

# Bema: A Multimodal Interface for Expert Experiential Analysis of Political Assemblies at the Pnyx in Ancient Greece

Kyungyoон Kim\*  
University of Minnesota

Bret Jackson  
University of Minnesota  
Macalester College

Stephen J. Guy  
University of Minnesota

Richard Graff  
University of Minnesota

Ioannis Karamouzas  
University of Minnesota

Moses Adeagbo  
University of Minnesota

Daniel F. Keefe  
University of Minnesota

## ABSTRACT

We present *Bema*, a multimodal user interface that enables scholars of Greek rhetoric and oratory to perform virtual reality studies of ancient political assemblies at the hill of the Pnyx. Named after the flat stone speakers' platform utilized at the Pnyx, the *Bema* interface supports the high-level task of gaining a nuanced understanding of what it would feel like to give or receive a speech together with as many as 14,000 Athenian citizens and, further, how this experience must have changed as a result of at least two massive renovations captured in the archaeological record. *Bema* integrates solutions for several low-level interaction tasks, including navigating in virtual space and time, adjusting data visualization parameters, interacting with virtual characters, and analyzing spatial audio and architecture. Navigation is accomplished through a World-in-Miniature technique, re-conceived to support multi-touch input within a 4-wall Cave environment. Six degree-of-freedom head tracking and a sound level meter are used to analyze speeches delivered by users. Comparative analysis of different historical phases and assembly sizes is facilitated by the use of crowd simulation to generate realistic spatial arrangements for the assemblymen and staged animated transitions that preserve context while comparing two or more scenarios. An evaluation with our team's scholar of ancient Greek rhetoric and oratory provides support for the most important design decisions and affirms the value of this user interface for experiential analysis.

**Index Terms:** I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Virtual reality

## 1 INTRODUCTION

In recent years, research in 3D user interfaces has matured to the point that application designers now have a host of possible techniques and solutions to address the “classic 5” tasks [5] of Selection, Manipulation, Navigation, Symbolic Input, and System Control. Despite these recent advances, application designers still struggle to understand how best to integrate individual interaction techniques into complete interfaces and applications. To address this gap, one thing that is needed is additional research and case studies that focus on high-level tasks (e.g., data understanding, experiential learning and analysis) as motivated by complex real-world problems and datasets.

To this end, we present *Bema*, a multimodal 3D user interface for use in virtual environments (Figure 1). *Bema* supports the high-level task of gaining a nuanced understanding of what it would feel like to give or receive a speech together with as many as 14,000 Athenian citizens gathered at the Pnyx and, further, how this experience must have changed as a result of at least two renovations



Figure 1: The *Bema* interface for experiential analysis of mass gatherings and political assemblies, applied here to study the speech giving and receiving on the hill of the Pnyx in ancient Athens.

captured in the archaeological record. *Bema* was iteratively designed together with a scholar of ancient Greek rhetoric and oratory. The data visualized in the current application of *Bema* to the Pnyx site are based upon archaeological evidence, ancient texts, and first hand observations. The high-level data analysis tasks come directly from the research questions of our team's scholar. Specific questions include:

- How many audience members could attend meetings in this structure and what would it feel like to speak to 6,000 audience members or more? (Ancient texts tell us that 6,000 is the special number required to reach a quorum.)
- What does the structure itself suggest about how auditors were arranged and speakers were positioned?
- What demands did the structure make on the speaker to be seen, heard, and understood, and on the auditors to see, hear, and understand?
- Finally, how do the answers to these questions change over time, i.e., in regards to the three different stages in the architectural evolution of the Pnyx?

The Pnyx site is of enormous historical importance; it is widely regarded as the birthplace of democracy, but despite decades of intensive study, many fundamental features of the site remain in dispute. Why then, do the “problem[s] of the Pnyx” [12], including the research questions above, remain open and highly debated? We argue that a key reason for this is that answering these questions requires *experiential analysis* – a first person method of data analysis that goes beyond reading texts and performing statistical tests and

\*e-mail:volcano@cs.umn.edu

is particularly elusive for studies that treat cultural performances and other intangible aspects of ancient history. This is the type of analysis that can, in theory, become one of the most exciting applications of visualization in head-tracked, stereoscopic virtual reality (VR) environments, but only if appropriate user interfaces can be coupled with the visualizations to enable experts to engage deeply with and question their data.

Bema enables experiential analysis of the Pnyx by integrating solutions for several low-level tasks, including navigating in space and time, adjusting data visualization parameters, interacting with virtual characters, and analyzing spatial audio and architecture. Navigation is accomplished through a World-in-Miniature technique [27] re-conceived and updated to utilize multi-touch input within a 4-wall Cave environment. Visualization settings and system control operations are also performed via the multi-touch device, which is located on a pedestal placed inside the Cave. Six degree-of-freedom head tracking and a sound level meter are used to analyze speeches delivered by users; this provides scholars with a method to experience, visualize, and understand the extreme effort it would have required to speak at a volume that could be understood by the assembly. Crowd simulation is utilized to generate realistic spatial arrangements for the assemblymen, a key for creating a realistic experience.

The study of poorly preserved ancient sites often requires engaging with competing scholarly hypotheses and reconstructions. This factor is multiplied in the case of the Pnyx, a structure that underwent two major alterations in the course of its use. Accordingly, the interface is designed specifically to facilitate comparing two or more scenarios, where scenarios are defined by three parameters: 1. the historical phase of the Pnyx (we call these Phase I, II, and III), 2. the size of the assembly (varies from 1,000 to 14,000 citizens), and 3. the crowd simulation method utilized to place the assemblymen (our contributions include a comparison of two alternative state-of-the-art techniques). One interesting component of the interface employed to facilitate this comparison is adapting “staged animation,” as studied in the information visualization research community, to the problem of maintaining context while traveling virtually through space and time.

The primary contribution of this paper to the 3DUI community is a case study – a detailed account of how 3D user interfaces can be adapted, extended, and combined to produce a useful data analysis environment. By discussing the iterative design process and lessons learned, we demonstrate that the application to the Pnyx is complex and requires careful design. We hope that others in the research community can benefit from our account of the process and design decisions made in developing Bema.

