Autonomous Agent Crowd Simulation in Immersive Virtual Reality

Massimiliano Pascoli[1,*,†]

[1]Human-Computer Interaction Lab, Department of Mathematics, Computer Science and Physics, University of Udine

Abstract

In the last decade, the simulation of credible and convincing crowd behavior has become increasingly important across a wide range of domains, including video games, social networks and the Metaverse, cinematic computer-generated imagery, urban planning, evacuation and riot simulations, military training, and cultural heritage preservation. This research aims to design techniques to simulate realistic crowds of autonomous agents in an efficient and scalable way, with particular attention to immersive Virtual Reality applications. The research focuses on increasing the variety and plausibility of agents relying on psychological metrics and profiles, appearance, individual behaviors, and cultural aspects without neglecting simulation performance and optimization. Key aspects of my research include modeling personality traits, emotional states, and their influence on agents' clothing and facial expressions to foster richer individualization. In parallel, my research will explore algorithms for resource-efficient, fast, and scalable crowd simulation. In particular, the research will employ data-driven methods, such as Machine Learning and Deep Learning, to support behavior modeling and the adaptability of agents through context understanding. The proposed approaches will be validated through applied scenarios, including user testing in serious games and training applications across various domains.

Keywords

Crowd Simulation, Crowd Appearance, Agents' Behavior, Artificial Intelligence, Autonomous Agents, Crowd Perception, Computer Graphics, User Studies

1. Introduction

Since the seminal work of Reynolds on flocking behavior [1], crowd simulation (CS) has evolved substantially, driven by both advances in hardware and rising demands for realism and interactivity across domains such as games, training, urban planning, and digital heritage [2, 3]. Today, CS supports diverse applications—from lifelike non-player characters in games to simulating pedestrian flow for safety and architectural analysis [4, 5, 6].

Despite its maturity, CS in virtual environments (VEs) still faces key challenges. These include improving behavioral realism through psychological models, supporting multiscale crowd visualization, applying CS in immersive Virtual Reality (iVR), and maintaining performance as fidelity increases [2]. High-fidelity simulations require complex algorithms for behavior, navigation, and rendering, each computationally demanding, particularly in resource-constrained iVR systems [7, 8]. To address this, efficient implementations often involve GPU acceleration and advanced rendering techniques [9, 10].

Balancing realism and scalability is central to current research, especially in path planning, which is often a bottleneck in large-scale simulations [11, 12]. In iVR, this challenge is amplified by the performance constraints of head-mounted displays, where stereo rendering and interaction require tight computational control.

To structure my ongoing research, three main areas can be identified: **Perception and Behavior**, focusing on agent psychology and user perception; **Performance and Optimization**, targeting efficient, scalable simulations; and **Graphics and Appearance**, concerned with visual diversity in crowds. A

CHItaly 2025: Technologies and Methodologies of Human-Computer Interaction in the Third Millenium, Doctoral Consortium, 6-10 October 2025, Salerno, Italy
*Corresponding author.
†PhD Student currently in the 2^{nd} year of the XXXIX cycle.
massimiliano.pascoli@uniud.it (M. Pascoli)
https://dmif.uniud.it/it/didattica/dottorato/iai/dottorandi/massimiliano-pascoli (M. Pascoli)
0009-0007-3257-2547 (M. Pascoli)

fourth, transversal area, **CS Applications**, integrates advances across the first three areas into practical applications.

The rest of this manuscript is organized as follows: Sec. 2 reviews the relevant literature, Sec. 3 summarizes results achieved so far, and Sec. 4 outlines ongoing and future research directions.

2. Related Work

The field of CS has a well-established history, with almost four decades of research [2] that highlights several persistent open challenges, particularly in immersive environments. A classification of CS approaches was provided in literature [3], proposing a taxonomy based on the granularity of control performed by the simulation on the agents. My research mainly focuses on mesoscopic systems, which aim to simulate interactions among dynamic agent groups. This approach offers a reasonable compromise between the benefits of complex single-agent behavior achieved by microscopic models and the performance of macroscopic systems that can handle thousands of agents with fewer individual details.

