Repopulating the Colosseum: Investigations into AI Simulation for Ancient Crowd Composition

ZACKARY HEGARTY, Indiana University, IN, USA

This study explores the application of artificially intelligent (AI) virtual agents to simulate the movement of spectators through a high-quality digital reconstruction of the ancient Colosseum. Building on the foundational work of Gutierrez et al. in "AI and Virtual Crowds: Populating the Colosseum," this research advances the understanding of crowd behaviors within this ancient spectacular structure by exploring how 3D agent-based modeling that accounts for historical context (class and gender) in a more granular way than is typically found in the literature might provide novel results. The simulation itself was generated in the popular 3D development platform, Unity, and utilized its built-in "AI Navigation" system. This article highlights the continued potential of AI-driven simulations to offer insights into ancient crowd behaviors and reconstructive archaeology and concludes by discussing future directions for research around AI simulations for within the domain of cultural heritage.

Keywords:
Artificial Intelligence, Crowd Simulation, Spatial Analysis, Digital Archaeology.

SDH Reference:
Hegarty, Zackary. 2024. Repopulating the Colosseum: Investigations into AI Simulation for Ancient Crowd Composition, SDH, 8, 2, 67-96.

https://doi.org/10.14434/sdh.v8i2.40246

1. INTRODUCTION

The use of virtual reality environments and artificial intelligence (AI) for exploring problems in cultural heritage is an existing but still developing area of research [Guidi and Frischer 2020]. These tools hold significant promise for analyzing archaeological sites and contexts from the past, offering insights into their historical functions. However, such digital-forward methods should not be seen as a standalone solution generator. Instead, they serve as valuable tools within the broader toolkit of cultural heritage practitioners and practitioners from allied disciplines. Their strength lies in adding to or transforming preexisting data or theories into new information, thereby enriching the array of insights that can be used to draw conclusions about historical sites, materials, and behaviors.

This paper aims to extend and update the foundational work of [Gutierrez et al. 2007], focusing on the question of how the large crowds at the ancient Colosseum utilized the space and whether there were sociocultural idiosyncrasies that influenced its function in ways not immediately apparent to modern eyes. The Colosseum is an excellent subject for such an examination because of its extensive documentation and the fact that much of the monument still exists. By revisiting the Colosseum simulation nearly twenty years after the initial publication, this research leverages new and more accessible technologies to update and expand their simulation. This modernization not only allows for easier repeatability, but it also introduces more sophisticated agent behaviors within the simulation, providing a more nuanced explanation of how the space functioned. Ultimately, this article hopes to demonstrate the continued value of AI-driven simulations and how the results generated by them can be useful to cultural heritage research and research in allied fields.

The paper is structured as follows. Section two reviews the historical facts of the Colosseum and subsequently situates the creation of a digital simulation within the broader context of cultural heritage work and AI simulation development. Section three explores the construction of the simulation, highlights its iterative design processes, and presents the results of the simulation. Section four discusses these results, focusing especially on their implications for cultural heritage practitioners. Finally, section five concludes by exploring areas for improvement and areas ripe for future development.

2. LITERATURE REVIEW

2.1 Historical Review

The Colosseum (more correctly, but less colloquially, known as the Flavian Amphitheatre) was the largest in ancient Rome, covering 3357 m^2 of the center of the city [Nielsen 2006]. The structure was constructed by the Flavian dynasty with its first stone laid by the Emperor Vespasian sometime around 72 CE and its inauguration celebrated by his son, the Emperor Titus, in 80 CE [Hopkins and Beard 2005]. This structure was designed as a space dedicated to public spectacle but was also loaded with political meaning not least of which because it was located on top of a portion of the *Domus Aurea* complex [Bomgardner 2021; Hopkins and Beard 2005; Lovatt 2016]. This complex was built by the Emperor Nero who, in turn, had demolished an even older sector of the city to build what was deemed by his critics to be nothing more than a pleasure palace. While the feeling among the masses about this imperial palace or its demolition is nebulous, the popularity of the games seems clear [Welch 2007]. Not only was the space host to exciting spectacles but it was also a locus for the emperor to dole out favors and gifts. The political function of this space also flowed the other direction; spectacles were one of the last places for the mob to appeal directly to the emperor. Spectacular spaces, such as this, did nothing if not embody the phrase, 'to see and be seen' [Hopkins and Beard 2005; Lovatt 2016; Welch 2007]

In terms of its design, the Colosseum was divided into five seating galleries and held an audience of approximately fifty thousand, quite large even by modern standards [Bomgardner 2021]. Seating organization was regulated by the *Lex Julia Theatralis* and was thus stratified based on class and gender by force of law, as is well described in Suet. *Aug.* 44 [Lovatt 2016; Rawson 1987; Rose 2005]. Moving out and up from the arena-side seats, the organization was generally as follows: Senators and high officials seated in the *podium*, the order of equestrians seated in the *maenianum primum*, monied/togate plebians seated in the *maenianum secundum immum*, non-togate plebians and likely foreigners/freedmen/slaves seated in the *maenianum secundum summum*, and finally any women in attendance were relegated to the uppermost gallery—the *maenianum summum in ligneis* (Fig. 1) [Bomgardner 2021; Welch 2007]. Although, I hasten to note that while this organization is rather well accepted, it is nonetheless interpretive and open to slight adjustment.

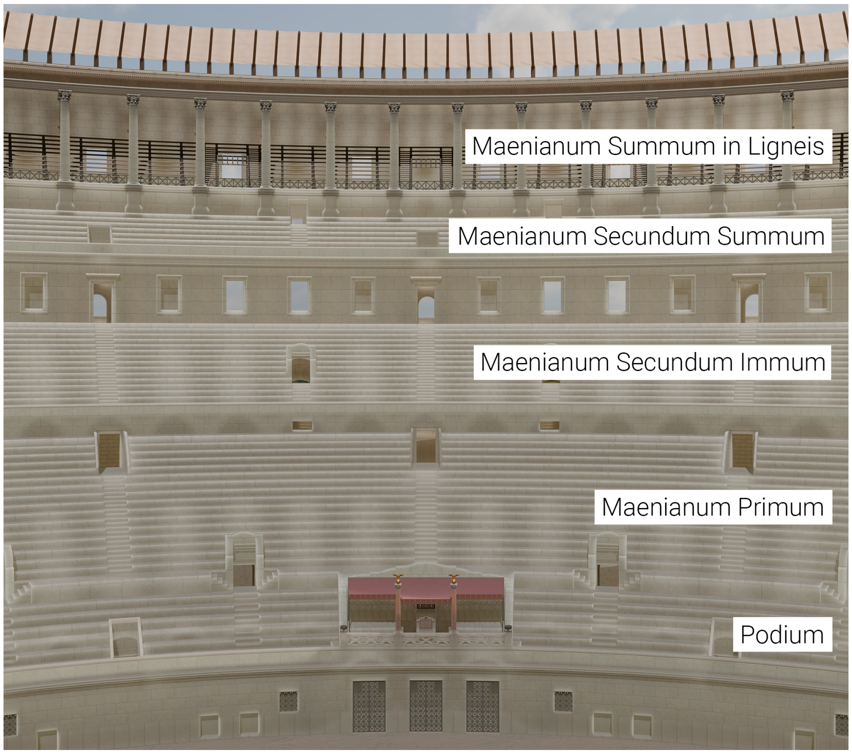

Figure 1. The seating galleries of the Colosseum; a digitally rendered and labeled interior view. Reconstruction is courtesy of and copyright by Flyover Zone, Inc., 2024. All rights reserved.

Beyond being impressively large for the day, it was also innovative in its design. It featured a multi-story, Greek-inspired façade with columns decorating the 80 arches which ringed its exterior [Welch 2007]. Additionally, it had extensive underground works for realizing complex events up on the arena floor which included tunnels, holding pens, and lifts [Hopkins and Beard 2005]. Additionally, in the uppermost reaches of the structure it had a massive awning, or *velarium*, used to shade spectators from the sun. This was operated by a specially trained group of seamen since it was, in essence, a massive sail [Alfano et al. 2015; Montilla 1969].

As nuanced as the infrastructural and backstage activities were, so too was the organization of the crowd of spectators seeking, as put by the author Juvenal in *Sat*. X, *panem et circenses*. A visitor to the games would begin their day by entering into the plaza surrounding the Colosseum seeking out their seat by using their *tessera*, here "ticket," which likely denoted at least the appropriate gallery, section, and row, if not also the gate, event, and event type [Bomgardner 2021]. While no tickets for the Colosseum survive, their use is strongly inferred from textual evidence and extant tickets from around the Empire [Futrell 1997; Golvin 1988; Hopkins and Beard 2005]. For example, a ticket from Arles notes: *CAV(ea) II, CVN(eus) V, GRAD(us) X, GLAD(iator) VELA(rium)* [Golvin 1988], "Second gallery, fifth section, tenth row, [for the] gladiator [fight to be held beneath the] awning." These tickets were probably doled out via the patron-client networks so critical to the function of Roman society

[Bomgardner 2021]. After examining this ticket and managing to move through the crowds, the visitor would navigate to the gate as indicated on their ticket to proceed within the structure. It was also possible that, beyond the obviously numbered gates on the exterior, the structure had signage perhaps integrated into the painted stucco likely to have decorated its interior to aid in spectator navigation; this is, however, heavily conjectural [Campisi 1987; Hopkins and Beard 2005; Pearson 2013; Zink 2014].

Whether or not these tickets were checked upon entry is unknown, but that the tickets were enforced is known. Ancient texts like Plaut. *Poen*. prol. 19 relate how *dissignatores,* in this context "ushers," were empowered to circulate through spectacular spaces such as this and ensure the right people were in the right place [Bomgardner 2021]. Interestingly, if one considers the ill-repute workers of this title suffered, this was perhaps the only way by which men of such a station could view the games up-close; moreover, it would have been one of only a few instances where they might have authority over men of potentially far higher station [Bond 2016; Rawson 1987]. Their job could have been executed either by examining the ticket or by inspecting the spectator, given the implementation of a rather rigorous dress code. In all spectacular spaces the law regulated dress and so the first three galleries were all directed to wear a version of the toga, creating a white-out in the first approximately two-thirds of the audience with splashes of purple from youths and the higher classes in attendance [Bomgardner 2021; Brøns et al. 2017; Zanker 1998]. The visitor would then consult their ticket to ensure they were in the correct section upon arriving at their intended vomitorium and would subsequently find their seat within the *cavea*.

The evidence presented above concerning the construction, social stratification, ticketing, etc. is all important to keep in mind when interpreting the flow of traffic within the Colosseum because it is, as the simulations which form the core of this research demonstrate, critical to understanding the experience of spectators in the space and the process of simulating that experience.