The remainder of the paper begins with a discussion of related work in user interfaces, immersive visualization, and applications to cultural heritage. We then present some additional background on the Pnyx and describe the data we utilize and generate for the study. Then, we describe the Bema interface in detail. Finally, we report on results and user feedback.

## 2 RELATED WORK

Our work builds upon recent research in 3D user interfaces and immersive data visualization and extends the use of these techniques in the context of cultural heritage.

### 2.1 Interaction Techniques and 3D User Interfaces

Bema’s navigational interface builds upon early 3D interaction techniques, such as the world-in-miniature (WIM) [27], but reinterprets the work within the modern context of a Cave display augmented with a 2D multi-touch tablet. This is similar to the multi-touch table and display wall configuration utilized by Coffey et al. [6], with the exception that the 3D display utilized for Bema is a 4-wall Cave that more completely encloses the user. This has

implications for the appropriate role for rotating the WIM, a topic discussed later in the paper.

In Bema, the multi-touch display is placed on a pedestal inside the Cave, and in this respect, the interface is also related to a variety of systems that have utilized smaller 2D displays in conjunction with larger 3D immersive displays. Over the years, the physical configurations of interfaces in this style have varied greatly, ranging from tracked transparent hand-held props [7] to pen-based tablets [4] to hand-held multi-touch smartphones [15]. We place the display on a pedestal for several reasons. First, this allows us to use a relatively large (and heavy) high-resolution multi-touch display that we can leverage to present 2D data as part of the integrated visualizations. Second, this leaves the user’s hands free to gesture, interact with colleagues, take notes, etc. Third, we hypothesize that standing in front of this physical pedestal actually positively contributes to the level of immersion experienced when delivering a speech to a virtual audience.

The visuals and touch interface itself are similar to the *view based* navigation technique presented by Ajaj et al. [1], but Bema differs in the complexity of the models and tasks analyzed, the multiple scenarios depicted together within the 2D plan views, and the use of an immersive display that almost completely surrounds the user. Our work also relates to a long-standing line of 3D user interface research that investigates how best to combine speech, gesture, 3D tracking, and other forms of computer input to create coherent multimodal interfaces (e.g., [3, 14, 19]). We believe that our use of speech is unique in this context. Rather than using speech as a means of issuing “key word” or natural language commands to the computer, we utilize speech (volume data coupled with head-tracking data) to create an interactive visualization of speech intelligibility metrics that are evaluated in real time for each assemblyman based upon speech input from the user.

## 2.2 Applications to Cultural Heritage

The rapidly increasing utilization of interactive 3D tools in cultural heritage studies has been well documented (see, e.g., [8]); however, much of this work has been limited to digital documentation, reconstruction, and preservation of tangible objects. There has been comparatively less work on visualizing ancient performance and performance spaces. Noteworthy exceptions include the Theatron project, focused on detailed modeling of stage constructions representing several distinct periods of theatre history [16]. The three-year (2003-2006) EU cultural heritage project ERATO likewise focused on ancient drama, but also musical performance, including simulating the acoustics of the ancient theaters and behavioral modeling of large virtual audiences [23]. Finally, the Virtual Paul’s Cross project is devoted to the visual and acoustical recreation of the occasion and setting for the public sermon delivered by John Donne in Paul’s Churchyard, London, on November 5, 1622 [30]. None of these projects utilize techniques that would be considered state-of-the-art from the standpoint of the 3D user interfaces research community.

The most closely related work to ours is that of the ARCHAVE project, which involved a reconstruction of the Great Temple of Petra [29]. In addition to providing a new opportunity for scholars to visit the site in a virtual context, one of the interesting findings from this research was the benefit to researchers who had previously visited the site in person and participated in excavations. Since the Cave display provided a first person view of the reconstruction, the researchers were able to match their spatial understanding of the site in VR with the perceptions they had formed in their time at the physical site. ARCHAVE utilized pinch gloves for 3D selection, virtual flying with input from a tracked six degree-of-freedom wand for navigation, and other 3D interaction techniques to enable a new process of immersive data analysis. Like ARCHAVE, but unlike most cultural heritage applications that



(a) Pnyx I

(b) Pnyx II

(c) Pnyx III

Figure 2: 3D models showing the evolution of the Pnyx assembly area and relative size of the three phases. Phase I was constructed in the early 5th century BCE, phase II is dated to ca. 404 BCE, and phase III to the second half of the 4th century BCE.

focus on preservation and presentation in museum settings, Bema explores the utility of a novel 3D user interface for *data analysis* performed by *expert scholars*.

### 3 STUDYING AND MODELING THE PNYX

Our user interface design decisions can be best illustrated with reference to our case. To more clearly isolate the “problem[s] of the Pnyx” [12], we present a brief survey of relevant scholarship. Then, we describe the modeling process and data that we utilize for both architecture and the assembly.

#### 3.1 Scholarship and Open Questions

Reforms enacted in Athens in 508/507 BCE increased the political authority of the *demos*, the People, giving rise to the earliest, and what would prove the most successful democracy in ancient Greece. Meetings of the Assembly (*ekklesia*), held on a regular basis and open to all adult male Athenian citizens, were occasions to debate and vote on measures governing foreign relations, state finances, and the management of religious affairs and state-sponsored festivals; in all such matters, the People had final say. At some time between about 500 and the 460s when further democratic reforms assigned still greater authority to the *demos*, the Athenians constructed a space for these large meetings on the northern slope of the Hill of the Pnyx, a few hundred meters to the south of the city’s civic and commercial center (*agora*). Here, 6,000 or more citizens would meet in the open air as often as forty times a year to deliver and hear speeches, and then vote, usually by show of hands, to determine the policy that seemed best to the majority [2].

