VIRTUAL DATA SPHERE: INVERSE STEREOGRAPHIC PROJECTION FOR IMMERSIVE MULTI-PERSPECTIVE GEOVISUALIZATION

Maxim Spur[1,*], Vincent Tourre[2], Guillaume Moreau[3], Patrick Le Callet[4]

[1]Scene Grammar Lab, Goethe University Frankfurt, Germany — spur@psych.uni-goethe.de
[2]AAU UMR CNRS/MCC 1563, Ecole Centrale de Nantes, France — vincent.tourre@ec-nantes.fr
[3]LAB-STICC, IMT Atlantique Bretagne-Pays de la Loire, Brest, France
[4]Group IPI Lab LS2N, LUNAM University, University of Nantes, France

Commission IV, WG IV/9

KEY WORDS: Geovisualization, Virtual Reality, Immersive Analytics, Multi-Perspective Views, Focus+Context

ABSTRACT:

Immersive geospatial visualization finds increasing application for navigation, exploration, and analysis. Many such require the display of data at different scales, often in views with three-dimensional geometry. Multi-view solutions, such as focus+context, overview+detail, and distorted projections can show different scales at the same time, and help place an area of interest within its surroundings. By inverting the principle of stereographic projection — projecting spatial features from a map onto a virtual sphere which surrounds the viewer — we present a novel technique for immersive geospatial focus+context that aims to mitigate problems with existing solutions. This sphere can intersect the map, dividing it into two parts: the inside of the sphere, which stays unchanged, and the outside, which gets projected to the surface, resulting in an inversion of the lens metaphor by distorting the context instead of the focus. This detail-in-context visualization maximizes the amount of context that can be legibly shown by the smooth compression inherent to the stereographic projection, and by utilizing otherwise unused screen space in the sky. The projection method allows for easy control over the projection and distortion characteristics by varying only two main parameters — the sphere's radius and its position. The omnidirectional nature of our system makes it particularly well-suited for immersive displays by accommodating typical immersive exploration and fully utilizing the additional visual space available. Applying our system to an urban environment, we were able to solicit positive reactions during feedback sessions with experts from urbanism.

1. INTRODUCTION

Many advances in visualization or visual analytics in geovisualization and urbanism are still making use of a flat map representation ((Masse and Christophe, 2015), (Karduni et al., 2017), and most examples listed of the literature review in (Zheng et al., 2016)), building on the traditional methods of cartography. However, with the increasing availability and accessibility of 3D technology and data, the use of 3D geovirtual environments (Pasewaldt, 2013) is finding growing advocacy in multiple geospatial domains. Be it to supplement 2D map views with better qualitative understanding (Brooks and Whalley, 2008), to reveal heterogeneous and multidimensional data that would be impossible to show otherwise (Graciano et al., 2017), or to utilize its attractive nature to support public participation (Brasebin et al., 2016), research is ongoing on how to improve the utility and usefulness of digital 3D maps.

The three-dimensional geometry of cities is not only the basis for measuring quantitative aspects (Ferreira et al., 2015) such as sky exposure, but vital for qualitative indicators (Ortner et al., 2016) of e. g., new developments and their visibility in the urban environment. In many aspects of urban studies and design, a representation of the urban morphology that is as close as possible to what can be seen in reality is desirable — providing the urbanist or stakeholder with a *sense of place* (Salerno, 2017).

Visibility itself does however become a problem with three-dimensional views. While geovirtual models naturally appear more realistic in 3D, and even more so with the stereoscopic vision immersive environments augmented and virtual reality (AR/VR) offer (Polys et al., 2018), they become sources of occlusion for objects lying behind them. The choice of perspective from which a scene is viewed thus lets the user set the trade-off between the clear visibility of a flat 2D map presentation and the realistic, but limited 3D view at eye level. Moving from one to the other point of view leads to an effect that can be seen as separating the scenes into two areas: the *focus* area, which remains legible throughout; and the *context* area, rendered illegible by occlusion and distortion.

The traditional definition of *focus* and *context* in visualization assumes simultaneous, but not necessarily equivalent, visibility of both areas, and various techniques have been developed to enable this (Cockburn et al., 2009). In 2D cartography, where the above drawbacks of 3D do not exist, focus+context is used to enhance the visibility of the focus area, usually by enlarging it at the expense of compressing part of the context, like a lens lying on top of the map (Pietriga and Appert, 2008).

In keeping with the spirit of Kevin Lynch's *Image of the City* (Lynch, 1960) by assuming the need to visualize an area of interest or focus within a city realistically, and preserving its connection to its surrounding context as well as internal continuity, we present a novel method for focus+context based on an inversion of the lens metaphor (Tominski et al., 2017) — the terrain outside the focus area is bent upwards for better visibility. It utilizes inverse stereographic projection of terrain to a sphere intersecting it, adapting the principle to allow for terrain with uneven elevations.

* Corresponding author

2. RELATED WORK

Focus+Context in the domain of geovisualization, also sometimes called *Detail-In-Context* (Keahey, 1998) was first applied and is still often used for 2D map views, usually in the form of lenses. Pietriga and Appert (Pietriga and Appert, 2008) explain the common types of such lenses, which are based on physical metaphors like magnifying glasses. The desired in-place magnification can be achieved by superimposing a magnified view of the focus over its origin in the overview, which leads to occlusion and therefore loss of the area just outside the focus. To avoid this discontinuity and loss of information, lenses can instead have a transition area between focus and context, in which a distortion occurs. Zhao et. al. (Zhao et al., 2012) presented a method for smooth magnification that uses conformal deformation, i. e., it minimizes distortion of smaller shapes by exactly preserving local angles, thus increasing legibility of features and the relationship between focus and context areas.