In the following, the relevant literature is reported, divided by research area already explored.

2.1. Perception and Behavior

In CS, the integration of emotions or personality traits, especially using the OCEAN (also known as Big Five) model, has shown promise in enhancing realism and behavioral diversity. Each dimension of OCEAN (Openness, Conscientiousness, Extraversion, Agreeableness, and Neuroticism) correlates with observable agent behaviors. Prior works have successfully incorporated OCEAN in scenarios such as emergency evacuations and social group behaviors [13, 14, 15]. Moreover, various personality and emotion models (e.g., HEXACO [16], PEN [17], PAD [18], SIR [19]) have been applied in CSs.

In the iVR domain, previous studies have mostly focused on user-agent interaction via spatial cues like collision avoidance and eye-gaze behavior [20, 21], rather than personality-driven agent behavior.

However, current CS implementations rarely integrate personality traits and, at the same time, support iVR. The added complexity of personality and emotional modeling in immersive contexts, where real-world scale, first-person perspective, and strict performance constraints become crucial, may limit their feasibility and scalability in practical applications.

The literature lacks systems capable of supporting unsupervised, personality-driven grouping or dynamically assigning goals based on agent traits and environmental features, enabling more nuanced and heterogeneous crowd dynamics. My research aims to bridge this gap by introducing a personality-aware mesoscopic CS framework specifically tailored for iVR, focusing on enhancing users' sense of social presence and realism while respecting the performance demands of iVR platforms.

2.2. Performance and Optimization

Path planning has been a foundational topic in algorithmic and robotics research for decades, long before the advent of modern game engines. Numerous algorithms have been proposed over the years (see [22, 23] for examples), each exhibiting different trade-offs in terms of execution time, memory consumption, and the quality or length of the computed path, depending on the structure of the graph and the heuristic employed. The literature offers a wide array of strategies, each tailored to excel in specific problem domains.

To address the computational demands of path planning, two primary optimization strategies can be used. The first focuses on improving the a priori efficiency in terms of time and memory of the data structures used by specific algorithms. The second strategy centers on selecting the most suitable algorithm from a portfolio based on the specific characteristics of a given problem instance: a challenge known as the algorithm selection problem [24]. Recent research has investigated automating this selection process using machine learning models [25]. In this context, a key challenge lies in balancing the flexibility offered by large algorithm portfolios with the increased complexity of identifying which algorithms to evaluate for a particular instance.

Several researchers have focused on evaluating and comparing the performance of different path planning methods, and have proposed selection algorithms that recommend the most suitable option for a given scenario [26]. Comparative studies have examined algorithms such as A* with the Manhattan heuristic, Dijkstra, and Breadth-First Search, assessing their efficiency in terms of execution time, number of expanded nodes, path length, and node count in the returned path [27]. The research community also performed path planning algorithm evaluation on different platforms, such as Android [28]. The literature has also explored path planning in 3D voxelized environments [29] and the use of machine learning to optimize path planning strategies [30, 31].

However, applying machine learning to algorithm selection typically requires a large and diverse dataset containing performance metrics for multiple algorithms across varied virtual environments, which poses a significant practical challenge.

3. Results So Far

In the first semester of my doctoral research, I focused on reviewing the state of the art in CS for extended reality, with particular attention to iVR. This work led to the development of a journal article [32] (Work 1) based on my MSc thesis, covering the perception and behavior area.

The paper proposes a novel, lightweight mesoscopic CS system that integrates synthetic agent personalities based on the OCEAN personality framework. Each agent is assigned a unique personality profile and autonomously forms groups using machine learning-based clustering. Additionally, agents express their personalities through probabilistic behavioral animations, with behaviors mapped along the five OCEAN dimensions treated as continuous spectrums. A within-subjects user study involving 40 participants demonstrated that including personality traits with the tested system significantly enhances users' sense of social presence, perceived realism, and the matching between agents and environment. To facilitate replicability and a wider application of the results, the code snippets of the proposed CS system are available on request.