The Pnyx served as the primary site for Athens’ political assemblies until at least the late fourth century BCE. Today, the site’s symbolic significance as the “cradle of democracy” is perhaps greater than ever. Yet even after decades of study, many aspects of the Pnyx remain enigmatic. After small-scale archaeological exploration in the mid-19th and early 20th centuries, comprehensive, though not entirely complete, excavation of the site was conducted in the 1930s [18]. The chief finding was that the meeting place had undergone two major reconstructions over the course of its use, and the combination of available evidence permitted the excavators, Kourouniotis and Thompson, to propose an account of the structure’s evolution. In each of its three phases, the assembly area was unroofed and roughly semi-circular in form. The auditorium proper was an expansive open space in the first phase (Pnyx I), up to 40 m deep and 60 m wide. In its second phase (Pnyx II), constructed at the end of the fifth century BCE, the orientation of the auditorium was reversed; in this new configuration, the audience, which had formerly faced north, now faced a speaker who would have addressed them from the south. Pnyx III retained this new orientation, but the auditorium, built up with an enormous quantity of earth fill,

was greatly enlarged to 60 m deep and almost 120 m wide, a structural feat made possible by the use of a massive stone retaining wall substantial parts of which are visible on the site today.

While most scholars accept the basic elements of this account, the ambiguous nature of the archaeological and literary evidence currently supports divergent views on surprisingly large number of fundamental matters. Among matters in dispute, questions concerning the size and capacity of the auditorium in its three phases have had special urgency, especially insofar as they figure in the efforts of political historians to determine levels of attendance and citizen participation — and their change over a 200 year span — in the world’s first and perhaps the most celebrated democracy. Another fundamental, possibly related question concerns the motivation for the dramatic and obviously difficult reorientation of the auditorium, and how this change and the generally expanding size of the venue — and, evidently, of the audience it held — affected its suitability as a venue for public address and collective political decision making.

#### 3.2 Architectural Modeling

Figure 2 depicts the three models of the Pnyx corresponding to the three phases described previously. These models are the result of years of work from members of our team. They are created to scale and use a common coordinate system so that they can be superimposed on top of each other. However, one of the most interesting changes from phase to phase is that the slope of the hillside was reversed in the later phases by filling it with earth. So, if superimposed, the Phase I model would actually lie underground in Phase II and Phase III. These architectural data serve as the foundation for the visualizations and interfaces developed.

#### 3.3 Crowd Simulation

The conventional method of estimating the maximum occupancy of a bounded space of even ground divides the total surface area of this space by the area occupied by each person. This method is used in the literature to debate the possible capacity of the Pnyx at different stages [10, 25]. It provides a reasonable first approximation; however, it assumes an unrealistic packing of people into the space and does not lend itself to visualization of an audience positioned within a 3D reconstruction.

To address this, we generate assemblymen datasets using crowd simulation. The Pnyx proves to be an unusual case study for crowd simulation techniques. It is unlike more formal gatherings where attendees have assigned seats (e.g., [24]), and it is unlike typical uses of crowd simulation for planning evacuations, reenacting historical battles, or simulating pedestrian movement in urban areas [20]. Thus, in order to create a realistic distribution of assemblymen, we had to modify existing models.

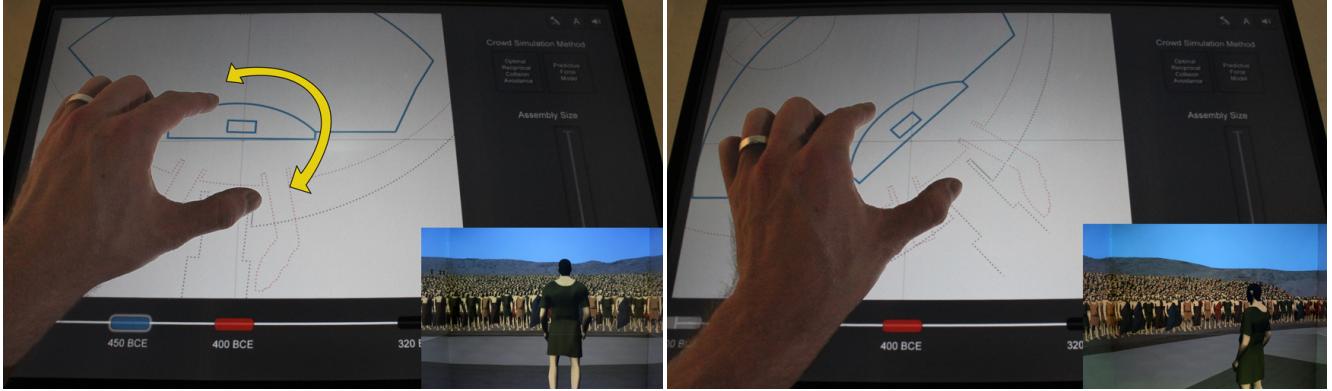


Figure 3: The user interface displayed on the multi-touch table. Left: Just before a rotation. Right: Just after using a 2-finger touch gesture to rotate the map; note that the 3D geometry displayed in the Cave rotates in correspondence with the map.

We adapted two state-of-the-art agent-based techniques: the *Optimal Reciprocal Collision Avoidance* (ORCA) navigation method proposed in [28] and the *Predictive Avoidance Model* (PAM) introduced in [13]. Both approaches account for the ability of agents to predict collisions in advance and adjust their paths to avoid these upcoming collisions. However, there are some key differences between them that influence the way in which assemblymen are distributed in the Pnyx, especially in dense scenarios. Because ORCA provides guaranteed collision avoidance between agents, agents can be viewed as ‘hard spheres’ of an exact radius and hence they can walk arbitrarily close to each other. In contrast, PAM agents rely on Newtonian forces to avoid collision. These forces scale as agents get closer to each other, resulting in a ‘soft sphere’ model. This encourages agents to keep some separation from each other unless they are being pushed together by the crowd pressure. Currently, we use these two simulation models only to compute a realistic distribution of assemblymen. In the future, we believe there is a great opportunity to extend this approach to also simulate audience behavior during a speech.

A principal parameter for our simulations is given by our assumption that an average ancient Athenian assemblyman occupies 60 cm in width (side-to-side shoulder width) and 30 cm in thickness (chest thickness). This is in reasonably good agreement with parameters utilized in previous literature. For example, Hansen [10] argues that the area occupied by each seated assemblyman was 0.4 m<sup>2</sup>, and Stanton [25] proposes an area of 0.23 m<sup>2</sup> per person.