In perspective rendering of 3D terrain models, occlusion can be caused by the landscape or built elements themselves when viewed at an oblique angle. Focus+context then, in the sense of increasing legibility of an area of interest within its context, does not necessarily need magnification as much as “disocclusion.” Wu and Popescu (Wu and Popescu, 2017) demonstrate an approach to effect this by rendering multiple perspectives inside one view: the original, primary perspective, where an object or region of interest (ROI) is hidden by occluding geometry before it, and the secondary perspective, chosen so that the ROI is visible from it. Blending these multiple views inside one results in a distortion of the area surrounding the ROI, which itself, along with the context outside this transition area does not appear distorted.

A simpler approach to multi-perspective views in virtual urban environments is presented by Lorenz et. al. (Lorenz et al., 2008) (used for navigation and wayfinding in (Möser et al., 2008), expanded on in (Pasewaldt et al., 2014) to apply to focus+context), in which one global deformation is applied to the entire geovirtual model, bending it along a parametric curve. The focus is the not-deformed area close to the viewer, which stays flat on the original plane. The transition area follows, in which the next part of the model is smoothly bent upward, following the deformation curve, to its target angle, at which the context is displayed as if viewed from an almost top-down perspective. Veas et. al. (Veas et al., 2012) used a very similar principle by utilizing a primary (first-person view) camera, and blending its view with that from a secondary camera that is rotated upwards about an axis at the distance of the transition between focus and context areas. In their application this view is rendered as transparent wireframes and laid over the image of a physical camera, for an Augmented Reality (AR) view. Veas et. al. term this as “extending overview,” which is also what can be said about the multiperspective view approach by Pasewald et. al. — both overcome the limitations a “long flat view” imposes through its inherent occlusion and perspective deformation by augmenting it with a view from above. Most recently, Chen et. al. (Chen et al., 2021) used geometric projection to achieve similar results, i. e., a terrain that bends upwards toward the user in one direction.

A limitation of these three systems lies in their directionality: the Multi-Perspective Views are deformed along one curve extending from the viewer, and the Extended Overview technique bends the terrain about one axis that is perpendicular to the viewing direction. The parameters for both systems can be optimized for the view directly ahead, but with increasing distance from the middle of the rendered images the projection deviates from that ideal. With a wide field of view (such as that in modern VR), the extreme left and right of the image would not benefit at all from the bending of the distant context. Mitigating this by moving the curve or rotation axis with every lateral rotation a viewer makes in an interactive setting would significantly alter how the environment is being deformed with every move, effectively rotating the image projected to the sky with the rotation of the viewer, pinching and stretching its sides.

One method that solves this problem of directionality is using a panoramic perspective view, which presents all directions around the camera simultaneously and equally by projecting on a cylinder instead of a rectangle as the viewing plane of the camera. Böttger et. al. (Böttger et al., 2008) presented a transformation similar in result to the panorama in that regard, but which also provides the benefits for focus+context that come with a more gradual compression of distant features than with perspective projection, which compresses infinitely toward the horizon. The result is an enlarged focus area (the origin on the original map) on the bottom of the picture, and a compressed context at the top end. If applied to a global map of the Earth, this method provides the whole global context around a focal point. This view suffers from the same drawback as the cylindrical panoramic projection, in that it bends straight lines that are not intersecting the origin. This effect can be mitigated by stretching the image horizontally, but that results in a loss of its conformal nature, rendering the context less recognizable.

3. DESIGN CONSIDERATIONS

We intend to develop a technique that strikes a balance between preserving legibility of the urban form and reshaping its visualization to allow for focus+context. On a higher level, this means presenting professionals in urbanism, as well as novice users with a legible and recognizable view of the city, which at the same time offers a new perspective on the relationship between a region in focus and its surrounding context, aiding in the visual analytics process (Zheng et al., 2016), exploration, and navigation in AR, or communication. On a lower level, this encompasses the following specific requirements we defined:

- **Present 3D city models in familiar perspectives:** In many urban applications, particularly in those that include participatory elements, a common perspective is one that shows a region of the size of a city block from above the rooftops, giving visibility to the shape of individual buildings without too much occlusion (Brasebin et al., 2016), but lower perspectives down to eye-level are also utilized (Polys et al., 2018). These should act as our focus views, without any distortion that could negate their benefits.

- **Reveal context hidden by occlusions or perspective:** The resulting focus+context view should not introduce discontinuities, and apply a deformation that minimizes the impact on legibility, so as to maintain the “image of the city,” and actually enhance the visibility of connections (physical or conceptual) between different regions.

- **Provide simple controls over focus+context:** The relative size of the focus area within the context and the amount of context made visible should be adjustable in real time, according to the needs of the user in exploring different views or finding an optimal balance.

This contribution has been peer-reviewed. The double-blind peer-review was conducted on the basis of the full paper.
https://doi.org/10.5194/isprs-annals-V-4-2022-235-2022 | © Author(s) 2022. CC BY 4.0 License.

- **Benefit from immersive technologies:** Immersive technologies are increasingly being considered for broader use, including in urban analytics, thanks to their rapid improvement (Chen et al., 2017). By their nature, they provide an omnidirectional canvas on which to project data, as opposed to the limited windows traditional monitors present — a view of every direction is available through natural head movement. If this space is to be utilized, our method needs to work equally well in any direction the user looks, without having to change its characteristics to accommodate rotations.