2.2 Technology Review

This study was born as a follow-up to the foundational work conducted by Diego Gutierrez and his team at the University of Zaragoza along with the digital archaeologist Bernard Fischer in the early aughts [Gutierrez, et al. 2007]. However, it is worthwhile to note that the use of digital agents in this way, i.e., within a virtual crowd, is situated as part of a broader practice in cultural heritage and allied disciplines for utilizing, broadly, digital characters to populate and otherwise serve within digital environments. Such digital characters include any, "…software entities that look and act like real or imaginary creatures in a computer-generated environment" [Vosinakis 2020].

Such a definition obviously encompasses a very broad range of implementations and as such can be divided by function. There are several classification schema which review and attempt to classify these digital entities in cultural heritage research, of which [Antunes and Correia 2022; Vosinakis 2020] are quite useful. To synthesize these schemas, there are generally four types of digital character: animated props, virtual guides, virtual agents, and learning virtual agents. One might consider this division like a sliding scale defined by character interaction and independence. Animated props are the least complex and can be seen as more akin to moving art pieces that do not interact with human actors nor each other in any meaningful way. A designer also completely predetermines their actions. While such props are quite basic, I would also mention that they can be critical for a digital scene by bringing life to an otherwise lifeless digital world. Virtual guides, meanwhile, not only bring life to a digital environment, but also allow for interactive experiences. They may provide culturally contextualized information concerning a scene, navigational/technical information, etc.; in essence, these characters are developed in a way that allows them to helpfully interact with a human user within a digital scene.

Differing from virtual guides, virtual agents have intentions of their own and do not primarily have an interest in interacting with human users. Rather, these agents have, as their name implies, agency. This agency means that they will navigate, take actions, decide on goals, etc. on their own accord, or better, based on conditions predetermined by the designer of the scene. A prime example in the domain of cultural heritage is a virtual simulation of Pompeii which contained virtual agents making decisions like shopping or navigating to socially appropriate sections of the city [Maïm et al. 2007].

Such agents will more truly act on their own accord when they are imbued with some artificial intelligence mechanism that allows for learning. These learning virtual agents will still consider goals given to them from the designer but will forge their own paths to complete them according to their experience. When this is conducted across multiple training sessions, the agents may develop unique solutions to given problems by learning from past mistakes and solutions. This is especially true when various costs for different decisions are included in the scenario. This does not necessarily mean that agents of this variety are going to find the same solution that a human would find. Rather, they are more likely to find the most optimal solution to a given problem, though with the possibility of being imperfect decision-makers like real humans depending on their design [Lake 2020]. It is not immediately clear if this type of agent is useful for cultural heritage practitioners because the agent behavior may not be explainable enough for a field that is forced to operate within an experimental space already so ripe with unknowns [Bickler 2021; Lakkaraju and Bastani 2020].

The current study, which focuses on agents who do not learn from previous experiences but still act independently, falls into the category of virtual agents. Using these virtual agents bolsters explainability since each experimental parameter can be traced throughout the simulation. This work is also a form of agent-based modelling. Generally, the use of agent-based modelling within the domain of cultural heritage, and especially archaeology, is reasonably common [Cegielski and Rogers 2016; Lake 2020]. However, it is almost exclusively applied to 2D environments and usually for landscape-level surveys. Applications to 3D CH spaces are not particularly widespread but have been considered since at least the late 1990s [Musse and Thalmann 1997]. Concerning the Colosseum, to the author's knowledge, three simulations using virtual agents have been conducted using digital reconstructions of this structure: that of [Gutierrez, et al. 2007] as mentioned above, a simulation conducted for a 2014 National Geographic documentary assisted by digitization experts at University of Arkansas and as reported in [Bomgardner 2021], and a more recent comparative simulation published in the fire safety journal *Fire* [Gravit et al. 2022].

[Gutierrez, et al. 2007] were the first to use advanced computer graphics and AI to simulate the navigational capacity of virtual crowds in cultural heritage sites, specifically focusing on the Colosseum. The researchers developed a detailed 3D model of the structure and utilized AI to create virtual agents that simulate crowd behaviors. These agents were governed by Hierarchical Finite State Machines, enabling them to perceive their environment, make decisions, and navigate through the structure. The study employed a bottom-up approach to crowd simulation, allowing unpredictable and dynamic interactions among the agents. The main aim was to test the ergonomics of the Colosseum, challenging the long-held belief that it was an extremely efficient people-mover by identifying potential bottlenecks through simulation; the results were that the structure could be filled/emptied in 15-20 minutes rather than far lower times as previously suggested. This innovative combination of AI and virtual reality provided the first example of a tool for scholars to test hypotheses and gain insights into the ergonomics of this space thereby laying the foundations for this sort of study in the cultural heritage domain. This work also paved the way for expanding the simulation in various directions, from altering the social characteristics of the crowd to refining the data collection strategy.

Later, the simulation conducted for the National Geographic documentary “Time Scanners” in 2014, focused on the Colosseum’s efficiency in comparison to a modern stadium, the Bird’s Nest in Beijing [Bomgardner 2021; Taylor 2014]. The project was conducted by structural engineer Steve Burrows, CBE, well-known for creating the Bird’s Nest Stadium in Beijing, the Allianz Arena in Munich, and the Etihad Stadium in Manchester, along with members of the Center for Advanced Spatial Technology at the University of Arkansas: Katie Simon, Adam Burns, and Rachael Opitz. The structure was digitized using terrestrial laser scanners, and one-eighth of the building was subsequently reconstructed for use in the simulation. Their simulation revealed that each building could be evacuated in approximately twelve minutes; in fact, the Colosseum was slightly faster. This is congruent with [Gutierrez, et al. 2007], which emptied just a few minutes slower. The simulation utilized the crowd simulation software MassMotion by Oasys to populate the scenario with agents seeking the fastest paths out of the building [Oasys 2015]. However, to the author’s knowledge, this simulation remains unpublished beyond the documentary and general notes from the software house, making it difficult to examine in greater detail. It is nonetheless useful to note that its findings—using industrial-grade crowd simulation software—are consistent with those of other published studies.

The third study simulating the crowds of the Colosseum compared it with the Gazprom Arena in St. Petersburg, Russia [Gravit, et al. 2022]. The study utilized the “Sigma FS (Russia)” software package to model and analyze pedestrian behavior during evacuation scenarios. When the researchers pitted these two structures against each other, they found that the Colosseum, despite being an ancient structure, has an efficient evacuation design, with a specific flow rate of 1.14 persons per second per meter, facilitated by its large staircases. In contrast, the Gazprom Arena, designed with modern infrastructure and commercialization in mind, showed a lower specific flow rate and requires additional organizational and structural elements to manage evacuation effectively. The authors suggest that the Colosseum meets current evacuation standards and can serve as a model for modern stadium design in terms of evacuation efficiency. However, a reasonable critique of this study is that their model of the Colosseum, although constructed from archaeological plans, lacks sufficient detail to conclusively suggest ancient crowd movement, especially in the upper reaches of the structure. While the evacuation time found in this study is consistent with previous studies, this limitation likely fails to fully capture the nuances of ancient pedestrian behaviors, potentially overstating the efficiency and safety of the space. Moreover, its focus on safety standards is probably less useful for cultural heritage practitioners.

With the above taken into account and by being situated within a broader framework of the use of digital characters in cultural heritage, this study aims to underscore the significant role these agents can play in bringing historical environments to life. It seeks to fill a gap in the existing literature regarding simulating the various social identities of the agents entering the Colosseum, since, as is noted in the historical review, the crowds were *de rigueur* socially stratified. It also hopes to examine seating capacity and temporal parameters to hone the accuracy of the simulation. By modifying the virtual agents in a way that captures the most salient identity data, this research aims not only to bring life back to the Colosseum but also to use these agents to explore various crowd configurations/behaviors at play in this ancient space. The ultimate goal is to understand how various simulable conditions might affect the ergonomics of the internal structure of the Colosseum.

3. THE SIMULATION

3.1 Foundations of the Simulation

The construction of a simulation capable of running virtual agents through a digital model of the Colosseum requires three ingredients: a digital reconstruction model of the Colosseum itself, a software platform to manage the digital scenario, and lastly an AI-driven navigation tool to run the virtual agents.

A precise digital model of the Colosseum was essential for the simulation. One of, if not the most comprehensive and highly faithful reconstruction models of the structure in existence has been created by Flyover Zone, Inc., led by Bernard Frischer, who has dedicated at least three decades to this work [Frischer 2024]. This model was obtained from Flyover Zone, Inc. for scholarly purposes for use within this research. Utilizing such a model is significant because it ensures a more accurate simulation by minimizing errors from the reconstructed environment. This reliability is crucial for studying crowd behaviors within the Colosseum and using a high-fidelity model such as this helps confirm its trustworthiness.

After selecting the model, the next most important element of the simulation was selecting a software environment within which the simulation might be managed and run. The development platform Unity, also known as a game engine, was chosen for its comprehensive documentation, community support, and open nature, making it (essentially) free to use and ready for immediate deployment [Unity 2024]. These features make it an ideal choice for achieving relatively quick, reliable, and repeatable results. A particular benefit of this platform is that in addition to its ability to manage the simulation, it also includes a built-in AI navigation system. It is an easy-to-use and almost fully plug-and-play system that manages both the mapping of the navigable surfaces and the navigation algorithms which drive the agents. These built-in navigation systems include: a navigation meshing system used to identify the digitally reconstructed environment, the A* algorithm for agent pathfinding, and the Reciprocal Velocity Obstacles algorithm for local collision avoidance [Unity 2024]. In addition, a previous phase of this work examined the functionality of Unity’s default AI navigation system in comparison to that conducted by [Gutierrez, et al. 2007] and found them to be substantially similar [Hegarty 2024]. This similarity allows for results to be compared to those generated by [Gutierrez, et al. 2007] and thereby the other experimental results as detailed in §2.2. After considering all of these factors, it was decided that Unity would function as an ideal tool for this work.

The first step in deploying this AI navigation system is mapping the navigable surfaces available to the virtual agents. The navigation surfaces are generated through a procedural remeshing of the target environment, in this case the Colosseum and the plaza of its surrounds. The result is a navigation mesh (navmesh) which provides the virtual agents information about where it is possible to move based on their own dimensions. This mesh is nothing novel to Unity and is, for example, used within the robotics community for identifying navigable surfaces for autonomous vehicles [Arkin 1989]. The navmesh is capable of discretizing a given space into polygons, which helps inform a navigation agent about navigable areas and the fastest way to generate a path to its destination. As is possible to see in Figure 2, the mesh is a series of squares and triangles which are assigned various costs and upon which an agent can freely move. Costs for moving over various patches of the navmesh can be modified by the user to reflect more or less difficult terrain. For example, the flat walkways surrounding the structure have a cost assignment of “1,” i.e. no cost advantage or disadvantage. Meanwhile, the staircases within the structure have a higher cost which leads agents to avoid them when possible, i.e., agents will avoid unnecessary ‘fatigue’ from the extra effort it would take to climb stairs. In this scenario, agents were slowed on stairs by a factor of three, generally following [Hinman et al. 2014].