We simulated crowds entering the Pnyx and finding a place to stand for all three historical phases (Pnyx I, II, and III) and for crowds ranging in size from 1,000 people to 14,000 people. In all of our simulations, each agent was given the goal of approaching the raised platform where the speaker stood. In Pnyx I, the agents entered the auditorium from two areas located at the left and right of the speaker’s platform. In Pnyx II, agents entered from the two staircases of the auditorium, and in Pnyx III agents entered from the large central staircase. During the simulation, each agent moved towards its goal while avoiding collisions with the other agents and the walls of the Pnyx.

Because the crowd density increased as the agents gathered around the speaker-area, only a few of them could actually reach their goals. To adapt our simulations to this situation and restrain agents from continually trying to penetrate dense areas, we devised the following approach. When an agent encountered an area of high density, it attempted to stay in place, rather than advance toward the speaker. This allowed the agents to be close to the speaker when the crowd density was low, and spread across the auditorium when the crowd density was high.

This entire simulation process resulted in a database of different

possible assemblies, parameterized by the historical phase of the Pnyx (I, II, III), the number of assemblymen (1,000 - 14,000 in increments of 1,000), and the crowd simulation algorithm utilized ([28] or [13]).

## 4 BEMA: A 3D USER INTERFACE FOR EXPERIENTIAL DATA ANALYSIS

This section describes the Bema interface in detail.

### 4.1 First Person Immersion

Bema is designed for use in a 4-wall Cave environment as shown in Figure 1. Scenes are rendered interactively using head-tracked perspective and quad-buffered stereo using a combination of custom virtual reality software and the OpenSceneGraph engine [21]. Although photorealistic rendering has not been our first priority, we do utilize several techniques to enable real-time rendering and realism. The assemblymen are rendered using an adaptive level-of-detail approach, which reduces the polygon count for models that are far away from the user. To form the assembly, a model for each assemblyman is first selected at random from a small set of key character models. Then, before placing the model at the assemblyman’s assigned position, a small perturbation is applied to the model to create some additional variation in the crowd. As shown briefly in the accompanying video, we have also experimented with applying randomized listening behavior animations to the characters. A multi-touch surface is mounted on a pedestal inside the Cave so that it is held at a comfortable height above the ground. The surface is also a display, so it provides both multi-touch input and a fifth display. In our current implementation, both the Cave walls and the multitouch display are driven by a single computer with two 2.5 GHz Intel Xeon E5-2640 processors and three NVidia Quadro K5000 graphics cards.

As shown in Figure 1, scholars typically stand behind the pedestal so all walls of the Cave are visible. (Although it is unlikely that ancient orators in the Pnyx used a lectern, the height of the pedestal and its position within the Cave workspace put the user in the position — quite familiar today — of standing behind a lectern to address an audience. In this respect, the arrangement well suits the visualization’s emphasis on an immersive experience of oratorical performance.) Users are not, however, limited to remaining set in a single position. The view is head-tracked (11 camera Optitrack system, NaturalPoint Inc., Corvallis, OR, USA); so, when turning to look to the side or stepping to a different section of the Cave, the perspective of the computer graphics is updated in real time.

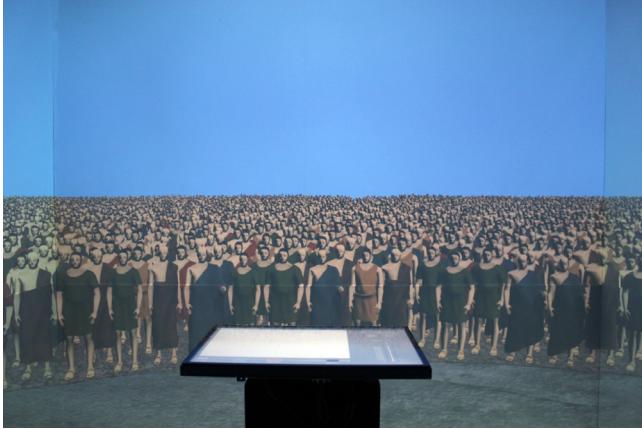


Figure 4: View from the speakers' platform in Pnyx II.

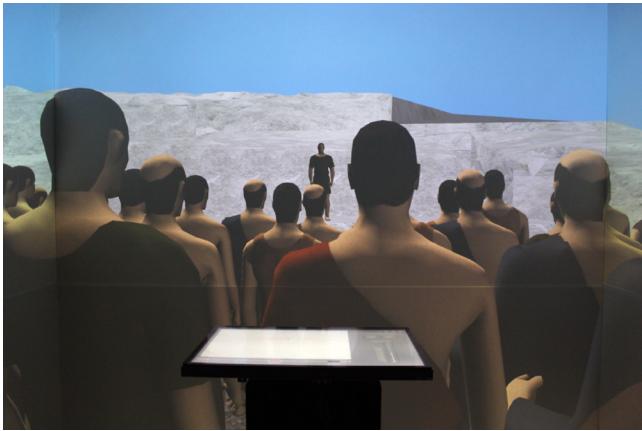


Figure 5: View from within the assembly in Pnyx II.

## 4.2 Navigating in Space and Time

Bema's navigational interface is a modern-day interpretation of the classic World-in-Miniature technique [27]. In addition to displaying the WIM on a multi-touch surface rather than as a projection in space, the WIM here is unique in that it displays multiple versions of the scene superimposed. The versions correspond to the three historical phases of the Pnyx. The currently-displayed phase is highlighted as shown in Figure 3, and dotted outlines depict the key architectural features of the other phases.

Users travel through the virtual space by reframing this "map" view through direct manipulation touch gestures. Following the typical multi-touch manipulation technique, touch with a single finger translates the map with a 1-to-1 correspondence between finger movement and map movement, and touch with multiple fingers does the same while also rotating the map if the finger positions indicate a rotation.