These requirements eliminate approaches such as MCVs or similar overview+detail techniques, which would introduce discontinuities. The lens metaphor only works in its inverse application, as we want the user to control how much and how the context is added, while the focus remains unchanged from its original form. Necessarily, an inverse lens approach without discontinuities results in some form of deformation. If an important part of maintaining legibility consists of the preservation of local shapes (Zhao et al., 2012), conformal deformation becomes necessary.

4. INVERSE STEREOGRAPHIC PROJECTION

4.1 Projecting 3D Maps to a Virtual Sphere

The *stereographic projection*, in its application for "unfolding of a spherical surface" was first described by Hipparchus, who used it to map stars from the "celestial sphere" (Neugebauer, 1949). It is a *geometric azimuthal* projection (Deetz and Adams, 1934) and a *conformal* projection, i. e., the shapes of "small" features appear correctly, or, more precisely, that relative angles at each point are kept (Snyder, 1987). Larger and more irregular structures will appear deformed, though, as the projections of features on the plane are more compressed with increasing distance from the sphere.

This has been applied to create projections of photographic panoramas (Germán et al., 2007), and in immersive visual analytics to project graph visualizations (Kwon et al., 2016). As a cartographic tool, it remains particularly useful in meteorological applications, where it is used to track moving weather features such as cyclones (Shenk et al., 1971). These examples show that the characteristics of this projection were not only found useful in many applications, but that this widespread use keeps familiarizing users with it, particularly within geospatial domains of expertise.

In our application, we want to use the properties of this projection to fold the context around the viewer. We thus create a virtual sphere surrounding the viewer, and project onto it geospatial data from their surroundings after a certain distance.

This distance depends on the scale of the objects (building, district) needed in the focus area that will remain detailed, and on the legibility of distant features of the context area. We call this the *focus view extent*, as it defines where the focus ends and the context begins. This approach implies that the focus area or region-of-interest (ROI) is a circle centered on the viewer, and thus having a geometry-based focus representation, as defined by (Trapp, 2013). This ROI immediately around the viewer is by definition close enough to remain legible and needs to present 3D geospatial data, which could contain uneven terrain if topographic features are relevant to the use case.

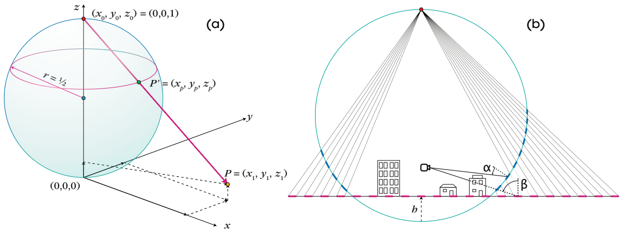

Figure 1. Inverse stereographic projection of a point P in space to P' on a sphere (a), and the compression and improvement in viewing angle from the projection in an urban environment (b).

Our application as outlined above is shown in fig. 1 *(a)*. Three main points differ from the classic stereographic projection:

1. the projection is *from* a secant plane to the sphere,
2. the projection only applies to points on the plane *outside* the sphere, and
3. the "plane" is actually a *terrain* with *different elevations* at each point.

Point 3 presents a challenge for the classic stereographic projection, which assumes a flat plane at a pre-defined distance from the projection origin on a pole of the sphere. Fig. 1 (a) shows the basis for our adapted form of projection, mapping any point in space to the sphere, not just those on one plane. Using its coordinate system and the points shown therein, the resulting formula for the projected point P' on the unit sphere centered around $(0, 0, 0.5)$ from P anywhere outside it is:

$$P' = \frac{1}{x_1^2 + y_1^2 + (z_1 - 1)^2} \cdot \begin{bmatrix} (x_1) \cdot 1 - z_1 \\ (y_1) \cdot 1 - z_1 \\ x_1^2 + y_1^2 \end{bmatrix} \tag{1}$$

4.2 Distortions from Projecting Terrain

Since we are not necessarily using flat planes for projection anymore, the conformity gets disturbed wherever there are differences in terrain elevation. Changes in elevation slopes result in changes to the angles of the lines connecting points on those slopes — they are no longer coplanar with the connecting lines of other points, meaning they do not share the same distortion characteristics as in the planar case and therefore do not preserve angles as lines on coplanar shapes do when projected.

The projection still is "conformal" in the strictest sense only for infinitesimal distances, but not anymore when comparing the projected image on the sphere to a flat map of the same region. A more direct comparison of cartographic features projected to distant parts of the sphere can be made if the terrain is first flattened before projection, eliminating the additional distortions elevation differences cause. Fig. 2 illustrates this approach, and its result can be seen in fig. 3.

There are clear distortions to the city's grid of streets when they lie on sloped terrain — features on an uphill slope from the sphere's center are stretched, while those downhill get compressed or even inverted. They are particularly noticeable when

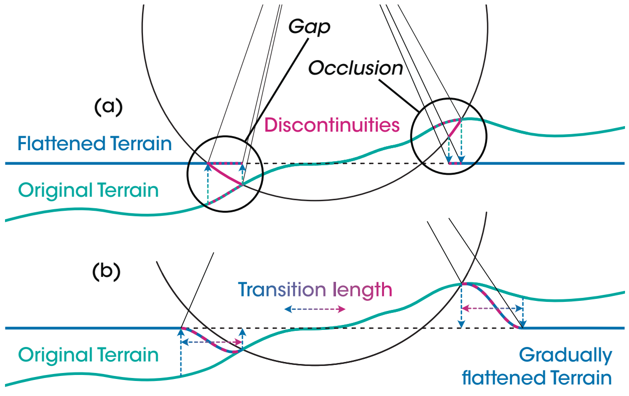

Figure 2. Inverse stereographic projection of elevation outside a sphere to its surface leading to discontinuities (a), and gradually flattening the terrain before projection to avoid them (b).

the terrain is being moved relative to the sphere, as this effect is dependent on the origin of projection relative to the terrain features and can change drastically by doing so. These distortions, while sacrificing some conformity, do however result in a form of the "kinetic depth effect," or *structure from motion* (Vezzani et al., 2014) during movement, which enables the perception of depth information from geometry projected on flat surfaces.

