[52] Hanan Samet, Jagan Sankaranarayanan, Michael D. Lieberman, Marco D. Adelfio, Brendan C. Fruin, Jack M. Lotkowski, Daniele Panozzo, Jon Sperling, and Benjamin E. Teitler. 2014. Reading News With Maps by Exploiting Spatial Synonyms. *Commun. ACM* 57, 10 (2014), 64–77. https://doi.org/10.1145/2629572

[53] Aswin C. Sankaranarayanan, Rob Patro, Pavan Turaga, Amitabh Varshney, and Rama Chellappa. 2009. Modeling and Visualization of Human Activities for Multicamera Networks. *EURASIP Journal on Image and Video Processing* 2009, 259860 (22 Oct 2009), 1–13. https://doi.org/10.1016/j.patrec.2012.07.005

[54] Jagan Sankaranarayanan, Hanan Samet, Benjamin E Teitler, Michael D Lieberman, and Jon Sperling. 2009. Twitterstand: News in Tweets. In *Proceedings of the 17th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*. ACM, ACM, New York, NY, USA, 42–51. https://doi.org/10.1145/1653771.1653781

[55] Arno Scharl, Alexander Hubmann-Haidvogel, Albert Weichselbraun, Gerhard Wohlgenannt, Heinz-Peter Lang, and Marta Sabou. 2012. Extraction and Interactive Exploration of Knowledge From Aggregated News and Social Media Content. In *Proceedings of the 4th ACM SIGCHI Symposium on Engineering Interactive Computing Systems*. ACM, ACM, New York, NY, USA, 163–168. https://doi.org/10.1145/2305484.2305511

[56] Noah Snavely, Steven M Seitz, and Richard Szeliski. 2006. Photo Tourism: Exploring Photo Collections in 3D. *ACM Transactions on Graphics (TOG)* 25, 3 (2006), 835–846. https://doi.org/10.1145/1179352.1141964

[57] Noah Snavely, Steven M Seitz, and Richard Szeliski. 2008. Modeling the World From Internet Photo Collections. *International Journal of Computer Vision* 80, 2 (2008), 189–210. https://doi.org/10.1007/s11263

[58] Anthony Stefanidis, Andrew Crooks, and Jacek Radzikowski. 2013. Harvesting Ambient Geospatial Information From Social Media Feeds. *GeoJournal* 78, 2 (2013), 319–338. https://doi.org/10.1007/s10708-011-9438-2

[59] Jose M Such, Joel Porter, Sören Preibusch, and Adam Joinson. 2017. Photo Privacy Conflicts in Social Media: A Large-Scale Empirical Study. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI)*. ACM, ACM, New York, NY, USA, 3821–3832. https://doi.org/10.1145/3025453.3025668

[60] Benjamin E. Teitler, Michael D. Lieberman, Daniele Panozzo, Jagan Sankaranarayanan, Hanan Samet, and Jon Sperling. 2008. NewsStand: a New View on News. In *Proceedings of the 16th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems (GIS '08)*. ACM, New York, NY, USA, 10. https://doi.org/10.1145/1463434.1463458

[61] Akihiko Torii, Michal Havlena, and Tomáš Pajdla. 2009. From Google Street View to 3D City Models. In *2009 IEEE 12th International Conference on Computer Vision Workshops, ICCV Workshops*. IEEE, IEEE, New York, NY, USA, 2188–2195. https://doi.org/10.1109/ICCVW.2009.5457551

[62] Matthias Uden and Alexander Zipf. 2013. Open Building Models: Towards a Platform for Crowdsourcing Virtual 3D Cities. In *Progress and New Trends in 3D Geoinformation Sciences*. Springer, New York, NY, USA, 299–314. https://doi.org/10.1007/978-3-642-29793-_17

[63] Andrea Vaccari, Mauro Martino, Francisca Rojas, and Carlo Ratti. 2010. Pulse of the City: Visualizing Urban Dynamics of Special Events. In *Proceedings of the 20th International Conference on Computer Graphics and Vision (GraphiCon)*. ACM, New York, NY, USA, 1–10. https://doi.org/10.1145/2700478

[64] Sarah Vieweg, Amanda L Hughes, Kate Starbird, and Leysia Palen. 2010. Microblogging During Two Natural Hazards Events: What Twitter May Contribute to Situational Awareness. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, ACM, New York, NY, USA, 1079–1088. https://doi.org/10.1145/1753326.1753486

[65] Chaolun Xia, Raz Schwartz, Ke Xie, Adam Krebs, Andrew Langdon, Jeremy Ting, and Mor Naaman. 2014. CityBeat: Real-Time Social Media Visualization of Hyper-Local City Data. In *Proceedings of the 23rd International Conference on World Wide Web*. ACM, ACM, New York, NY, USA, 167–170. https://doi.org/10.1145/2567948.2577020

[66] Jianxiong Xiao, Tian Fang, Peng Zhao, Maxime Lhuillier, and Long Quan. 2009. Image-Based Street-Side City Modeling. *ACM Transactions on Graphics (TOG)* 28, 5, Article 114 (2009), 12 pages. https://doi.org/10.1145/1661412.1618460

[67] Weiwei Yang, Xuetong Sun, and Ruofei Du. 2015. Learning Depression Patterns from MyPersonality and Reddit.

[68] Nick Yee. 2006. The Demographics, Motivations, and Derived Experiences of Users of Massively Multi-User Online Graphical Environments. *Presence: Teleoperators and Virtual Environments* 15, 3 (2006), 309–329. https://doi.org/10.1007/1-4020-3898-_9

[69] Jie Yin, Sarvnaz Karimi, Andrew Lampert, Mark Cameron, Bella Robinson, and Robert Power. 2015. Using Social Media to Enhance Emergency Situation Awareness. *IEEE Intelligent Systems* 27, 6 (2015), 4234–4238. https://doi.org/10.1109/MIS.2012.6