As illustrated in Figures 3-5 and the accompanying video, the virtual position of the Cave relative to the Pnyx models updates in real time based on these gestures. The Cave position is indicated on the map by the crosshairs at the center of the map window. Note that these crosshairs never move, so in this interaction technique, the user moves the map relative to a fixed Cave icon, rather than moving a Cave icon relative to a fixed map. There is conflicting evidence in the literature regarding which of these two possible approaches to designing map-based navigational techniques is preferable (e.g., [1]). Our intuition is that for the general case users might find it more natural to move a "you are here" marker on a map rather than

the other way around. However, for this particular application, we believe the design choice we have made is appropriate. As in Ajaj et al.'s view-based method [1], when users rotate the virtual map in Bema, the virtual model rotates in the same way so that the user is never required to perform a mental rotation in order to relate the two complementary displays.

Our justification for this design decision begins by recognizing that we are using a semi-immersive display. If we were working with a fully-immersive display, we could leverage the well-known advantage of these displays [17] and rely upon head-tracking alone in order to view the whole scene – there would be no need for a special rotation widget or interface since the user could simply turn around to view the scene from any direction. With our 4-wall semi-immersive Cave display, some interface to rotate the scene is required in order to look behind us. For analysis of the Pnyx site in particular, we reason that this scene rotation operation will be frequent – recall, the speaker's position flips by 180 degrees when moving from Phase I and II. Since our application relies heavily upon maintaining good spatial awareness and large rotations of the scene will be required frequently, we argue that reducing the cognitive load that would be required for the user to match the orientation of the map with the orientation of the virtual world is an important design goal. Our interface meets this goal by enforcing the constraint that the two displays always stay in alignment.

As the user navigates, the interface also adjusts the vertical position of the model so as to maintain the constraint that the user's physical feet touch the virtual ground. To keep the motion of the virtual model smooth, the constraint is implemented as a gradual animation. Thus, when moving over rough terrain, the model gradually "settles" to the correct height rather than "jumping" up and down throughout the navigation.

Navigation in time (across the different phases of the Pnyx) is accomplished via a touch-activated timeline widget that is displayed horizontally just below the map as shown in Figure 3.

## 4.3 Interacting with the Assembly

Additional visualization system controls on the right side of the multi-touch display include a slider to adjust the size of the virtual assembly gathered at the Pnyx and a toggle to select the crowd simulation method utilized to position the assemblymen.

The most innovative interaction to perform with the assembly is an interactive visualization of speech intelligibility. Motivated by the potential to utilize VR for exposure therapy, controlled studies have previously demonstrated that VR users feel a strong sense of presence when performing tasks in front of virtual audiences [22, 26]. Informally, we notice a similar reaction – standing on the speakers' platform and looking out at the Pnyx assembly elicits a sense of how daunting it must have been to deliver a speech to such a large assembly. To reinforce this and provide an opportunity for experiential analysis of the actual volume with which one would have to speak in order to be heard in the various phases of the Pnyx, the Bema interface analyzes the user's speech and provides an interactive visualization of speech intelligibility via a color map applied to the assemblymen models as shown in Figure 6.

In the current implementation, a simplified model for speech intelligibility is utilized. (In a parallel research thrust we are developing a set of more complex and precise acoustic simulations for each phase of the Pnyx that we plan to substitute as they come online.) The current model treats the speaker as a directional sound source and is a function of several parameters: speech level as recorded from the sound level meter ( $s$ ), distance from the speaker ( $d$ ), the speaker's heading vector as reported by the head tracker ( $h$ ), and the vector from the speaker to a listening position ( $l$ ). For a given listening position defined by  $d$  and  $l$ , the model returns  $SI$ :

$$SI = \frac{s(\vec{l} \cdot \vec{h})}{d^2}. \quad (1)$$

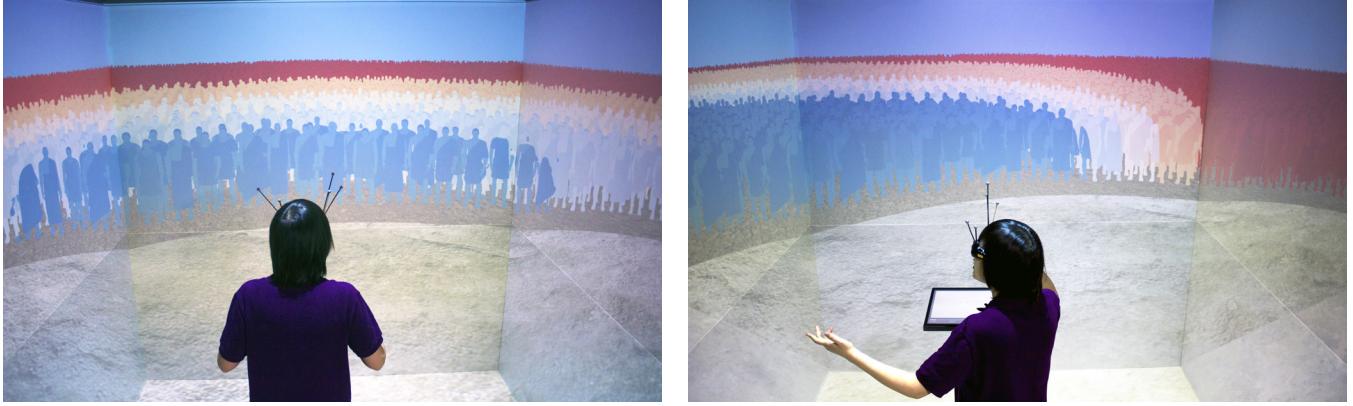


Figure 6: The first person real-time visualization of speech intelligibility shown within Pnyx III.

Figure 6 shows the result of evaluating this equation once for each member of the assembly and applying a sequential color map to the results. Our implementation utilizes an I-CubeX Loud v1.3 sensor attached to a Wi-microDig v6.1 (Infusion Systems Ltd., Montreal, Canada) for input, and  $s$  is recalculated at each rendering frame based on a weighted average of the most recent few seconds of sensor readings.