4.3 Flattening Terrain before Projecting to the Sphere

If the distortions in the projection resulting from terrain are unwelcome — if the distortions are too intense, or the projection needs to be as similar to a regular 2D map as possible — but an elevation model is still required in the focus area inside the sphere, the terrain needs to be flattened on a secant plane before projection. We suggest taking the plane that passes through the point on the terrain which is intersected by the vertical axis of the sphere, as it is the reference position of the user.

However, if the terrain is simply flattened to that secant plane, gaps or occlusions can appear at the interface of sphere and terrain wherever it is not originally flat, as shown in fig. 2 (a). This compromises the visual continuity that is essential in e. g., perceiving paths in a city (Lynch, 1960). It is therefore necessary to introduce a transition zone between original and flattened terrain to counter this, in which the elevation is gradually morphed to that of the secant plane (shown in fig. 2 (b)). This can be understood as deforming data to avoid the much more disruptive discontinuity, similar to how smooth focus+context lenses work (Pietriga and Appert, 2008). Gaps and occlusions can thus be avoided with a sufficiently large transition zone. It can also serve as a more gradual way to ease into a flattened context view from the focus area, depending on the function used for flattening. Fig. 3 shows this implemented as in fig. 2 (b).

5. VISUALIZATION TECHNIQUE

Our main goal is to improve the legibility of the urban environment at multiple scales, especially in focus+context situations. This relies on an implementation of a multi-perspective visualization that is easy to understand and simple to use. With the following parameters, we aim to enable users to adapt the visualization to diverse requirements and intuitively interact with the projection method in real time, as well as to quickly navigate any map. We developed our prototype with the possibilities of VR input devices in mind, however, our interaction schemes work with any pointing device.

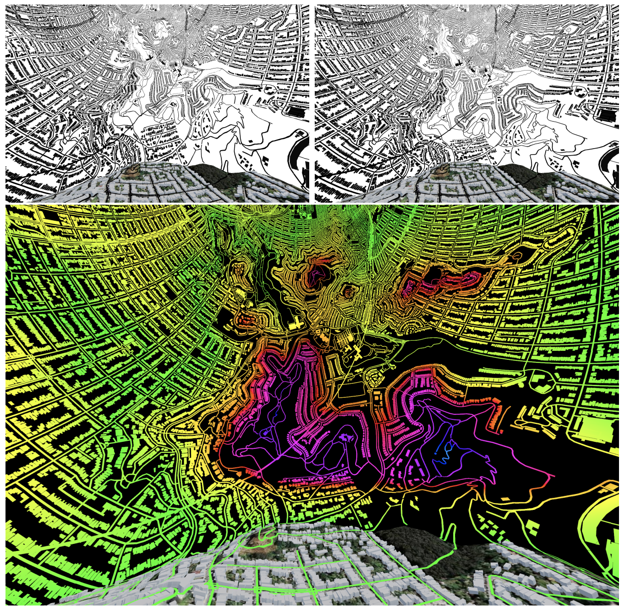

Figure 3. Projecting a terrain (*upper left*), its flattened version (*upper right*), and the differences between the two versions (lower); green for no change, purple for strong distortion.

5.1 Focus vs Context: What to Project

The virtual projection sphere divides all spatial data into two parts by intersecting it: the data inside (the focus view), and the data outside of it (the context view). Inside it, most data can remain as it is — after all, the focus view should already be clearly legible. Outside the sphere we could let some data remain completely unchanged, or alter its appearance, e. g., by dimming it to reduce distraction from the projected context view. In fig. 4 we show two such possibilities: in the image on the left the buildings are projected as transparent shapes, while the road network (with color-coding of elevation) shines through. The original scene is not altered, instead the projection just gradually emanates from it. The image on the right demonstrates a diminished view of the original context, with a simplified, footprint-only view of the buildings projected on the sphere, along with roads and rivers. This can be useful to provide the viewer with a hint at the context area in a more familiar manner. Otherwise, it could be completely culled and not projected, which could make sense for data that is only relevant for the focus view and may distract from the projected data of interest to the viewer. An example of this is shown in fig. 3. Context data that *is* projected to the sphere could be altered in appearance from how it is in the focus view, to e. g., make it more legible as done in (Pasewaldt et al., 2011), or to just signify and highlight the boundary between original and projected view, potentially to avoid misunderstandings about the introduced curvature.

5.2 Projection Parameters: How to Project

Using a sphere as our surface for stereographic projection simplifies the main projection parameters down to the radius of the sphere, and the position of its center relative to the terrain. We set the origin of the projection to always be at the zenith of the sphere. The only other remaining factor is the position of the camera relative to sphere, and this we also constrain to always be fixed to the vertical axis of the sphere — the user can look around freely and move the camera up and down without changing the sphere's location, but horizontal movement translates

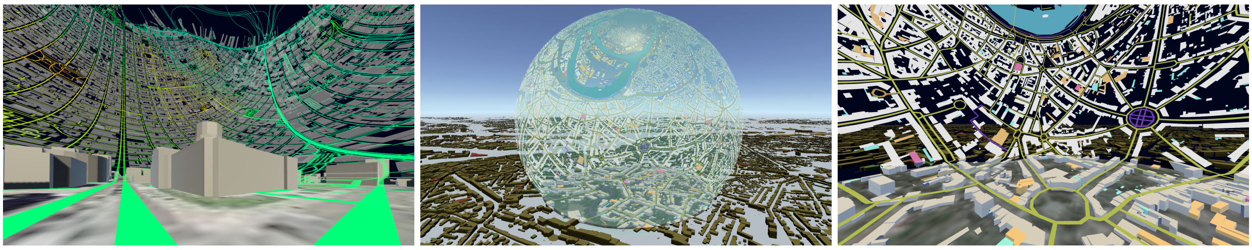

Figure 4. Different configurations of the data sphere (middle) for urban focus+context multiviews (left and right)

the sphere accordingly, ensuring a consistent projection in all directions and at all positions.