During the construction of the navmesh, it was noticed that some areas of the Colosseum were not generating walkable surfaces as expected. This was particularly obvious in the depths of the structure, i.e., within the first- and second-floor mezzanines. These two public-facing concourses, the only two without exterior access to the open-air arcades of the façade, have the tightest staircases and hallways [Gutierrez, et al. 2007]. As such, to force the simulation to build navigable surfaces in these rather difficult spaces, two modifications were made. First, additional geometry was added to some staircases within the two mezzanines to aid the procedural generation of the navmesh. This was done simply by adding a primitive cube, modified to be placed over top of some of the stairs; this provided smoother geometry for the mapping procedure and resulted in a very regular, and thus very useable, navmesh. Second, the agent height, originally set to 1.5 meters, was halved; notably, this pertains to the dimensions used to create the navigation mesh, not to the agents themselves, who retained their original dimensions when instantiated into the space. Although this is not an ideal situation, it reasonably simulates the ability of real people to maneuver their bodies to gain greater freedom when moving through tight spaces.

Since this study took [Gutierrez, et al. 2007] as its starting point, the simulation restricted itself to a quarter of the Colosseum. As the structure is, essentially, radially symmetrical, this allowed for fewer computational resources to be needed to complete the work. However, the use of only a quarter of the model was interpreted in a perhaps non-obvious way. Namely, instead of simply digitally slicing the geometry of the structure and removing it from the simulation completely, the navigable surface was manipulated in such a way so that the only “walkable” surfaces within the structure were within the quarter in question, in this case, the northwest quarter between gates XX and XXXVIII.

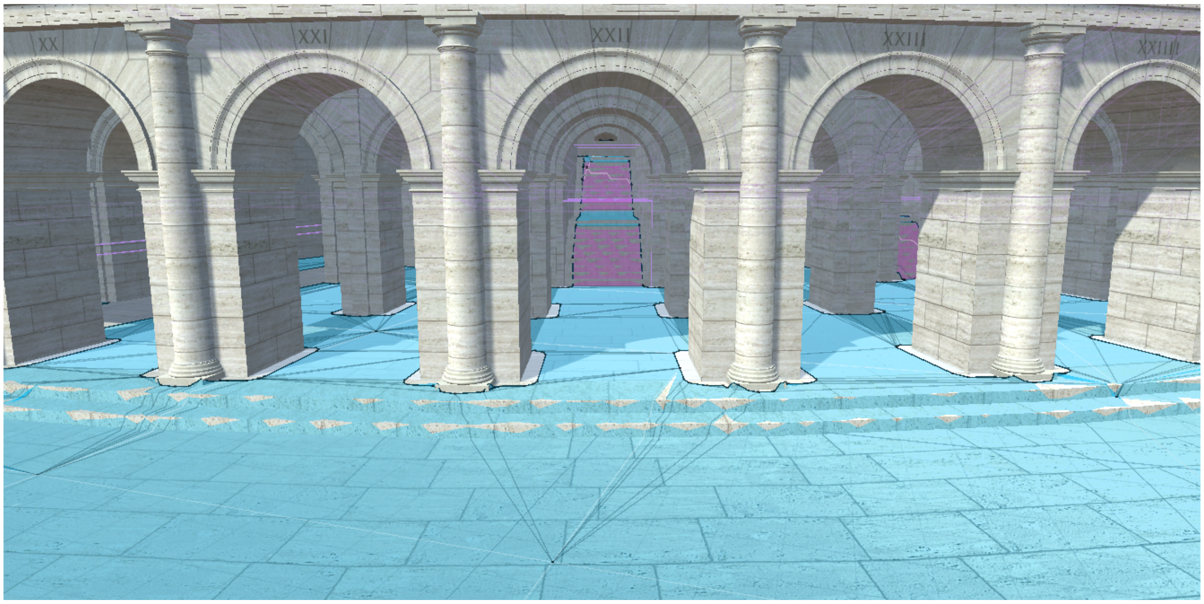

Figure 2. Example of the navmesh as implemented in Unity. Blue mesh denotes walkable surfaces with a cost of "1." Magenta mesh denotes walkable surfaces with a cost of "3." Reconstruction is courtesy of and copyright by Flyover Zone, Inc., 2024. All rights reserved.

The next step was to insert the agents themselves into the digital environment. They have a number of property fields that can be manipulated to fit the needs of a given simulation. This includes, for example, the size and function of the agent. In this simulation, the agents were generated as 1.5 meters tall by 0.75 meter diameter capsules; in Unity, the capsule primitive is the default agent shape. It should also be noted that the 0.75 meter diameter pertained to its collider, an otherwise invisible element which handles the agent's physical collisions with the environment and other agents. This allowed for each individual agent to have a small zone of 'personal space' around their rendered mesh body, which was 0.5 meters in diameter. This helped facilitate, although it did not perfectly ensure, a collision-free simulation.

A number of other agent properties relate to the functioning of the navigation system, one of which is its "priority." An agent's priority is a means to allow certain agents to ignore other agents with a lower priority while navigating. Functionally, this allows high priority agents to, for example, push through to the front of a crowd or queue. Instead of assigning the agents a specific integer value for their priority property, they were assigned a number within a random range using Unity's implementation of an xorshift random number generator [Marsaglia 2003; Unity 2024]. This approach ensured that agents almost never found themselves stuck at an impasse, locked in place by another agent with the exact same priority value. Since agents did not treat every other agent they encountered equally, their consideration of other agents in the crowd was not overly deterministic. Each agent was assigned a priority value within a specific range, creating a controlled yet dynamic situation that resulted in a more believable crowd, with agents rarely getting unrealistically stuck.

Therefore, throughout all simulation runs, these value ranges were used to divide the crowd in various ways. As will be discussed later, this property became one of the key variables of the simulation set examined by this research.

Additionally, changes were made to the agents' speed properties, starting with the introduction of random variation in walking speed. This reflects a more realistic situation where the crowd would no longer simply walk in lockstep towards their destinations. Instead, upon awakening, agents were randomly set to ±25% of an average walking speed of 1.3 m/s; generally following [Schimpl et al. 2011].

The agents are also, by default, programmed to intelligently avoid nearby agents on their way to their destination. For this, they operate by using the Reciprocal Velocity Obstacles (RVO) algorithm. This algorithm works by enabling each agent to assume that other agents it encounters are operating under the same avoidance principles [Van Der Berg et al. 2008]. It avoids collisions by inferring clear paths through calculations on the velocity vectors of nearby agents with equal or higher priority. This algorithm is robust and powerfully replicates the local pathfinding behaviors of pedestrians [Lemonari et al. 2022; Van Der Berg et al. 2008]. Within Unity, the algorithm is slightly restricted because it can only track up to eight nearby agents at a time [Unity 2024]. This limitation affects its functionality in very dense crowds. However, it was retained for approximating crowd behavior due to its ease of use as part of Unity’s default navigation system and its reasonably robust performance.

Another key algorithm implemented in Unity’s AI navigation system is the A* algorithm. This is a widely used algorithm which is capable of generating a least-cost path for a given agent [Lemonari, et al. 2022]. In brief, it is a pathfinding and graph (i.e., navmesh) traversal algorithm that efficiently finds the shortest path between two nodes [Delling et al. 2009]. It also estimates the cost to reach the goal from each node, optimizing the search process by prioritizing paths with the lowest total cost. In this case, agents targeted the mouth of each vomitorium based on their identity and location within the specific quarter of the Colosseum. For example, an agent identified as a Senator would navigate to one of the three podium vomitoria, and so on moving up the structure.

Some changes were made to the simulation to refine the operation of the functionality of the AI Navigation system implemented in Unity. This included modifying the pathfinding logic so that each agent in the simulation would incrementally verify its path toward its destination. This allowed for more believable behavior because, without this modification, an agent would too strictly follow their original path. If, for example, they were pushed by pressure from the crowd, they were observed trying to backtrack to their original path rather than recognizing that they had been pushed closer to their destination. Iteratively re-pathing, in this case every two seconds, resolved this issue and enabled agents to use the crowd's pressure to move toward their target destination.

A second change was related to the costs associated with the navmesh. By default, these costs only affect the A* algorithm’s calculations for the most efficient path and do not affect the actual speed of an agent. Since this was a desired element of the simulation, i.e. the slowing of agents on high-cost surfaces like the stairs, a direct association between navmesh area cost and agent speed was made. This allowed not only for the agent to realistically slow down on staircases but also provided a less arbitrary way by which to determine the cost of the particular navmesh areas; i.e., costs were fixed to the slowing factor of a given zone.

Another change related to the ability of the agents to verify their final destination. It consisted of adding zones to each vomitorium entrance which provided a means for the agents to check if they were in the correct place. If an agent found themselves at an improper vomitorium, which might happen if an agent was pushed by the pressure of the crowd, their priority would be temporarily reset to a very high level so they might be able to push out and correct the situation; a reasonable behavior considering the scenario.

An additional mechanism in the simulation focused not on the agents, but on determining their number and density in a given area. Panels which could sense the agents were added to areas of the first- and second-floor mezzanines and also the ground-floor arcade. The mezzanines in particular were areas determined by [Gutierrez, et al. 2007] to be areas where the most congestion occurred. They also hypothesized that adding a heat map to the simulation would be useful for demonstrating in a more concrete fashion how exactly bottlenecks might be forming within the Colosseum. Accordingly, the simulation in the current study incarnated this idea as an additional method for assessing the space's ergonomics. A more detailed discussion of this system is provided in §3.3; however, it is worth noting that it played a significant role in the development of the simulation.

In addition to setting the simulation parameters, it was necessary to devise a method by which temporal and crowd composition related data might be collected from the scene (Fig. 3). Ultimately, the decision was made to collect this information from a number of virtual cameras situated around the environment. In parallel, an agent counter would be set to track the number of agents (both total and based on target gallery) present in the scene, updating frame by frame and reporting this information to a digital canvas. This canvas was recorded and saved as an MP4 video file, which was then processed to extract single frames at specific intervals to produce analyzable data. In this case, frames were extracted every two seconds, aligning with the agents' path recalculation period. These videos can be examined at: https://vwhldata.sitehost.iu.edu/.