During the second semester, my focus shifted to performance optimization in CS systems. We wrote a paper, yet to be published, exploring (i) the trade-offs between memory utilization and time taken by path planning algorithms, and (ii) the similarity between the algorithmically found paths and the paths drawn by users (Work 2).

This work presents a framework for generating VEs, extracting navigation meshes, and benchmarking path planning algorithms across 1.5M automatically generated problems spanning three map types. It evaluates execution time and memory usage in single- and multi-threaded settings, showing that Best-First Search offers a resource-efficient alternative to A*, while Breadth-First Search balances path quality and speed. An exploratory user study compared participant-drawn paths to algorithm-generated ones, revealing significant effects of map type and algorithm on the similarity between the user-drawn path and the algorithm-generated ones. The framework supports future work on automated algorithm selection for resource-constrained iVR systems; consequently, the produced dataset and the framework will be available to other researchers to facilitate further expansion on the topic of path planning.

Recently, I had the opportunity to work on the CS applications area. The research is yet to be published and focuses on the effects of self and others' locomotion for cultural heritage learning applications in iVR (Work 3).

In this user study, native Italian speakers learned vocabulary related to rural objects from the Natisone Valleys while accompanied by a crowd of virtual agents simulating a shared experience. The rural objects were scattered around the VE, and the participants had to find them. Three locomotion conditions were tested: teleportation, arm-swing, and a hybrid "lerport" (teleportation with continuous movement visualization for others). The simulated crowd behavior is visualized as if every agent were controlled by a different user, accounting for how the locomotion technique set for the experiment condition would be shown in the represented VE. In particular, the crowd explores the VE using a state-based pseudo-random algorithm, aiding the user in the rural object research task. Teleportation significantly improved vocabulary retention and reduced simulator sickness compared to arm-swing, and also enhanced presence over lerport.

4. Planned Research

Aligned with my core research areas, I have outlined additional research proposals that will contribute to the development of my PhD thesis, advancing CS in iVR across theoretical and applied dimensions.

Bootstrapping the Metaverse Building on Work 3, this research proposal falls in the CS Applications area, and will explore how CS can populate interactive VE with lifelike agents when few real users are present in a Metaverse. By integrating insights from performance, behavior, and visual realism, the proposed solution aims to enhance user immersion in large-scale, networked Metaverse applications. The efficacy of the proposed solutions will be tested with user studies.

Performance Prediction for Path Planning Algorithms Expanding on Work 2, this research proposal will develop predictive models to estimate algorithm performance under varying computational constraints and VE types. The goal of the proposed solution is to enable optimal path planning algorithm selection in resource-constrained iVR systems, contributing to both CS optimization and practical deployment, thus falling into the Performance and Optimization area.

Crowd Appearance and Behavior This research proposal will investigate how visual saliency, through clothing, expressions, or posture, interacts with agent behavior to influence user perception. The proposed solutions aim to reduce visual monotony in large crowds while maintaining performance, and their efficacy will be evaluated with a user study. The results will contribute to the Graphics and Appearance area of CS.

Fighting Crowds In an upcoming project spanning the CS Applications and Perception and Behavior areas, a simulation system will be developed to model conflict dynamics in crowds. This will support a deep learning model predicting the likelihood of fights based on behavior and context. The system will generate synthetic training data and support research in safety, public space design, and multi-agent interaction. Simulation realism and fidelity of the system will be tested with a user study.

The findings from these research proposals will form the foundation of my thesis, contributing toward a robust framework for iVR crowd simulation, addressing performance, perception, visual fidelity, and real-world applications.

5. Declaration on Generative AI

During the preparation of this work, the author used Grammarly in order to perform grammar and spelling checks, paraphrase and reword (in particular to summarize the content of this paper), and improve writing style – in accordance with the generative AI usage taxonomy that can be found at https://www.ceur-ws.org/GenAI/Taxonomy.html. After using this tool, the author reviewed and edited the content as needed and takes full responsibility for the publication's content.