5.2.1 Intersection Ratio To account for changes in terrain elevation during horizontal navigation and to prevent them from changing the projection characteristics, we define an *intersection ratio* h (fig. 1 *(b)*), at which a theoretically flat terrain intersects the sphere. Its range is $]0, 0.5]$ — from being almost tangential to the terrain, where the focus area is but a point, to being bisected at the equator by the terrain and thus acting as a dome over the focus area. Above 0.5 the *intersection ratio* causes increased perceived compression of the projected features, as the projection surface would now be tilted toward the viewer. The actual vertical position of the sphere is calculated from its diameter and this intersection height.

During navigation, the terrain is vertically adjusted so that its point through which the sphere's vertical axis goes stays at a preset height, thus keeping the sphere's position stable as well. This leaves the users with only two parameters to adjust its projection characteristics: its diameter and intersection ratio. Permitting changes to the diameter while leaving the sphere in place would change its intersection ratio, so we implemented a system that automatically adjusts the sphere's position to preserve its current ratio. Adjusting the intersection ratio is just a scaled adjustment of its position, so the diameter is already preserved in this interaction.

5.2.2 Focus View Extent Exposing radius and intersection height and automatically adjusting the underlying or complementing parameters allows for full control over the sphere and its projection characteristics, however this mode of usage remains less than fully intuitive. Changing either parameter results in a change of projection characteristic (compression) *and* the distance at which the sphere intersects the secant plane, as measured horizontally from the point on the plane below the camera, at the same time.

For a more intuitive interaction scheme we therefore introduce the measure of *focus view extent* as another abstraction that can be directly manipulated by the user or automatically kept fixed when the diameter is being adjusted. In this mode, interacting with the diameter changes how much context is seen in the field of view, with varying compression and angle of intersection, as the sphere's position is automatically adjusted to keep the focus view extent — the intersection area — fixed. The inverse of keeping the diameter fixed while changing the focus view extent results in relatively little change in compression.

We expose controls over these projection parameters to the user via a GUI with sliders and checkboxes in the desktop implementation. For the VR version, we map the two arbitrary parameters (intersection and diameter, or diameter and focus view extent) to the vertical and horizontal axis of one of the controllers' touchpad. We scale the input for diameter and focus view extent geometrically for a more natural-feeling and fast interaction, allowing a user to quickly and precisely adjust the projection parameters to their needs at any time.

5.3 Navigation

We designed the methods for traversing the 3D environment with a focus on VR. This meant taking advantage of the availability of controllers that are tracked in 3D space along with their rotation, as well as considering the well-known effects of simulator sickness that inappropriate methods of travel can cause (Kolasinski, 1995). Navigation in 3D environments is accomplished either gradually, or instantaneously — through steering or teleportation, respectively. The merits of these systems and their variations are still being thoroughly studied (Clifton and Palmisano, 2019), but generally speaking steering is a more natural method that maintains a high degree of spatial awareness, but very often causes motion sickness, while teleportation generally avoids simulator sickness, but disrupts *spatial updating* — the "'process that automatically keeps track of where relevant surrounding objects are while we locomote, without much cognitive effort or mental load" (Riecke, 2003).

Steering is usually accomplished in an indirect manner, by e. g., tilting a joystick or pressing buttons to change directions and speed or acceleration, however more direct methods have been proposed, such as an approximation of the functionality of a hamster ball (Hurtado et al., 2018). In teleportation, the user usually needs to find a target, and point at it to be immediately transported to (or close to) its position. This can pose a challenge when longer distances and occlusions are involved, similar to the challenge of focus+context visualization. The data sphere can by its nature serve as an aid in both methods, as it envelops the viewer with legible targets to point to or otherwise interact with.

Including the case of travel by walking across the room with tracking of the VR headset's position (the sphere's axis remains attached to the camera's position as discussed above), we implemented a number of methods to move and scale the terrain relative to the viewer and therefore the sphere, with all but one of them being based on ray casting. The user can press a button on one of the controllers, casting a visible ray, or pointer, from its tip, which can intersect either the unchanged terrain inside the sphere or its image on the projection surface.

Building on that, we have implemented direct and indirect "grabbing" methods to translate the environment in a natural way with a measure of inertia, zooming and rotation, steering, instant teleportation, and a "fly-over" method, in which the target position is being gradually eased in. Utilizing the buttons, touchpads and triggers of two HTC Vive controllers, we were able to present all these navigation methods, as well as the interaction with the projection parameters at the same time to a user, allowing them to explore and compare each method quickly.