Between the high-fidelity reconstruction model, the default Unity AI navigation systems, and the few modifications as detailed above, preliminary tests of the simulation demonstrated that it was able to produce quite believable crowds with only a vanishingly small proportion of agents halting or pathfinding in a non-believable manner.

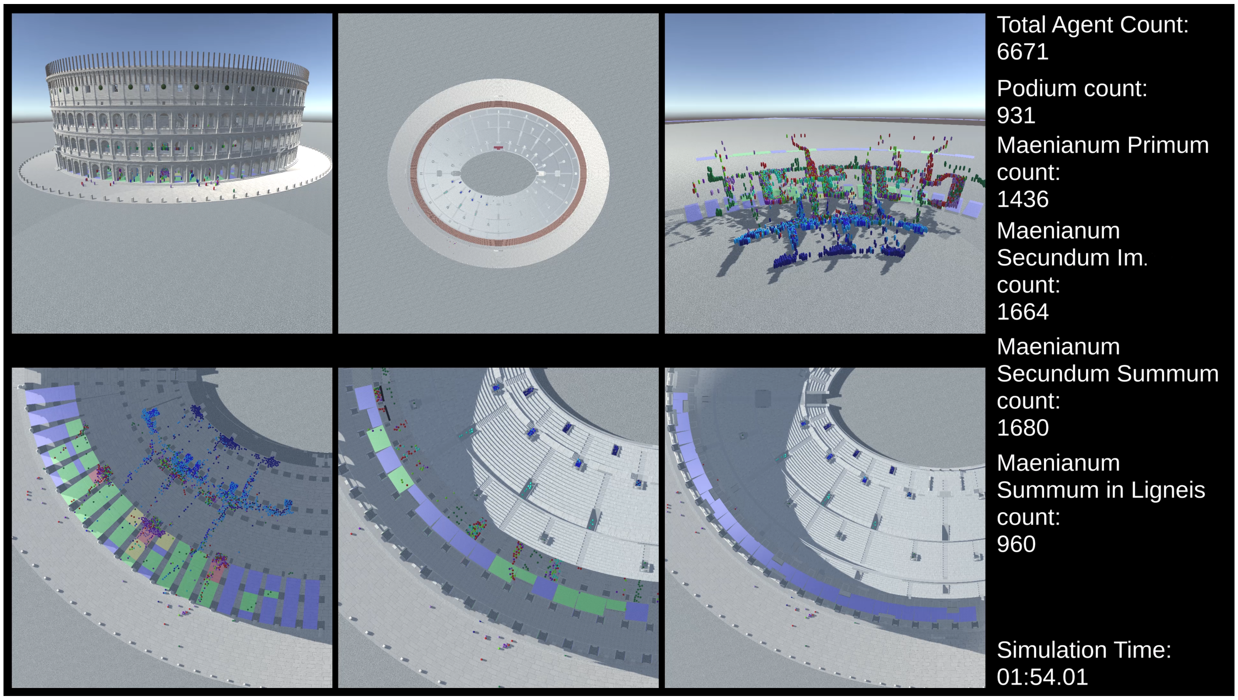

Total Agent Count:
6671
Podium count:
931
Maenianum Primum count:
1436
Maenianum Secundum Im. count:
1664
Maenianum Secundum Summum count:
1680
Maenianum Summum in Ligneis count:
960
Simulation Time:
01:54.01

Figure 3. Example of an extracted frame from the simulation—specifically run four, take three (more fully described below). Images from the top row represent (from left to right): an exterior perspective view, an orthographic top view, an interior perspective view with the mesh of the Colosseum model hidden allowing the viewer to see into the heart of the structure. Images from the bottom row represent (from left to right): an orthographic top view with the camera clipping the Colosseum model to reveal the ground-floor with heat maps, an orthographic top view with the camera clipping the Colosseum model to reveal the first-floor mezzanine with heat maps, an orthographic top view with the camera clipping the Colosseum model to reveal the second-floor mezzanine with heat maps. Note that the virtual agents are color coded to represent the level of the vomitorium to which they are navigating. Also note this image has been modified to correct a spelling error. Reconstruction is courtesy of and copyright by Flyover Zone, Inc., 2024. All rights reserved.

3.2 Simulating Navigational Priority

With the foundation of the simulation in place, it was ready to serve as a testbed for experimentation. However, before beginning this discussion, it is important to understand how the experimental variables were developed. These emerged through iterative processes; as the simulation grew in complexity, the questions it was best suited to address arose organically from the work itself. This process gave rise to two primary questions. First, how would manipulating the "priority" agent-property change the behavior of the crowd (and could this be done with an eye toward historical reality). Second, how would changes to the instantiation rate of the agents and the number of agents per gallery change the crowd density throughout the internal structure of the Colosseum.

In developing a strategy for exploring these questions, an important point to note to return to is that the simulation took as its starting point the simulation developed by [Gutierrez, et al. 2007] whose study instantiated 7669 agents. As such, the first simulated runs of the environment attempted to closely match the scenario established in this earlier work. In the present simulation, the agents were divided into 32 groups, each of which were tasked with navigating to a different vomitoria within the quarter of the structure under study. These agents were then replicated and arranged in arcing phalanxes around the outer arcade of the Colosseum in a manner consistent with figure 11 from [Gutierrez, et al. 2007]. This served as each agent's starting position. This resulted in 7680 agents, all of whom would awaken at the start of the simulation and navigate through the crowds to their various destinations.

Dividing the agents into 32 groups ensured that the same number of agents would pass through each vomitorium and only slightly changed the total number of agents instantiated here relative to [Gutierrez, et al. 2007]. However, unlike in [Gutierrez, et al. 2007], the agents were not tasked with moving to their specific seats. This decision was made for several reasons, the most salient being that it would have added unnecessary complexity to the research. The primary focus was on the navigation of agents within the bowels of the structure (i.e. the majority of the navigation distance), rather than within the cavea. By concentrating on this particular space, the complex social nuances involved in navigating through a seating gallery were avoided as was the generation of a appropriately high-quality navmesh in the *cavea*. This made the simulation overall less complex to construct and interpret. Additionally, to enhance the reliability of the simulation results, data was collected three times. This occurred each time the simulation parameters were changed. These simulation repetitions will hereafter be referred to as *takes*. Under these experimental conditions, four runs were conducted, resulting in a total of twelve takes.

Having organized the simulation starting parameters in a manner consistent with [Gutierrez, et al. 2007] and having identified the pertinent questions that could be simulated, it was determined that manipulating the priority value-ranges of the agents would be useful for simulating social stratification within the crowd. The first run kept all of the agents' priority property fields equal, i.e. all set to the same integer value range, since this was not a property varied in [Gutierrez, et al. 2007]. This also provided a controlled situation concerning the stratification of the crowd where an agent had a random chance to encounter an agent of either lower or higher priority. The results of this simulation run demonstrate that an unstratified crowd, operating under these experimental conditions, generally takes nine to eleven minutes for every member of a given gallery to reach their appropriate vomitorium (i.e., "Time of Last Agent Arrival"); this excludes those agents navigating to the *podium* who by the nature of the design of the Colosseum have a naturally fast route to their intended destination (tab. 1).

Table 1. Agent Temporal Data, Run 1

Take	Gallery	Time of First Agent Arrival	Time of Last Agent Arrival (Only Successful Agents)	Total Agents in Group	Number of Agents Halted	Percent Halted
1	P	00:42.02	05:06.02	1680	0	0.000%
1	MP	00:48.02	10:10.01	1680	0	0.000%
1	MSI	01:18.02	11:08.04	1680	0	0.000%
1	MSS	02:16.01	11:18.05	1680	1	0.060%
1	MSL	02:34.02	10:46.03	960	1	0.104%
2	P	00:46.02	05:18.01	1680	0	0.000%
2	MP	00:48.02	09:11.99	1680	0	0.000%
2	MSI	01:18.02	10:18.02	1680	1	0.060%
2	MSS	02:18.01	10:18.02	1680	2	0.119%
2	MSL	02.40.02	10:28.02	960	0	0.000%
3	P	00:42.02	05:18.01	1680	0	0.000%
3	MP	00:48.02	10:08.01	1680	0	0.000%
3	MSI	01:22.02	09:50.00	1680	0	0.000%
3	MSS	02:10.01	10:34.03	1680	0	0.000%
3	MSL	02:36.02	10:58.04	960	0	0.000%

Note the abbreviations for the galleries follow Latin naming conventions for tables 1-5. Times listed as hours (where applicable):minutes:seconds.milliseconds for tables 1-3, 5.

After completing this run, it was decided to change the priority value-range of every agent in a way that reflects the class and gender stratification of an ancient Roman crowd in three different ways. The goal was to determine if including a class/gender factor in the functioning of the agents would produce a situation different from a randomized crowd. If successful, it would demonstrate that not only the architecture of the Colosseum stratified the crowd, but also the crowd's own pre-existing social organization. This was done based on the assumption that the class of a spectator directly correlated with their seating assignment, as discussed in the historical review above. On account of this, three additional runs of the simulation were conducted that: 1) divided those navigating to the *podium* from the rest of the crowd, 2) divided the agents by gallery, and 3) divided the agents by galleries except for those navigating to the *maenianum summum in ligneis* who were instead divided into four equal groups which were then sequentially assigned the one of the priority values as the other four galleries.

The first of these additional scenarios differentiated the crowd by assigning virtual Senators—agents navigating to the *podium*—a higher priority value range, while all other agents were assigned a uniform, lower range. This was intended to reflect a similar but slightly different aspect of the [Gutierrez, et al. 2007] simulation wherein the Senators were provided with access into the structure via particular gates. Though special access paths were not provided in this simulation, it was speculated that providing the Senators a more robust way to push through the crowd would nonetheless change their ability to enter the structure. The results do indeed demonstrate that this increased priority value-range allowed these agents to enter more quickly and only slightly increased the entrance time for the rest of the crowd (tab. 2).

Table 2. Agent Temporal Data, Run 2

Take	Gallery	Time of First Agent Arrival	Time of Last Agent Arrival (Only Successful Agents)	Total Agents in Group	Number of Agents Halted	Percent Halted
1	P	00:42.02	03:40.02	1680	0	0.000%
1	MP	00:48.02	09:44.00	1680	0	0.000%
1	MSI	01:24.02	10:16.02	1680	0	0.000%
1	MSS	02:10.01	10:02.01	1680	3	0.179%
1	MSL	02:44.02	10.08.01	960	0	0.000%
2	P	00:44.02	03:44.02	1680	0	0.000%
2	MP	00:46.02	09:48.00	1680	0	0.000%
2	MSI	01:18.02	11:20.05	1680	0	0.000%
2	MSS	02:18.01	11:44.06	1680	0	0.000%
2	MSL	02:38.02	12:04.07	960	0	0.000%
3	P	00:42.02	03:38.02	1680	0	0.000%
3	MP	00:50.02	09:21.99	1680	0	0.000%
3	MSI	01:20.02	11:02.04	1680	0	0.000%
3	MSS	02:28.01	11:18.05	1680	0	0.000%
3	MSL	02:38.02	11:48.06	960	0	0.000%

The second of these additional scenarios provided different priority value-ranges based on the gallery to which an agent was navigating; by proxy this also captures the class and gender of the spectators. For example, agents identified as men belonging to the togate Plebian class, i.e. the agents navigating to the *meanianum secundum immum*, were assigned a random value between thirty and forty, while agents identified as women who were navigating to the *maenianum summum in ligneis* were assigned a random value between fifty and sixty (tab. 3).