In addition to the visual presentation of speech intelligibility, Bema can also operate in a virtual orator mode wherein a virtual character recites in ancient Greek the text of Demosthenes' Third Philippic. The audio is rendered in the virtual environment via an eight speaker spatial audio system (one speaker at each corner of the Cave structure). The speech audio level is adjusted according to equation 1, but, in this case, the speaker's position is set to that of a virtual orator standing on the speaker's platform and the listening position is set to the position of the user in model space coordinates. The resulting speech audio is mixed with crowd and wind noise. Currently, these are assumed to be at a constant level throughout the model; however, accurate simulation of these aspects of the model is also part of future research plans.

#### 4.4 Facilitating Comparative Analysis

Comparative analysis is a fundamental problem in data visualization, one that has received renewed interest in recent years. However, much of this work has focused on 2D information visualizations. While visual superposition, juxtaposition, and hybrids are often useful for comparative 2D visualization [9], these techniques are less directly applicable to first-person immersive 3D visualization since it is difficult to stand inside two or more virtual worlds simultaneously. Therefore, the Bema interface adopts the approach of registering all three models of the Pnyx within the same coordinate system but displaying one model at a time. To facilitate comparative analysis, particular attention is paid to supporting the cognitive task of maintaining spatial context while transitioning between phases.

Maintaining context during transitions is accomplished via a “staged” animated transition. Staged animations have proven useful for maintaining context during transitions between 2D statistical graphics [11], and Bema introduces a reinterpretation of this technique for transitioning between co-registered 3D models that differ in the altitude of the ground plane. The first stage of the animated transition is a cross-dissolve (implemented via alpha blending) that causes the previous model to fade out of view and then, the next model to fade into view. During this short transition, which lasts two seconds, key geometric features on each of the models are simultaneously visible, enabling users to create a mental model of the correspondence. Importantly, the viewer's height relative to the virtual model does not change during this first transition. This makes

it easier for users to understand the spatial relationships, but since Pnyx I, II, and III have different altitudes for the ground plane of the model (due to renovations between phases), this also means that the user may be “underground” or “floating in the air” after the first stage of the transition. The second stage of the transition is then to gradually “settle” the user to the ground, animating the vertical position of the virtual model relative to the Cave until the virtual ground plane aligns with the physical position of the user's feet. During both stages of the animation, a virtual surveyor's meter stick is displayed in front of the user. This further helps to compare the architecture of the two phases and makes explicit the dramatic changes to the hillside of the Pnyx.

### 5 EXPERT USER EVALUATION

Since Bema was designed specifically to support expert users, our evaluation is structured as a discussion of the series of insights obtained through user feedback with our team member who is an established professor and scholar of ancient Greek rhetoric and oratory.

#### 5.1 Methodology

Bema has been developed over the course of almost three years, and the interdisciplinary research team met regularly during this period to review and critique the interface. The use of the multi-touch display and the navigational interface both underwent two major design-evaluate-redesign cycles; so our results include comparative evaluation with the alternative interaction techniques that were first implemented.

#### 5.2 Results and User Feedback

The accompanying video provides several examples of the interactive visualization system. The user feedback presented here is organized based upon the most important design decisions made in developing the user interface.

##### 5.2.1 Use of Immersive First-Person Visualization

Preliminary analysis of the 3D models depicted in Figure 2 was conducted via desktop-based visualization; the first aspect of the interface to assess is, therefore, the design decision to utilize an immersive, head-tracked stereoscopic visualization. In addition to immediately identifying a modeling error that was missed throughout months of desktop-based visualization, we found that the 3D head-tracked stereoscopic view facilitated a variety of new spatial analyses of the site. One example comes from an analysis session where our Pnyx scholar was examining the slope of the hill after the renovation from Pnyx I to II, “This is a 4-degree slope ... and we used it [when we created the 3D model of this Phase] because that

is what the excavators suggested ... you can see here, at 4 degrees it makes perfect sense. You're preserving a sight line with a minimal amount of effort." We found this type of spatial analysis more immediate, natural, and accurate in the Cave as compared to desktop environments.

We were also curious about the extent to which scholars who had visited the Pnyx site in the past would relate their prior real-world experiences with the virtual experience provided by Bema, something cited as a key advantage of the ARCHAIVE system discussed in related work. We see strong evidence for this, for example, during one Bema session our Pnyx scholar felt compelled to relay to the group what he had observed at the physical Pnyx site with reference to the virtual model, "I can stand right here, and these benches [pointing with his hand and looking to the left] are deteriorating, and these here [now pointing and looking to the right] are almost gone – you can only make out a couple of cuts [in the stone bench], and we know [from archaeological photos and drawings] it was better 50 years ago, so it's degrading." Here, the utility of a head-tracked, immersive view is obvious. Our experience is that this is a powerful catalyst for new discussion and thinking.

### 5.2.2 Navigation and Context-Preserving Transitions

The navigational interface described in Section 4.2 was designed in response to feedback on an earlier 3D wand-based interface for virtual flying. Our Pnyx scholar reported, "this is for me many, many times better... specialists are very familiar with these superimposed plans." Together, we reason that there are distinct advantages of the current design. First, since the plan views in the WIM illustrate multiple phases superimposed in the same image, they provide more information compared to the alternatives of navigating without the use of a WIM or displaying a more traditional WIM that is rendered as 3D miniature instance of the virtual world. Second, the touch-based interface is easier to control and less fatiguing than a wand, especially for users who have not previously worked regularly with 3D user interfaces in VR.

The staged animations described in Section 4.4, which were designed to preserve a sense of context while transitioning between scenarios, were successful. What's more, they were utilized during analysis sessions for a style of comparative analysis that we had not anticipated during design. Our Pnyx scholar developed a strategy of repeatedly transitioning back and forth between phases. In one instance he reported, "what I am doing right now is testing/checking different things about the model. The stairs right here [spoken at the beginning of the cross dissolve stage of the animation] relate to the stairs in this phase that cut in right here [spoken at the end of the cross dissolve]." He continued, "these stairs don't exist, it's been carved away [at the beginning of another cross dissolve], but on this surface, there is a slight trace of steps [at the end of the cross dissolve]." Thus, repeatedly transitioning back and forth between phases served even more of a role in comparative analysis than we had intended. Rather than simply aiding the process of maintaining context during transitions, this proved to be a useful direct technique for visual comparison.