6. IMPLEMENTATION AND PRELIMINARY EVALUATION

3D video game engines are finding increasing use for constructing visualizations (Buyuksalih et al., 2017) and simulations (Cristie et al., 2015) in urbanism as well as other fields, and are particularly well suited for developing applications for immersive analytics (Sicat et al., 2018, Cordeil et al., 2019). For our development we used the Unity 3D engine, and its SDK for Mapbox, a platform providing cartographic data and tools to author, edit and stylize geospatial datasets. Furthermore, Unity 3D supports immersive technology such as the HTC Vive VR HMD used to test our prototype. Using the Mapbox SDK, we access geographical vector and raster data stored as *Tilesets* from Mapbox and use these to construct and deform textured or colored geometry in Unity at runtime.

During our iterative development of the data sphere system, we asked professionals and researchers in urbanism for their feedback and informal evaluation of most elements we implemented. Their input helped inform our choices of the parameters for the navigation methods, and how the projection parameters should be constrained and controlled, as well as aesthetic choices for colors and transparencies.

As the implementation neared a more complete state, we invited seven participants from the same group to perform a series of exploration tasks with the system, to gauge their reaction to the visualization and the interaction methods. We presented them with a 3D terrain of the city they live in, textured with aerial imagery, which visualized all 3D buildings and the transportation networks in a stylized way, similar to the right image in fig. 4. We explained the use of the controllers, and instructed them to find and navigate to certain locations.

The procedure was as follows: a gradual introduction to the navigation commands, tested by the task of locating and navigating towards their own house in a zoomed-out way — the city was fully inside the focus area, not intersecting the sphere and thus just a basic 3D map at first. We then asked them to locate the building they are currently in, and zoom into it, thereby extending the city far beyond the focus and into the context, projecting it on the sphere. They were then asked to locate and navigate to a landmark in close proximity, using the navigation commands on the sphere. After that, we introduced the commands for changing the projection parameters, enabling them to see the context in ways that allowed them to more easily find a third, much more distant landmark.

All participants quickly mastered the interaction methods, and we could see preferences for some navigation controls over others: the grab methods, especially the pointer-based one were strongly favored over the steering methods, which sometimes prompted signs of cybersickness. Somehow surprisingly, the fly-over method was deemed acceptable once we calibrated the acceleration curves to the demands of the participants. Shorter travel within the focus view was preferably executed with the indirect grabbing method, as it does not require attention to where the controller is pointing towards, whereas longer travel benefited from the context view and the ability to directly point toward the desired destination on the projection surface. Regarding our novel focus+context view, most participants — even those with extensive experience with 3D city models — were positively surprised at seeing the city in this way, with one professional in urbanism expressing that it "completely changes [her] way of understanding urban spaces."

7. DISCUSSION

Utilizing stereographic projection of the terrain to a sphere intersecting it (fig. 4), we preserve the (elevated) first-person perspective for the focus view inside the sphere, and project the top-down perspective — as seen from the zenith of the sphere — to its surface, resulting in a form of multi-perspective visualization by blending these two views into one.

7.1 Projection Characteristics

The characteristics of this projection allow for a transition between these perspectives, the smoothness of which depends solely on the radius of the sphere and at which height it intersects the terrain. Furthermore, they enable us to capitalize on the advantage of immersive environments offer with their large virtual display size by utilizing the normally unused region above the horizon of the perspective view.

Being axially symmetric, it works and looks the same in all horizontal directions, as opposed to solutions which deform the terrain in only one direction (Pasewaldt et al., 2014, Chen et al., 2021). Seamless exploration of the whole environment is thus possible, instead of only having one static perspective, or creating an entirely new projection with every rotation that is inconsistent between viewing angles.

The characteristics of the deformation and the virtual environment allowed us to introduce a number of navigation methods which allow a user to traverse large distances in the terrain intuitively and precisely. By making them all available to the user at the same time utilizing only the VR controllers, we aimed to enable users to quickly try out and compare each method, to find one that they are personally most comfortable with. The teleportation methods allowed very quick travel to specific points, while the steering and grabbing methods were helpful for smaller-scale exploration, and in particular enhancing the structure-from-motion effect that the projection of elevated terrain enabled. The option to flatten this elevation before projection gives the users the choice to forgo this effect in order to get a less distorted image of the city that looks more familiar to regular map projections.

7.2 Potential Applications and Use Cases

The immersive nature of our implementation in virtual reality made the data sphere system more engaging than a representation on a desktop screen, both for the more immediate visibility of the projection surface actually surrounding the viewer (making its nature easier to grasp) and also because at this point immersive systems are still mostly a novelty. Most people having only experienced VR in games, if at all, we could observe a higher proclivity for playful exploration of the urban scenes we presented — our participants expressed strong engagement with the projection and the navigation methods. This suggest that a system like this could find application in communicating new proposals for urban developments — an area that already pays a lot of attention to engaging stakeholders as much as possible.

Other possible areas of application are in cases of urban planning where work is focused on one (small) region/district in a 3D view, and connections (transports, vegetation, river, global city map, etc.) to other regions need to be seen, or other, distant regions need comparison among themselves and to the region in focus. In planning tasks that are more local, our focus+context method can be used to visualize e. g., legal requirements for 3D zoning in detail for the city block in focus, while keeping an overview of the simplified 2D land use in its context.

For use as an aid in geographic exploration, the virtual data sphere can offer a different way to judge distances of and the connections to points of interest in the context of an area in focus. In flat terrain projections, distances from the center of the focus are exactly mapped to their vertical position on the sphere and can thus be intuitively compared without the aid of rulers or grids. This could be of particular interest if isochrones are added to the map (O'Sullivan et al., 2000), making the comparison between distance and time easier.

Finally, another obvious application for the inverse stereographic projection would be a navigation aid in augmented reality, showing a user the highlighted path among the whole street network. It can visually connect to the current street they are on, and smoothly compress the entire path on the map, while also indicating its orientation and hinting at distances.