The last scenario followed the same logic as above with one additional change. Instead of placing all agents navigating to the *maenianum summum in ligneis* (i.e., women) into the lowest value-range group to reflect their assignment to the highest gallery of the structure, they were distributed among the four higher priority groups (tab. 3). This was achieved by dividing these agents vertically, meaning each agent was assigned to one of the four vomitoria in this gallery. Consequently, each of these four sub-groups of women was allocated value-ranges corresponding to those assigned to men. The logic for this being that agents classified as women might be better modelled with the class divisions of their families (i.e. the social rank of the men in their families); it is reasonable to imagine that the wife of a Senator could use their higher social station to move through a crowd of lower ranking men or women. Thus, the agents identified as women were divided from the men, based on their gallery, and divided from each other, based on class assignment/vomitorium location. While there is very little historical information about the division of this uppermost seating gallery, it is easy enough to imagine a scenario where it was divided vertically in a manner similar to the stone seats further down the *cavea*. In such a case, one might expect intra-gallery partitions and that particular vomitoria would serve just one group/class of women [Bomgardner 2021].

The results of these two runs demonstrate that a crowd operating under these highly stratified experimental conditions generates more hierarchical temporal data (tab. 4). Also, the crowd is generally more efficient at reaching their desired destination (more fully detailed in §4.1). Between these two runs, however, there is an interesting emergent property. Namely, when the agents identified as women are themselves divided by class, the agents navigating to the *maenianum secundum immum/summum* (i.e., men ranked as plebians) take slightly longer to reach their desired vomitorium. Additionally, under these conditions, those navigating to the *meanianum secundum summum* (i.e., low-class plebians, etc.) and the *maenianum summum in ligneis* (i.e., women) take nearly the same amount of time—barring the outlier in run four, take three. This stands in contrast with run three wherein these two groups are consistently distinguished by nearly a minute.

Table 3. Priority Assignments for Runs 3 and 4

Run	**Gallery**	**Priority Value-Range**
3	P	10-20
3	MP	20-30
3	MSI	30-40
3	MSS	40-50
3	MSL	50-60
4	P	10-20
4	MP	20-30
4	MSI	30-40
4	MSS	40-50
4	MSL	10-20, 20-30, 30-40, 40-50

Table 4. Agent Temporal Data, Runs 3 and 4

Run	Take	Gallery	Time of First Agent Arrival	Time of Last Agent Arrival (Only Successful Agents)	Total Agents in Group	Number of Agents Halted	Percent Halted
3	1	P	00:42.02	03:48.02	1680	0	0.000%
3	1	MP	00:48.02	04:54.02	1680	0	0.000%
3	1	MSI	01:20.02	06:00.00	1680	0	0.000%
3	1	MSS	02:24.01	09:48.00	1680	0	0.000%
3	1	MSL	03:04.02	10:34.03	960	0	0.000%
3	2	P	00:44.02	03:34.02	1680	0	0.000%
3	2	MP	00:48.02	05:18.01	1680	0	0.000%
3	2	MSI	01:26.01	06:00.00	1680	0	0.000%
3	2	MSS	02:18.01	09:29.99	1680	0	0.000%
3	2	MSL	02:54.02	10:18.02	960	0	0.000%
3	3	P	00:44.02	03:50.02	1680	0	0.000%
3	3	MP	00:46.02	05:16.01	1680	0	0.000%
3	3	MSI	01:20.02	06:00.00	1680	0	0.000%
3	3	MSS	02:24.01	09:09.98	1680	0	0.000%
3	3	MSL	02:50.02	10:36.03	960	0	0.000%
4	1	P	00:42.02	03:52.03	1680	0	0.000%
4	1	MP	00:46.02	05:06.02	1680	0	0.000%
4	1	MSI	01:22.02	06:30.00	1680	0	0.000%
4	1	MSS	02:28.01	09:40.00	1680	2	0.119%
4	1	MSL	02:30.02	09:52.00	960	22	2.292%
4	2	P	00:44.02	03:32.02	1680	0	0.000%
4	2	MP	00:48.02	04:58.02	1680	0	0.000%
4	2	MSI	01:28.01	07:43.98	1680	0	0.000%
4	2	MSS	02:28.01	09:46.00	1680	0	0.000%
4	2	MSL	02:28.01	09:48.00	960	0	0.000%
4	3	P	00:42.02	03:30.02	1680	0	0.000%
4	3	MP	00:50.02	05:06.02	1680	0	0.000%
4	3	MSI	01:26.01	06:26.00	1680	0	0.000%
4	3	MSS	02:24.01	11:14.04	1680	11	0.655%
4	3	MSL	02:36.02	09:42.00	960	0	0.000%

3.3 Extending the Simulation

Following this set of four runs, three additional simulation parameters were identified as potentially worthy of further study. First, the specific number of agents navigating toward a given vomitorium entrance was, in the previous runs, equal to the total number of vomitoria entrances located in a given gallery. This was useful for generating a scenario wherein each vomitorium entrance would be handling the same number of agents during runtime. However, this might not be the most realistic scenario. The number of spectators navigating to a particular vomitorium can be more accurately simulated by considering the total number of spectators seated in each gallery and distributing them among the vomitoria in that gallery.

Peter Rose's research provides information on the total number of spectators in each gallery [Rose 2005]. He calculated this by comparing the possible seating area of the gallery with various seat sizes. For this particular simulation, his calculation for 0.5 by 0.8 meter seats was selected which resulted in the following number of spectators per gallery: 4100 (*Podium*), 9300 (*Maenianum Primum*), 16200 (*Maenianum Secundum Immum*), 8000 (*Maenianum Secundum Summum*), and 8100 (*Maenianum Summum in Ligneis*).

Two other parameters were related to the instantiation of agents into the simulation. First, it became obvious that the instantiation of all the agents into the simulation at once was probably an unlikely imitation of reality; every member of the crowd was probably not waiting directly outside for the gates to open. This is substantiated by modern guidelines for stadium management which recommend one hour as the smallest reasonable timeframe to consider for the filling of such a space [Sports Grounds Safety Authority 2024]. The second of these parameters had to do with the directionality of the entrance of the agents. Of course, considering the ancient urban fabric within which the Colosseum was embedded, agents headed to this particular quarter of the structure would probably not only enter from directly adjacent districts. Rather, one would expect spectators to flow around all sides of the Colosseum and from every nearby district of the city, only entering the building once they arrived close to their gate. One might also expect that they would not be able to enter the structure before arriving to their intended quadrant since the major and minor axes of the structure are generally believed to have been special-use gates and were likely impassable by the crowd, at large [Hopkins and Beard 2005].

Having identified these three parameters (seating capacity, time of instantiation, location of instantiation) from the previous simulation conditions, a fifth run of the simulation was developed. This run took as its starting point the third simulation run by implementing priority value-ranges divided only by seating gallery. During that run, the agents cleared the simulation, on average, the fastest and resulted in the fewest number of agents being halted by chance obstacles (tab. 4); i.e., this was the most efficient and stable parameter set. Next, the agents were removed from the area directly outside the gates and instead were set up so that they would be instantiated at a short distance away and would be instantiated into the scene from a random location all around the structure. This instantiation would also be prolonged over the course of the recommended one hour as per [Sports Grounds Safety Authority 2024]; their numbers were equally distributed over this time period.

These gallery capacity numbers were derived from those calculated by Rose, quartered (since only a quarter of the Colosseum was going to be simulated as with the earlier runs), and then distributed across the appropriate number of vomitoria per gallery. Note also that since distributing these numbers across the appropriate number of vomitoria resulted in some non-integers, these figures were normalized to result in whole integer values since fractions of an agent are, of course, not possible to simulate. This resulted in a less than two percent deviation from Rose's calculated values.

Table 5. Agent Temporal Data, Run 5

Take	Gallery	Time of Last Agent Arrival (Only Successful Agents)	Total Agents in Group	Number of Agents Halted	Percent Halted
1	P	01:05:46.92	1008	0	0%
1	MP	01:07:11.08	2324	0	0%
1	MSI	01:07:29.12	4032	0	0%
1	MSS	01:08:17.20	1988	0	0%
1	MSL	01:07:55.17	2016	0	0%
2	P	01:06:14.97	1008	0	0%
2	MP	01:05:10.85	2324	0	0%
2	MSI	01:07:33.12	4032	0	0%
2	MSS	01:08:39.12	1988	0	0%
2	MSL	01:09:28.92	2016	0	0%
3	P	01:04:52.81	1008	0	0%
3	MP	01:06:26.99	2324	0	0%
3	MSI	01:08:15.20	4032	0	0%
3	MSS	01:09:05.01	1988	0	0%
3	MSL	01:08:01.18	2016	0	0%

The goal of this simulation run was to understand how the flow of the crowd would change under these more realistic spatial and temporal conditions. After running this simulation, it was clear that the elongated temporal period during which the spectators entered the structure dramatically reduced the crowds throughout the space and in fact led to approximately no more than one thousand agents being present in the scene at any time. Another interesting behavior to note was that, when the agents did form substantial crowds, they formed at the gates nearer the major and minor axes of the quarter under study. This could indicate, as in [Gutierrez, et al. 2007], that specific gate entrances may have been assigned to various gallery or section assignments to forcibly even out the crowd. This seemed to be because most of the agents, having been instantiated in areas other than those directly in front of this quarter, navigated around the outer arc of the Colosseum and entered quickly upon reaching the first set of open gates of the quarter under study. However, even in these areas, the crowd density as recorded by the heat map indicated that the density of the crowd never exceeded two agents per meter. No take of this run of the simulation saw agents unnaturally halted. This more sparce distribution of virtual agents also had an impact on the class distribution in comparison to previous runs. Under these conditions, each group of agents cleared the interior of the structure within 4.6 minutes of each other across all takes—an approximately 35% smaller spread than the average of the takes from run three (tab. 5).