### 5.2.3 Comparing Different Assemblies

The ability to compare assemblies using different viewpoints, phases, and assembly sizes integrated into a single interface was evaluated as critically important: "having the variability available for *all* of these – all within your control with an intuitive way to move between them – is fantastic. I didn't realize that going into the project." Figure 7 shows example views from one analysis session. Viewing Phase I with 5,000 assemblymen provided a visual check that the reconstruction (both architecture and assemblymen) makes sense; the assembly fits within the available space. With 6,000 assemblymen, we observe that the assembly is full, some people must stand on the hillside and on the walls – just what would have hap-

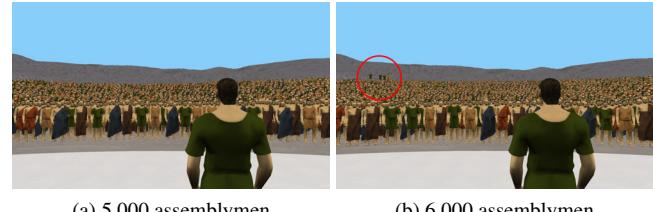


Figure 7: Pnyx I with different size assemblies as simulated with PAM. The circled area shows that it has reached its capacity.

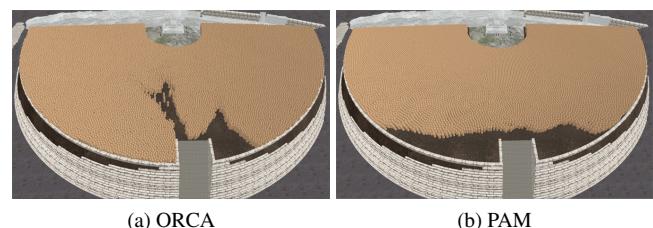


Figure 8: Comparison of a maximum capacity arrangement of the assembly as calculated by the two crowd simulation methods tested – 14,000 assemblymen in Pnyx III.

pened in ancient Athens. Upon seeing this, our Pnyx scholar reported, "I'm getting a lot of confidence in what this shows."

This statement of confidence is an example of an outcome from the high-level task of analyzing through experience how different assemblies may have fit within the space of the Pnyx. What we observed during analysis sessions is that insights come, new hypotheses are formed, and confidence builds over time as the various user interfaces described here are utilized in combination. Often, a first step is to explore the space and architecture, navigating through space and between different phases of reconstruction, then attention often turns toward analysis of the assembly and implications for how speeches must have been delivered and received.

Figure 8 illustrates an example of differences observed between the two crowd simulation algorithms used to realistically arrange the assemblymen. We observed that both methods are a significant improvement relative to arrangement via a regular grid, random sampling, or Poisson disk sampling, all of which were explored in prior iterations. The crowd simulation techniques result in an organic, natural arrangement that appears physically plausible in comparison to other techniques. In the example shown in Figure 8, the Predictive Avoidance Model approach produces a more plausible result, but we found that this was not the case for all scenarios. We are unable to conclude from this work that one technique regularly outperforms the other; however, we can identify several interesting paths for future work based upon lessons learned.

## 6 FUTURE WORK AND CONCLUSIONS

In future work, we are most excited to extend the use of crowd simulation within Bema. In our current implementation, we assume that each simulated agent tries to reach a predefined goal position close to the speaker area. A more realistic approach may be to partition the Pnyx auditorium into predefined seating zones — it is not clear that this was done in practice, but the possibility has been noted in the literature (e.g. [18]). Spreading out the goals could improve the flow rate and relieve "constant pushing behavior" that we have observed in some of the current simulations.

Another interesting and more challenging problem is to dynam-

cally assign seats to the members of the assembly. We have reasons to believe that the Athenians did not remain stationary throughout the entire assembly proceedings. It is likely many would move in response to environmental changes in order to ensure their comprehension of the speaker. This was very important, as natural phenomena, like the wind, could significantly impair the hearing, especially in tightly-packed gatherings or when the speaker had a quieter voice. We believe that goal-seeking *proxy* agents could be added to our simulation framework in order to accomplish this; these agents would dynamically explore the scene and assign appropriate seats to the virtual attendants following the approach of Yeh et al. [31].

The Bema interface has already indicated early impact in the application domain. From this work, we hope that other researchers can gain insight about possible methods for combining, adapting, and extending 3D user interaction techniques to create full interactive systems for data analysis and exploratory visualization. Overall, we see great potential for continued use of advanced 3D user interfaces for cultural heritage, digital humanities, and experiential data analysis applications and encourage the use of case studies applied to real immersive data analysis challenges as an important continued mode of 3D user interface research.

## ACKNOWLEDGEMENTS

Thanks to Sarit Ghildayal for assistance with the implementation and video and to Azidah Rabbani for developing architectural models. This work was supported in part by a grant from the National Science Foundation (IIS-1054783) and donations from Nvidia and Autodesk Research. This work utilizes the VRPN library maintained by UNC with support from NIH/NCRR and NIH/NIBIB award 2P41EB002025.