During the preliminary evaluation we noticed how the projected image was sometimes the only part used by some participants. Instating some limitations that guide users to view and use the sphere for context and the terrain inside it as focus made our participants understand and appreciate the paradigm much quicker. While some of the sentiment of "changing one's view of the city" that participants expressed could be attributed just seeing it in immersive virtual reality, we think that the enveloping nature of the sphere heightens this feeling of being immersed in the data in a literal way, though this is an avenue to be investigated in future work.

8. CONCLUSION AND FUTURE WORK

Through our implementation of the inverse stereographic projection of terrain to a virtual sphere surrounding the viewer, we created a technique that satisfies our design considerations for a focus+context tool for use in geovisualization, particularly of urban environments, and improved on the idea of continuous, multi-perspective deformations. We synthesized our new approach to geospatial focus+context exploration by adapting and amalgamating long-established, as well as cutting-edge methods from cartographic projection, multi-perspective view deformation, and immersive analytics. Implementing it with the goal of being intuitively usable and interactive, we were able to elicit positive feedback and find first acceptance from potential users, and showing promise for diverse potential use cases.

Due to the novelty of this concept and the broad possibilities of using and interacting with it, the inverse stereographic projection on a virtual sphere opens up many avenues for future work. An obvious one would be to extend the same principle to different radially axial projection surfaces, like the cylinder, cone, ellipsoids, etc. The projection could be further refined by e. g., using interpolated multiscale maps (Dumont et al., 2016) at different distances, and adapted to work in AR settings.

Investigation is called for into its actual use not just for urban data, but also for larger scales, such as geological or climatological visualizations, especially in the shape of formal evaluations and comparisons. Three such avenues are easily identified: comparing different projection parameters, such as the projection shapes itself (sphere vs. cone vs. cylinder etc.) and their settings; comparing this projection system against other multi-perspective implementations that ostensibly share similar goals (Pasewaldt et al., 2014, Chen et al., 2021); and comparing against more distant paradigms which serve similar purposes, such as overview+detail systems.

REFERENCES

Böttger, J., Preiser, M., Balzer, M., Deussen, O., 2008. Detail-in-context visualization for satellite imagery. *Computer Graphics Forum*, 27number 2, Wiley Online Library, 587–596.

Brasebin, M., Christophe, S., Jacquinod, F., Vinesse, A., Mahon, H., 2016. 3D Geovisualization & stylization to manage comprehensive and participative Local Urban Plans. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 4, 83.

Brooks, S., Whalley, J. L., 2008. Multilayer hybrid visualizations to support 3D GIS. *Computers, Environment and Urban Systems*, 32(4), 278–292.

Buyuksalih, I., Bayburt, S., Buyuksalih, G., Baskaraca, A., Karim, H., Rahman, A. A., 2017. 3D Modelling and Visualization Based on the Unity Game Engine–Advantages and Challenges. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 4, 161.

Chen, S., Miranda, F., Ferreira, N., Lage, M., Doraiswamy, H., Brenner, C., Defanti, C., Koutsoubis, M., Wilson, L., Perlin, K. et al., 2021. UrbanRama: Navigating Cities in Virtual Reality. *IEEE Transactions on Visualization and Computer Graphics*.

Chen, Z., Wang, Y., Sun, T., Gao, X., Chen, W., Pan, Z., Qu, H., Wu, Y., 2017. Exploring the design space of immersive urban analytics. *Visual Informatics*, 1(2), 132–142.

Clifton, J., Palmisano, S., 2019. Effects of steering locomotion and teleporting on cybersickness and presence in HMD-based virtual reality. *Virtual Reality*, 1–16.

Cockburn, A., Karlson, A., Bederson, B. B., 2009. A review of overview+ detail, zooming, and focus+ context interfaces. *ACM Computing Surveys (CSUR)*, 41(1), 1–31.

Cordeil, M., Cunningham, A., Bach, B., Hurter, C., Thomas, B. H., Marriott, K., Dwyer, T., 2019. Iatk: An immersive analytics toolkit. *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, IEEE, 200–209.

Cristie, V., Berger, M., Bus, P., Kumar, A., Klein, B., 2015. Cityheat: Visualizing cellular automata-based traffic heat in unity3d. *SIGGRAPH Asia 2015 Visualization in High Performance Computing*, SA '15, Association for Computing Machinery, New York, NY, USA.

Deetz, C. H., Adams, O. S., 1934. *Elements of map projection*. 4 edn, US Government Printing Office.

Dumont, M., Touya, G., Duchene, C., 2016. A comparative study of existing multi-scale maps: what content at which scale? *International Conference on GIScience Short Paper Proceedings*, 1number 1.

Ferreira, N., Lage, M., Doraiswamy, H., Vo, H., Wilson, L., Werner, H., Park, M., Silva, C., 2015. Urbane: A 3d framework to support data driven decision making in urban development. *2015 IEEE conference on visual analytics science and technology (VAST)*, IEEE, 97–104.

Germán, D. M., Burchill, L., Duret-Lutz, A., Pérez-Duarte, S., Pérez-Duarte, E., Sommers, J., 2007. Flattening the viewable sphere. *Computational Aesthetics*, Citeseer, 23–28.

Graciano, A., Rueda, A. J., Ortega, L., Feito, F. R., 2017. Towards a hybrid framework for the visualization and analysis of 3d spatial data. *Proceedings of the 3rd ACM SIGSPATIAL Workshop on Smart Cities and Urban Analytics*, 1–8.