3.4 Visualizing Bottlenecks

In addition to the temporal data generated above, information about crowd density was generated using heat maps. As briefly mentioned above, these were arranged in the concourses as noted by [Gutierrez, et al. 2007] to exhibit bottlenecks, namely, the first- and second-floor mezzanines. Heat maps were also added to the exterior-facing ground-floor arcades, since this is the one place where all the agents are forced to go through. As previously mentioned, a series of cameras and agent-counting fields recorded the state of the simulation during runtime and simultaneously captured the state of the heat maps (Fig. 1). Recording information about the emergence of bottlenecks within the structure is important, as it helps identify the specific locations and times where agents were delayed. This delay is, in turn, useful for understanding the experience of an ancient spectator who may have found themselves stuck in a given location, waiting for a bottleneck to clear. For example, some agents of a particular gallery may be forced to linger in cramped surrounds more than others.

The heat maps utilized in this simulation are helpful for locating bottlenecks because they generate visual feedback about the crowd density for a given location. Crowd density is estimated based on a modified version of [Sports Grounds Safety Authority 2024] guidelines which recommend "…that the available floor space of the concourse should be able to accommodate [spectators] at a density level of no more than 2 persons per square meter…"; it is their opinion that two persons or less per square meter is the "optimum density for general concourse areas." This recommendation was slightly manipulated to fit the current circumstance, where heat maps signaling blue represent zero agents, green represent fewer than 1 agent per meter, yellow represent between 1 and 2 agents per meter, and red represent more than 2 agents per meter. The placement of the heatmaps was done to closely follow the architecture of the areas under question as much as possible. While these were placed regularly, it was important to transform their shape to match the nuances of the architecture so that each heat map panel would appropriately cover a discrete location in a concourse. Nonetheless, while these panels are not of the same size, they all use the same threshold values for crowd density.

In the present study, this information was mostly used anecdotally to visually demonstrate the difference between each of the experimental runs and to inform developmental decisions. However, in future phases of this work it is intended to be more fully analyzed for its potential to generate more nuanced density data, especially over time. Just such a preliminary analysis was conducted for run four, take three. Here, the state of each segment of the heat map was simply assigned a numerical value between zero (blue, zero crowd density) and three (red, maximum crowd density). Assigning these values allowed for the aggregation of crowd density for, in this case, the first-floor mezzanine (Fig. 4-5). Visualizing the data in this way produces a very readable mapping of the crowd density in this location of the Colosseum simulation and thus the formation of bottlenecks that under these conditions slow the crowd's passage through this space. However, I must emphasize that this is just one of many possible visualizations. Others may prove to be more useful as this aspect of the study evolves in future research phases.

Visualizing crowd pressure data in this and other related ways is anticipated to be a productive method for examining bottlenecks in simulations like the one here. In fact, visualizing crowd density and pedestrian flow is standard practice for research in this area, however, there are no standards in place for its application in CH [Duives et al. 2015; Gravit, et al. 2022; Moussaid et al. 2018]. Therefore, it is crucial for future research to focus on developing standard tools and techniques that are most useful for visualizing data in this field.

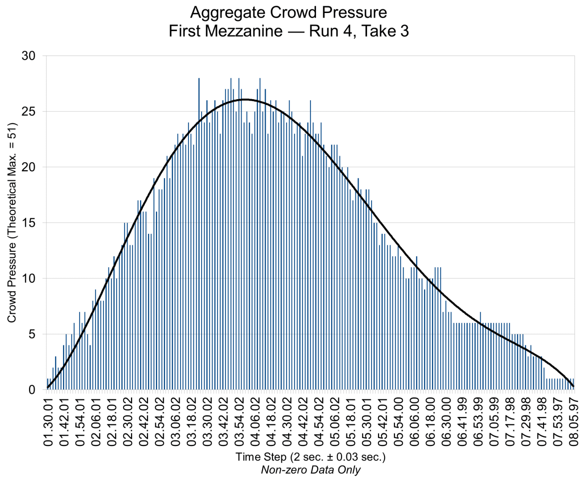

Figure 4. Chart showing the aggregate crowd pressure seen in the first mezzanine of the Colosseum simulation during run four, take three. Note this is the same data as figure 5.

Figure 5. Color-gradient graph showing the aggregate crowd pressure seen in the first mezzanine of the Colosseum simulation during run four, take three. Note this is the same data as figure 4.

4. DISCUSSION

4.1 Interpretation of Results

Taken together, the results from the five simulation runs presented above demonstrate the importance of considering more than just spatial parameters when simulating an ancient crowd. These parameters should also include, at minimum, both temporal factors and the social identity of spectators. The results of the simulations provide significant insights into the behavior of crowds under varying conditions of stratification and density.

The first set of simulations, consisting of runs one through four, involved nearly 8,000 agents. These simulations demonstrated key differences in the "Last Agent Arrival Time" based on the stratification of the crowd by the class and gender identity assigned to the virtual agents via the priority agent property. This piece of information is particularly interesting because it shows how long it took every member of a given gallery to make their way to their destination. When the crowd was not stratified, agents showed little variation in their “Last Agent Arrival Time” with the exception of agents navigating to the lowest level vomitoria. This resulted in a rather homogeneous crowd experience, with the overwhelming majority of agents arriving to their seats within a few dozen seconds of each other. In contrast, when the crowd was stratified, the “Last Agent Arrival Time” varied much more noticeably among the galleries. This stratification mirrored the expected social hierarchy with the ancient space, indicating that social structures within the crowd influenced movement patterns. The results highlighted that while certain aspects of the design of the structure, like the seating galleries, exhibit stratifying patterns, the interior components, such as the concourses, did not significantly impact stratification. This finding implies that the internal architecture of the structure plays a lesser role in influencing crowd stratification compared to the social hierarchy imposed on the agents based on their ticket. This is to say, the social experience of a visitor to the Colosseum is more impacted by their pre-existing social class/gender versus the architectural design of this space, though architecture is clearly a reinforcing factor.

Another interesting aspect of the results of this simulation is that when the agents used their greater social advantage (or were disadvantaged by this) the crowds as a whole navigated to their intended vomitorium with greater ease. This tended to shorten the “Last Agent Arrival Time” of the entire crowd on average (tab. 6). If one considers that navigational conditions that improve efficiency are more likely to reflect real-world scenarios, then these results indicate that runs with stratified crowds are more likely to reflect the ancient reality, since simulations with greater social stratification showed a lower average "Last Agent Arrival Time."

Table 6. Average Time Across Each Gallery for First and Last Agent Arrival (in seconds)

Run	Average Time for First Agent of a Gallery to Arrive	Average Time for Last Agent of a Gallery to Arrive
1	92.42	563.35
2	94.02	558.56
3	98.15	418.41
4	94.82	427.08
5	—	4041.84

In the second phase of the simulations, i.e. run five, the concept of stratification was applied to a more realistic scenario where virtual agents were observed over a one-hour timeframe and with more historically accurate crowd numbers. This phase revealed that in lower density conditions, the crowds tended to become less stratified. This behavior can be attributed to the reduced interactions among agents, which in turn diminishes the opportunities for agents to exert social dominance over one another. The findings here suggest that social stratification is more pronounced under high-density conditions where frequent interactions among agents reinforce hierarchical organization. It also shows that the long instantiation period largely overrode the social factors implemented in this scenario.

These results reinforce that crowd stratification and social hierarchy are dynamically influenced by crowd density and pre-existing social structures. In high-density scenarios, increased interaction among agents leads to a clearer manifestation of historical social stratification, aligning with expected behaviors based on their assigned priority values. Conversely, in lower-density scenarios, less frequent interactions reduce the impact of social hierarchies, resulting in a more egalitarian movement pattern that only superficially considers class priority in relation to movement. These findings underscore the importance of considering the social identity of an agent as much as the reconstructed environment of a simulation when analyzing crowd behavior in cultural heritage simulations. The stratification observed in high-density conditions reveals the significant influence of social hierarchies on movement patterns, while the more egalitarian nature of lower-density crowds highlights how factors like the period of time designated for agent instantiation can partially dissolve social hierarchies.

4.2 Technical Limitations

When interpreting a simulation such as the one presented here, it is crucial to understand its limitations. Digital simulations, while powerful tools for generating otherwise non-observable information, can often be perceived as overly authoritative due to their computational origins [Richards-Rissetto 2022]. This perception is inaccurate as these models have inherent constraints and should be understood within their specific contexts. As such, when thinking about the limitations of this simulation of the ancient Colosseum, there are three technical limitations that warrant discussion.

The first limitation lies in the imprecision of the pathfinding system. Using the A* algorithm for pathfinding inherently relies upon the generation of a high quality navmesh. The navmesh, as described above, was not perfectly complete due to the complex structure it was applied to, as briefly discussed in §3.1. The primary issue arose from the types of polygons used to discretize the mesh during this automatic, procedural process—i.e., squares and triangles. Although generally effective, these shapes cause problems when they are generated at a density necessary for an environment as complex as the Colosseum. The square in particular created a situation wherein paths that are axial (i.e., paths that move from the center of a given square to the midpoint of an edge) were prioritized over diagonal paths (i.e., paths that move from the center of a given square to a corner) that would obviously be globally faster. This is likely because the navigation system does not step out far enough in terms of its calculations of least-cost to see the shorter diagonal path.

This issue was mitigated by using parameters to create a less dense navmesh, but this adjustment came at the expense of some structural granularity; for example, this led to less mesh in the *cavea* and irregular mesh on the stairs of the mezzanines. While this lack of granularity was overcome, were this simulation expanded to examine navigation in the *cavea* an improved process for generating the navmesh would be required. Ultimately, the odd pathfinding behavior derived from the navmesh may have slightly lengthened some agents' paths, however, most of these odd paths were removed from the simulation via the change in navmesh density.

The second limitation involves improper collisions resulting from the default implementation of the RVO algorithm, which in its implementation within Unity can only consider eight nearby neighbors at a time [Unity 2024]. While typically not a problem, this restriction sometimes allowed agents to collide and improperly occupy each other's space. This was particularly obvious when there were high density groups of agents, e.g., around the staircases leading between the ground- and first-floors. While this was partially addressed by increasing the agent collider radius, this solution is less systematic than would be hoped. Moreover, in the highest density crowds, some improper collisions were still noticed. This improper motion does not invalidate the results of this simulation since the results by-and-large agree with previous studies' results which show that all spectators could navigate through the structure and clear it in approximately 12-20 minutes. Nonetheless, it could mean that the current simulation results may have slightly shortened the overall navigation time, since some agents did not have to wait for their path to clear and instead were pushed through the crowd in a non-realistic manner.