## REFERENCES

- [1] R. Ajaj, F. Vernier, and C. Jacquemin. Navigation modes for combined table/screen 3D scene rendering. In *Proceedings of the ACM International Conference on Interactive Tabletops and Surfaces*, ITS '09, pages 141–148, New York, NY, USA, 2009. ACM.
- [2] C. W. Blackwell. Athenian democracy: an overview. In *Demos: Classical Athenian Democracy The Stoa: A Consortium for electronic publication in the humanities* ([www.stoa.org](http://www.stoa.org)), 2003.
- [3] R. A. Bolt. Put-that-there: Voice and gesture at the graphics interface. *SIGGRAPH Computer Graphics*, 14(3):262–270, July 1980.
- [4] D. A. Bowman. Interaction techniques for immersive virtual environments: Design, evaluation, and application. *Journal of Visual Languages and Computing*, 10:37–53, 1998.
- [5] D. A. Bowman, E. Kruijff, J. LaViola, and I. Poupyrev. *3D User Interfaces: Theory and Practice*. Addison-Wesley Professional, Boston, 2005.
- [6] D. Coffey, N. Malbraaten, T. Le, I. Borazjani, F. Sotiroopoulos, A. G. Erdman, and D. F. Keefe. Interactive slice WIM: Navigating and interrogating volume datasets using a multi-surface, multi-touch VR interface. *IEEE Transactions on Visualization and Computer Graphics*, 18(10):1614–1626, 2012.
- [7] L. M. Encarna  o, O. Bimber, D. Schmalstieg, and S. D. Chandler. A translucent sketchpad for the virtual table exploring motion-based gesture recognition. *Computer Graphics Forum*, 18(3):277–286, 1999.
- [8] A. E. Foni, G. Papagiannakis, and N. Magenat-Thalmann. A taxonomy of visualization strategies for cultural heritage applications. *ACM Journal on Computing and Cultural Heritage*, 3(1), 2010.
- [9] M. Gleicher, D. Albers, R. Walker, I. Jusufi, C. D. Hansen, and J. C. Roberts. Visual comparison for information visualization. *Information Visualization*, 10(4):289–309, Oct. 2011.
- [10] M. H. Hansen. Reflections on the number of citizens accommodated in the assembly place on the pnyx. In *The Pnyx in the History of Athens*, pages 23–33. ed. Bi  n Fors  n and Greg Stanton. Finnish Institute at Athens, 1996.
- [11] J. Heer and G. Robertson. Animated transitions in statistical data graphics. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1240–1247, 2007.
- [12] C. L. Johnstone. Greek oratorical settings and the problem of the Pnyx: Rethinking the Athenian political process. In *Theory, text, context: Issues in Greek Rhetoric and Oratory*, pages 97–127. ed. Johnstone, Christopher Lyle. Albany: SUNY Press, 1996.
- [13] I. Karamouzas, P. Heil, P. van Beek, and M. Overmars. A predictive collision avoidance model for pedestrian simulation. In *Motion in Games*, volume 5884 of *LNCS*, pages 41–52. Springer, 2009.
- [14] D. F. Keefe, D. A. Feliz, T. Moscovich, D. H. Laidlaw, and J. J. LaViola Jr. CavePainting: A fully immersive 3D artistic medium and interactive experience. In *Proceedings of I3D 2001*, pages 85–93, 2001.
- [15] D. F. Keefe, A. Gupta, D. Feldman, J. V. Carlis, S. Krehbiel Keefe, and T. J. Griffin. Scaling up multi-touch selection and querying: Interfaces and applications for combining mobile multi-touch input with large-scale visualization displays. *International Journal of Human-Computer Studies*, 70(10):703–713, Oct. 2012.
- [16] King's Visualisation Lab, King's College London. Theatron: Theatre history in Europe. <http://www.kvl.cch.kcl.ac.uk/theatron.html>, 2009.
- [17] J. Kjeldskov. Interaction: Full and partial immersive virtual reality displays. In *Proceedings of Information Systems Research Seminar in Scandinavia, IRIS24*, pages 587–600, 2001.
- [18] K. Kourouniotes and H. A. Thompson. The Pnyx in Athens. *Hesperia*, 1:96–217, 1932.
- [19] J. LaViola. MSVT: A virtual reality-based multimodal scientific visualization tool. In *Proceedings of the Third IASTED International Conference on Computer Graphics and Imaging*, pages 1–7, 2000.
- [20] J. Ma  n, S. Haegler, B. Yersin, P. Mueller, D. Thalmann, and L. Van Gool. Populating ancient Pompeii with crowds of virtual Romans. In *8th International conference on Virtual Reality, Archaeology and Intelligent Cultural Heritage VAST*, pages 109–116, 2007.
- [21] R. Osfield, D. Burns, et al. Open Scene Graph. <http://www.openscenegraph.org>, 2004.
- [22] D. P. Pertaub, M. Slater, and C. Barker. An experiment on public speaking anxiety in response to three different types of virtual audience. *Presence-Teleoperators and Virtual Environments*, 11(1):68–78, 2002.
- [23] J. H. Rindel. Roman theatres and revival of their acoustics in the ERATO project. *Acta Acustica united with Acustica*, 99(1):21–29, 2013.
- [24] W. Shao and D. Terzopoulos. Populating reconstructed archaeological sites with autonomous virtual humans. In *Intelligent Virtual Agents*, pages 420–433. Springer, 2006.
- [25] G. R. Stanton. The shape and size of the Athenian assembly place in its second phase. In *The Pnyx in the History of Athens*, pages 7–21. ed. Bi  n Fors  n and Greg Stanton. Helsinki: Finnish Institute at Athens, 1996.
- [26] C. Stinson and D. A. Bowman. Feasibility of training athletes for high-pressure situations using virtual reality. *Visualization and Computer Graphics, IEEE Transactions on*, 20(4):606–615, 2014.
- [27] R. Stoakley, M. J. Conway, and R. Pausch. Virtual reality on a WIM: interactive worlds in miniature. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 265–272. ACM Press/Addison-Wesley Publishing Co., 1995.
- [28] J. van den Berg, S. J. Guy, M. Lin, and D. Manocha. Reciprocal n-body collision avoidance. In *Robotics Research: The 14th International Symposium ISRR*, volume 70 of *Springer Tracts in Advanced Robotics*, pages 3–19. Springer, 2011.
- [29] E. Vote, D. A. Feliz, D. H. Laidlaw, and M. S. Joukowsky. Discovering petra: Archaeological analysis in VR. *IEEE Computer Graphics and Applications*, 22(5):38–50, 2002.
- [30] J. N. Wall et al. Virtual paul's cross project: A digital re-creation of John Donne's gunpowder day sermon. <http://vpcp.chass.ncsu.edu>, 2013.
- [31] H. Yeh, S. Curtis, S. Patil, J. van den Berg, D. Manocha, and M. Lin. Composite agents. In *ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, pages 39–47, 2008.