Hurtado, J., Albuquerque, E., Radetic, D., Cherullo, R., Silva-Calpa, G. F. M., Raposo, A., 2018. Floating hamster ball: A locomotion method for free flight in virtual environments. *2018 20th Symposium on Virtual and Augmented Reality (SVR)*, IEEE, 183–191.

Karduni, A., Cho, I., Wessel, G., Ribarsky, W., Sauda, E., Dou, W., 2017. Urban space explorer: a visual analytics system for urban planning. *IEEE computer graphics and applications*, 37(5), 50–60.

Keahey, T. A., 1998. The generalized detail in-context problem. *Proceedings IEEE Symposium on Information Visualization (Cat. No. 98TB100258)*, IEEE, 44–51.

Kolasinski, E. M., 1995. *Simulator sickness in virtual environments*. 1027, US Army Research Institute for the Behavioral and Social Sciences.

Kwon, O.-H., Muelder, C., Lee, K., Ma, K.-L., 2016. A study of layout, rendering, and interaction methods for immersive graph visualization. *IEEE transactions on visualization and computer graphics*, 22(7), 1802–1815.

Lorenz, H., Trapp, M., Döllner, J., Jobst, M., 2008. Interactive multi-perspective views of virtual 3d landscape and city models. *The European Information Society*, Springer, 301–321.

Lynch, K., 1960. *The image of the city*. 11, MIT press.

Masse, A., Christophe, S., 2015. Homogeneous geovisualization of coastal areas from heterogeneous spatio-temporal data. *The International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences*, 40(3), 509.

Möser, S., Degener, P., Wahl, R., Klein, R., 2008. Context aware terrain visualization for wayfinding and navigation. *Computer Graphics Forum*, 27number 7, Wiley Online Library, 1853–1860.

Neugebauer, O., 1949. The Early History of the Astrolabe. Studies in Ancient Astronomy IX. *Isis*, 40(3), 240–256.

Ortner, T., Sorger, J., Steinlechner, H., Hesina, G., Piringer, H., Gröller, E., 2016. Vis-a-ware: Integrating spatial and non-spatial visualization for visibility-aware urban planning. *IEEE transactions on visualization and computer graphics*, 23(2), 1139–1151.

O'Sullivan, D., Morrison, A., Shearer, J., 2000. Using desktop GIS for the investigation of accessibility by public transport: an isochrone approach. *International Journal of Geographical Information Science*, 14(1), 85–104.

Pasewaldt, S., 2013. Comprehensible 3D Maps and Building Panoramas. *Proceedings of the 6th Ph. D. Retreat of the HPI Research School on Service-oriented Systems Engineering*, 76, 117.

Pasewaldt, S., Semmo, A., Trapp, M., Döllner, J., 2014. Multi-perspective 3D panoramas. *International Journal of Geographical Information Science*, 28(10), 2030–2051.

Pasewaldt, S., Trapp, M., Döllner, J., 2011. Multiscale visualization of 3d geovirtual environments using view-dependent multi-perspective views.

Pietriga, E., Appert, C., 2008. Sigma lenses: focus-context transitions combining space, time and translucence. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 1343–1352.

Polys, N., Newcomb, C., Schenk, T., Skuzinski, T., Dunay, D., 2018. The value of 3d models and immersive technology in planning urban density. *Proceedings of the 23rd International ACM Conference on 3D Web Technology*, 1–4.

Riecke, B. E., 2003. *How Far Can We Get with Just Visual Information?: Path Integration and Spatial Updating Studies in Virtual Reality*. Logos Verlag.

Salerno, R., 2017. Mapping urban environment by geometry (es) and perception (s). *Urban Design and Representation*, Springer, 11–21.

Shenk, W. E., Powell, H., Salomonson, V. V., Bandeen, W. R., 1971. Meteorological uses of the stereographic horizon map projection. *Journal of Applied Meteorology*, 10(3), 582–589.

Sicat, R., Li, J., Choi, J., Cordeil, M., Jeong, W.-K., Bach, B., Pfister, H., 2018. Dxr: A toolkit for building immersive data visualizations. *IEEE transactions on visualization and computer graphics*, 25(1), 715–725.

Snyder, J. P., 1987. *Map projections–A working manual*. 1395, US Government Printing Office.

Tominski, C., Gladisch, S., Kister, U., Dachselt, R., Schumann, H., 2017. Interactive lenses for visualization: An extended survey. *Computer Graphics Forum*, 36number 6, Wiley Online Library, 173–200.

Trapp, M., 2013. Interactive Rendering Techniques for Focus + Context Visualization of 3D Geovirtual Environments. PhD thesis.

Veas, E., Grasset, R., Kruijff, E., Schmalstieg, D., 2012. Extended overview techniques for outdoor augmented reality. *IEEE transactions on visualization and computer graphics*, 18(4), 565–572.

Vezzani, S., Kramer, P., Bressan, P., 2014. *Stereokinetic Effect, Kinetic Depth Effect, and Structure from Motion*. Oxford University Press.

Wu, M.-L., Popescu, V., 2017. Efficient VR and AR navigation through multiperspective occlusion management. *IEEE transactions on visualization and computer graphics*, 24(12), 3069–3080.

Zhao, X., Zeng, W., Gu, X. D., Kaufman, A. E., Xu, W., Mueller, K., 2012. Conformal magnifier: A focus+ context technique with local shape preservation. *IEEE Transactions on Visualization and Computer Graphics*, 18(11), 1928–1941.

Zheng, Y., Wu, W., Chen, Y., Qu, H., Ni, L. M., 2016. Visual analytics in urban computing: An overview. *IEEE Transactions on Big Data*, 2(3), 276–296.