The third limitation pertains to the replication of social pressures by means of the priority agent property. In the simulation, priority often translates to pushing behavior, which might not accurately reflect reality. In a real-life scenario, individuals of higher social status, such as senators, would likely be allowed to pass through a crowd on sight rather than having to push through it. This discrepancy highlights a gap between the simulated behaviors and real-world social behaviors, suggesting that the model may oversimplify the interactions within a crowd. One can imagine that agents of lower priority should have slowed prematurely to leave a clear path for a higher priority agent. Perhaps this would have further slowed the crowd. However, implementing such a change would require a very sophisticated algorithm (or set of algorithms) to control social factors and visual cues for each agent. Likely, this would lead to much higher computational costs with fewer benefits than would be hoped for with such an effort.

Some of these limitations arise from constructing the simulation using Unity and its default AI navigation system. Although this system is a limiting factor, its relative ease of use facilitates verification through repetition by the research community. This advantage is due to the plug-and-play nature of the tools and their implementations of openly accessible algorithms, which are published and robust for such navigation tasks. By using these tools and refining their use in the future, this work aims to enhance the repeatability of agent-based simulations, addressing a significant limitation of data so generated for cultural heritage practitioners [Gietl et al. 2007; Herzog 2020].

Understanding these limitations is essential for accurately interpreting the results and for guiding future improvements in simulation methodologies. Despite these constraints, the simulations offer valuable insights into crowd behavior and social stratification, providing a foundation for further research and refinement. Moreover, perfectly realistic human motion is exceptionally messy and not perfectly simulable [Amos and Webster 2022]. Even if it were, it may not be the most useful for research in this area since such a simulation may become too complex to distill its most important parameters. Without the ability to abstract, results would become equally messy and ultimately less useful—if useful at all.

4.3 Implications for Practitioners

The question that any interested practitioner should have at this point is, how can a simulation such as this be used for research given the limitations outlined above. An answer to this is that it can be understood as a kind of visualization—not just a way of displaying data, but a complex process that also generates new data in the act of visualizing [Guidi and Frischer 2020]. Any visualization is ultimately a product of data curation, or in the language of agent-based modelling, system bounding [Lake 2020]. As long as this curation is done properly and transparently, as per guidelines like the Seville Principles for virtual archaeology, it can produce a very powerful interpretation of a given data set that may be impossible to otherwise interpret from systems which are too complex [Brughmans et al. 2019; Icomos 2017]. As for the Colosseum simulation presented here, it too curates the data that is being presented and, in this case, also generates additional information from that curation.

A clear example of this curation at work is that this simulation assumes that the crowd will have a variation in speed ±25% of an average walking speed of 1.3 m/s. While a seemingly innocuous simplification, there is the possibility that it misses nuances in relation to the composition of the crowd. One could imagine that the richest members of society who are likely to be healthier would be on average faster walkers, i.e., there could be class divisions at play. This is in fact a very reasonable piece of information to consider; differences in ability are a notable factor in modern crowd management [Geoerg et al. 2022]. However, the question is, how exactly would one then find out the appropriate values for an ancient context such as the Colosseum. Ancient pedestrian behaviors are, of course, not directly recoverable. So, the only way forward is to turn to modern examples. Average walking speed is one thing, but trying to convert modern information about class and health to its ancient Roman counterparts would be so complex as to be meaningless or worse misleading. This serves as a salient example of the simulation's utility in abstracting the multifaceted nature of the problem and distilling it into a more comprehensible form. This curated simplicity is exactly what drives its usefulness and potential as a tool for research in cultural heritage.

It is also essential to recognize that this simulation, like any, reflects its inputs and the biases of its designers. For example, this simulation can be biased by conforming to standard interpretations within the field. It assumes the *podium* in the Colosseum was exclusively for Senators, aligning with conventional understanding. Such biases are generally not problematic and serve as useful assumptions for generating meaningful results. This bias, nonetheless, highlights the nature of the simulation: it is a tool for generating results and questions for further investigation by building on existing scholarship in an interdisciplinary way. Rather than providing definitive solutions, simulations like this should be seen as explorations of potential hypotheses that arise during other research phases. This perspective allows them to be evaluated for their accuracy and used to drive further development.

Trustworthiness is also key for utilizing simulations like the one presented here. Trusting a simulation of this nature is justified due to its faithful reconstruction of the archaeology and architecture involved [Guidi and Frischer 2020]. The use of well-established least-cost path algorithms for navigation, of which the A* algorithm is an example, is already widely employed in agent-based modeling (albeit primarily in two-dimensional contexts) and reinforces this justification [Herzog 2020]. Additionally, simulating human behavior using independently operating agents, as done here, is a valid method for understanding group behaviors, even though it may be less precise for predicting the actions of specific individuals [Sports Grounds Safety Authority 2024; Warren 2018]. Nevertheless, the social characteristics inherent to the individuals that comprise crowds has been shown to influence the navigational behaviors of modern crowds [Larsson et al. 2021]. Thus, simulating a crowd as demonstrated here can be considered a realistic reflection of observed pedestrian behavior.

Cultural heritage practitioners can also use the results of this simulation to inform their own research. These simulations should be considered as part of a broader constellation of data points supporting a spatially focused argument. For example, the results of this simulation could support a hypothesis suggesting that the social hierarchies inherent in Roman society are reflected not only in the architecture of structures like the Colosseum but also in the formation and interactions of crowds within these spaces. To illustrate, as has briefly been done here, one might combine textual evidence about the legal organization of the space with the idea that the Colosseum was a venue where one's social standing could be displayed and reinforced. The simulation results, which indicate that social stratification was beneficial to the crowd during high-density events, can then be integrated into this broader argument. This approach can allow practitioners to contextualize the simulation findings within the larger framework of class and gender dynamics, shedding light on how these factors influenced public events in ancient Roman society.

One might also iterate on this simulation to understand how the Colosseum functioned during particular events, for example, Titus' 100-day inaugural games [Hopkins and Beard 2005]. In a case such as this, one might expand the simulation parameters over time. This would encompass things like the fact that spectators would be expected to filter in and out of the structure in unequal intervals, creating the need for two-lane navigation throughout the concourses. Another particular event that could be examined is the peculiar incidences when the emperor Commodus was engaged in the fighting himself. In this case, the higher social classes were very much expected to be present, while, as per the historian Cassius Dio, "…common people… had more choice than their betters and were much more inclined to give the proceedings a miss…" [Hopkins and Beard 2005]. Here, the simulation could be reset so as to skew the gallery numbers to match just such a scenario.

By so augmenting the simulation, practitioners can gain insights into the practical implications of the social organization in the Colosseum to develop new questions and explore whether particular hypotheses might be fruitful for further research. Importantly, this work builds upon prior digitally-based research that has examined social hierarchies within the context of cultural heritage research [Gutierrez, et al. 2007; Maïm, et al. 2007]. Moreover, by using Unity as the platform for the simulation, making changes on the fly suited to the particular needs of cultural heritage simulations can be easily accomplished.

In sum, the use of such simulations by cultural heritage practitioners should be multifaceted and holistic with respect to available evidence and to the intended hypotheses being tested. This can contribute to a greater understanding of the historical context surrounding the ancient Colosseum with respect to the human experience within the space. By situating this simulation's results within a wider array of evidence and interpretations, practitioners can enhance their insights into the functioning of this spectacular space and perhaps even into the experiences of its spectators. An integrated approach such as this can ensure that the simulation is employed as a powerful tool for hypothesis testing and theory building rather than as isolated or conclusive evidence.

5. CONCLUSIONS

The study presented here demonstrates the valuable insights that digital simulations can offer in understanding crowd behavior and social stratification in the ancient Colosseum. Despite certain limitations related to pathfinding algorithms, collision detection, and the simplification of social interactions, this simulation remains a robust tool for approximating real-world, otherwise inaccessible, behaviors. By using a high-fidelity reconstruction of the architecture and employing tried and tested algorithms, the simulation provides a faithful representation of how social hierarchies likely influenced crowd movements within the Colosseum. Additionally, it has provided guidance for cultural heritage practitioners to leverage simulations setups such as the one presented here not as a definitive problem-solver, but as part of a broader analytical framework that includes textual evidence, archaeological evidence, and other data points. The ability to visualize and explore the impact of social structures on crowd behavior enhances our understanding of historical spaces and their use. Moreover, it offers significant value as a highly readable visualization, making complex social and architectural interactions more easily appreciable, while maintaining its essential character as a 3D agent-based model.

Future stages of this research intend to explore three items of interest. First, there will be an examination of the possibility of implementing new AI navigation systems. Specifically, it is important to be able to explore the possibility of both using newer versions of the RVO algorithm which are proven to be less likely to allow collisions between agents and also using different polygon geometries for the navmesh which may allow for more believable paths. Second, it is important to implement a more granular system for the generation of heat maps within the concourses of the Colosseum. While the system currently in place is useful, something more along the lines of that used in [Gravit, et al. 2022] may be more telling of the scale, shape, and duration of bottlenecks. Finally, understanding how to apply this simulation methodology to other digital reconstructions will be an important step toward making it a more generalizable tool for cultural heritage practitioners and practitioners from allied disciplines. Especially important will be understanding how to interpret the data it generates from reconstructions that have less sure footing than the one representing the Colosseum as used within this study, for instance models such as those found in [Micoli et al. 2024].

While there are areas for improvement and many paths for future research, this simulation provides a reasonable and insightful approximation of historical crowd behavior. It is hoped that this study, and its implementation of an essentially free-to-use simulation platform, might act as a blueprint for the construction of more agent-based 3D simulations in the domain of cultural heritage. Ultimately, their integration into this research area can allow for a more nuanced and comprehensive understanding of ancient societies, reinforcing the importance of interdisciplinary approaches in analyzing the past.

6. ACKNOWLEDGEMENTS

The author wishes to thank Dr. Gabriele Guidi, who provided the initial inspiration for this work by proposing the use of 3D agents to address some unresolved issues in the 3D reconstruction of the Roman Circus of Milan [Micoli, et al. 2024]. Special thanks to Dr. Bernard Frischer for his ongoing guidance and his prompting to examine and expand upon his collaborative work from the early 2000s, which explored the use of virtual agents for ergonomic studies in the Colosseum [Gutierrez, et al. 2007]. Additional thanks to Dr. Julie Van Voorhis, Dr. Samantha Wood, and two anonymous reviewers whose observations have improved this work. The author wishes to thank Flyover Zone, Inc. for the gratis use of the Colosseum digital reconstruction model within the simulation presented here. The author gratefully acknowledges the Virtual World Heritage Laboratory and the Luddy School of Informatics, Computing, and Engineering of Indiana University for their institutional support.

7. CONFLICTS OF INTEREST

The author is an Editorial Assistant for *Studies in Digital Heritage*. To avoid potential conflicts of interest, the author has not been involved in the editorial handling or decision-making processes for this manuscript. The author reports no other conflicts of interest.

8. REFERENCES

Alfano, F.R.D.A., Iannace, G., Ianniello, C. and Ianniello, E. 2015. "Velaria" in Ancient Roman Theatres: Can they have an Acoustic Role? *Energy and Buildings 95*, 98-105.

Amos, M. and Webster, J. 2022. Crowd-Sourced Identification of Characteristics of Collective Human Motion. *Artificial Life 28*, 401-422.

Antunes, R.F. and Correia, L. 2022. Virtual Simulations of Ancient Sites Inhabited by Autonomous Characters: Lessons from the Development of Easy-Population. *Digital Applications in Archaeology and Cultural Heritage 26*, e00237.

Arkin, R.C. 1989. Navigational Path Planning for a Vision-Based Mobile Robot. *Robotica 7*, 49-63.

Bickler, S.H. 2021. Machine Learning Arrives in Archaeology. *Advances in Archaeological Practice 9*, 186-191.

Bomgardner, D. 2021. *The Story of the Roman Amphitheatre*. Routledge, London and New York.

Bond, S.E. 2016. *Trade and Taboo: Disreputable Professions in the Roman Mediterranean*. University of Michigan Press.

Brøns, C., Skovmøller, A. and Gisler, J.-R. 2017. Colour-Coding the Roman Toga: The Materiality of Textiles Represented In Ancient Sculpture. *Antike Kunst 60*, 55-79.

Brughmans, T., Hanson, J.W., Mandich, M.J., Romanowska, I., Rubio-Campillo, X., Carrignon, S., Collins-Elliott, S., Crawford, K., Daems, D., Fulminante, F., De Haas, T., Kelly, P., Del Carmen Moreno Escobar, M., Paliou, E., Prignano, L. and Ritondale, M. 2019. Formal Modelling Approaches to Complexity Science in Roman Studies: A Manifesto. *Theoretical Roman Archaeology Journal 2*, 1-19.

Campisi, M. 1987. Intonaci Coloriture e Rivestimenti Edilizi in Età Romana. *Bollettino d'Arte*.

Cegielski, W.H. and Rogers, J.D. 2016. Rethinking the role of Agent-Based Modeling in Archaeology. *Journal of Anthropological Archaeology 41*, 283-298.

Delling, D., Sanders, P., Schultes, D. and Wagner, D. 2009. Engineering Route Planning Algorithms. In *Algorithmics of Large and Complex Networks: Design, Analysis, and Simulation*, Lerner, J., et al. Eds. Springer Berlin Heidelberg, Berlin, Heidelberg, 117-139.

Duives, D.C., Daamen, W. and Hoogendoorn, S.P. 2015. Quantification of the Level of Crowdedness for Pedestrian Movements. *Physica A: Statistical Mechanics and Its Applications 427*, 162-180.

Frischer, B. 2024. Rome Reborn 4.0: A Virtual Tour into the Heart of the Eternal City. In *Archaeomatica: Technologie per i Beni Culturali*, Fasolo, M. Ed., 14-26.

Futrell, A. 1997. *Blood in the Arena: the Spectacle of Roman Power*. University of Texas Press, Austin.

Geoerg, P., Schumann, J., Boltes, M. and Kinateder, M. 2022. How People with Disabilities Influence Crowd Dynamics of Pedestrian Movement through Bottlenecks. *Sci Rep 12*, 14273.

Gietl, R., Doneus, M. and Fera, M. 2007. Cost Distance Analysis in an Alpine Environment:

Comparison of Different Cost Surface Modules. In *Layers of Perception: Proceedings of the 35th International Conference on Computer Applications and Quantitative Methods in Archaeology (CAA)*, Posluschny, A., et al. Eds. Bonn: Rudolf Habelt, Berlin, 1-9.

Golvin, J.-C. 1988. *L'amphithéâtre Romain: Essai sur la Théorisation de sa Forme et de ses Fonctions*. Diffusion de Boccard, Paris.

Gravit, M., Kirik, E., Savchenko, E., Vitova, T. and Shabunina, D. 2022. Simulation of Evacuation from Stadiums and Entertainment Arenas of Different Epochs on the Example of the Roman Colosseum and the Gazprom Arena. *Fire 5*, 20.

Guidi, G. and Frischer, B.D. 2020. 3D Digitization of Cultural Heritage. In *3D Imaging, Analysis and Applications*, Liu, Y., et al. Eds. Springer, 631-697.

Gutierrez, D., Frischer, B., Cerezo, E., Gomez, A. and Seron, F. 2007. AI and Virtual Crowds: Populating the Colosseum. *Journal of Cultural Heritage 8*, 176-185.

Hegarty, Z. 2024. Navigating Antiquity: Digital Insights into Ancient Architecture with AI Agents. In *IU Graduate AI Expo* Indiana University, Bloomington.

Herzog, I. 2020. Spatial Analysis Based on Cost Functions. In *Archaeological Spatial Analysis: A Methodological Guide*, Gillings, M., et al. Eds. Routledge, London and New York, 333-358.

Hinman, M.R., O'connell, J.K., Dorr, M., Hardin, R., Tumlinson, A.B. and Varner, B. 2014. Functional Predictors of Stair-Climbing Speed in Older Adults. *J Geriatr Phys Ther 37*, 1-6.

Hopkins, K. and Beard, M. 2005. *The Colosseum*. Harvard University Press, Cambridge.

Icomos 2017. Principles of Seville: International Principles of Virtual Archaeology.

Lake, M. 2020. Spatial Agent-Based Modelling. In *Archaeological Spatial Analysis: A Methodological Guide*, Gillings, M., et al. Eds. Routledge, London and New York, 247-272.

Lakkaraju, H. and Bastani, O. 2020. "How do I fool you?": Manipulating User Trust via Misleading Black Box Explanations. In *Proceedings of the Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, New York, NY, USA2020 Association for Computing Machinery, 79–85.

Larsson, A., Ranudd, E., Ronchi, E., Hunt, A. and Gwynne, S. 2021. The Impact of Crowd Composition on Egress Performance. *Fire Safety Journal 120*.

Lemonari, M., Blanco, R., Charalambous, P., Pelechano, N., Avraamides, M., Pettré, J. and Chrysanthou, Y. 2022. Authoring Virtual Crowds: A Survey. *Computer Graphics Forum 41*, 677-701.

Lovatt, H. 2016. Flavian Spectacle: Paradox and Wonder. *A Companion to the Flavian Age of Imperial Rome*, 361-375.

Maïm, J., Haegler, S., Yersin, B., Mueller, P., Thalmann, D. and Van Gool, L. 2007. Populating Ancient Pompeii with Crowds of Virtual Romans. In *Proceedings of the Proceedings of the 8th International conference on Virtual Reality, Archaeology and Intelligent Cultural Heritage*, Brighton, UK2007 Eurographics Association, 109–116.

Marsaglia, G. 2003. Xorshift RNGs. *Journal of Statistical Software 8*, 1 - 6.

Micoli, L., Umair, M. and Guidi, G. 2024. Virtual Reconstruction of no longer Existing Archaeological Structures in Highly Urbanized Areas. *Studies in Digital Heritage 8*, (In publication).

Montilla, R.B. 1969. The Awnings of Roman Theatres and Amphitheatres. *Theatre Survey 10*, 75-88.

Moussaid, M., Schinazi, V.R., Kapadia, M. and Thrash, T. 2018. Virtual Sensing and Virtual Reality: How New Technologies Can Boost Research on Crowd Dynamics. *Front Robot AI 5*, 82.

Musse, S.R. and Thalmann, D. 1997. A Model of Human Crowd Behavior : Group Inter-Relationship and Collision Detection Analysis. In *Computer Animation and Simulation '97*, Thalmann, D. and M. van de Panne Eds. Springer Vienna, Vienna, 39-51.

Nielsen, I. 2006. Colosseum. In *Brill's New Pauly* Brill.

Oasys 2015. Modern-Day Engineers vs. The Romans Oasys, https://www.oasys-software.com/news/modern-day-engineers-vs-romans/.

Pearson, J. 2013. *Arena*. Bloomsbury Reader, London.

Rawson, E. 1987. Discrimina Ordinum: The Lex Julia Theatralis. *Papers of the British School at Rome 55*, 83-114.

Richards-Rissetto, H. 2022. Technological Challenges to Practicing 3D Ethics in Archaeology. In *Digital Heritage and Archaeology in Practice: Data, Ethics, and Professionalism*, Watrall, E. and L. Goldstein Eds. University Press of Florida, Gainesville, 163-193.

Rose, P. 2005. Spectators and Spectator Comfort in Roman Entertainment Buildings: A Study in Functional Design. *Papers of the British School at Rome 73*, 99-130.

Schimpl, M., Moore, C., Lederer, C., Neuhaus, A., Sambrook, J., Danesh, J., Ouwehand, W. and Daumer, M. 2011. Association Between Walking Speed and Age in Healthy, Free-Living Individuals Using Mobile Accelerometry—A Cross-Sectional Study. *PLOS ONE 6*, e23299.

Sports Grounds Safety Authority, U.K. 2024. Guide to Safety at Sports Grounds (Green Guide) Sports Grounds Safety Authority.

Taylor, J. 2014. Colosseum. In *Time Scanners*, Acaster, L. Ed. National Geographic.

Unity 2024. Unity Documentation.

Van Der Berg, J., Lin, M. and Manocha, D. 2008. Reciprocal Velocity Obstacles for Real-Time Multi-Agent Navigation. In *IEEE International Conference on Robotics and Automation* IEEE, Pasadena, CA, 1928-1935.

Van Der Berg, J., Patil, S., Sewall, J., Manocha, D. and Lin, M. 2008. Interactive Navigation of Individual Agents in Crowded Environments. In *Proceedings of the Symposium on Interactive 3D Graphics and Games (I3D)*2008.

Vosinakis, S. 2020. The Use of Digital Characters in Interactive Applications for Cultural Heritage. In *Applying Innovative Technologies in Heritage Science*, Pavlidis, G. Ed. IGI Global, Hershey, PA, USA, 109-137.

Warren, W.H. 2018. Collective Motion in Human Crowds. *Curr Dir Psychol Sci 27*, 232-240.

Welch, K.E. 2007. *The Roman Amphitheatre: From Its Origins to the Colosseum*. Cambridge University Press, Cambridge.

Zanker, P. 1998. *The Power of Images in the Age of Augustus*. University of Michigan Press, Ann Arbor, MI.

Zink, S. 2014. Polychromy in Roman Architecture Colours: Colours, Materials, and Techniques. In *Transformations: Classical Sculpture in Colour*, Østergaard, J.S. and A.M. Nielsen Eds. Ny Carlsberg Glyptotek.