References

[1] C. W. Reynolds, Flocks, herds and schools: A distributed behavioral model, in: Proceedings of the 14th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '87,

Association for Computing Machinery, New York, NY, USA, 1987, pp. 25–34. doi:10.1145/37401.37406.
[2] S. R. Musse, V. J. Cassol, D. Thalmann, A history of crowd simulation: The past, evolution, and new perspectives, The Visual Computer 37 (2021) 3077–3092. doi:10.1007/s00371-021-02252-w.
[3] S. Yang, T. Li, X. Gong, B. Peng, J. Hu, A review on crowd simulation and modeling, Graphical Models 111 (2020) 101081. doi:10.1016/j.gmod.2020.101081.
[4] M. Zyda, From visual simulation to virtual reality to games, Computer 38 (2005) 25–32. doi:10.1109/MC.2005.297.
[5] G. Aschwanden, J. Halatsch, G. Schmitt, Crowd simulation for urban planning, in: Proceedings of eCAADe, volume 2008, 2008.
[6] C. K. Lim, K. L. Tan, A. A. Zaidan, B. B. Zaidan, A proposed methodology of bringing past life in digital cultural heritage through crowd simulation: A case study in George Town, Malaysia, Multimed Tools Appl 79 (2020) 3387–3423. doi:10.1007/s11042-019-07925-2.
[7] D. Thalmann, S. R. Musse, Crowd Simulation, Springer Science & Business Media, 2012.
[8] G. Ryder, A. M. Day, Survey of Real-Time Rendering Techniques for Crowds, Computer Graphics Forum 24 (2005) 203–215. doi:10.1111/j.1467-8659.2005.00844.x.
[9] A. Bleiweiss, Gpu accelerated pathfinding, in: Proceedings of the 23rd ACM SIGGRAPH/EUROGRAPHICS symposium on Graphics hardware, 2008, pp. 65–74.
[10] L. Kavan, S. Dobbyn, S. Collins, J. Žára, C. O'Sullivan, Polypostors: 2D polygonal impostors for 3D crowds, in: Proceedings of the 2008 Symposium on Interactive 3D Graphics and Games, ACM, Redwood City California, 2008, pp. 149–155. doi:10.1145/1342250.1342273.
[11] M. Kapadia, N. Pelechano, J. Allbeck, N. Badler, Virtual Crowds: Steps Toward Behavioral Realism, Synthesis Lectures on Visual Computing: Computer Graphics, Animation, Computational Photography and Imaging, Springer International Publishing, Cham, 2016. doi:10.1007/978-3-031-02586-0.
[12] L. Jaulin, Path planning using intervals and graphs, Reliable computing 7 (2001) 1–15.
[13] F. Durupinar, U. Gudukbay, A. Aman, N. I. Badler, Psychological Parameters for Crowd Simulation: From Audiences to Mobs, IEEE Trans. Visual. Comput. Graphics 22 (2016) 2145–2159. doi:10.1109/TVCG.2015.2501801.
[14] R. Zhou, Y. Ou, W. Tang, Q. Wang, B. Yu, An Emergency Evacuation Behavior Simulation Method Combines Personality Traits and Emotion Contagion, IEEE Access 8 (2020) 66693–66706. doi:10.1109/ACCESS.2020.2985987.
[15] S. J. Guy, S. Kim, M. C. Lin, D. Manocha, Simulating heterogeneous crowd behaviors using personality trait theory, in: Proceedings of the 2011 ACM SIGGRAPH/Eurographics Symposium on Computer Animation, ACM, Vancouver British Columbia Canada, 2011, pp. 43–52. doi:10.1145/2019406.2019413.
[16] K. Lee, M. C. Ashton, The HEXACO Model of Personality Structure, in: V. Zeigler-Hill, T. K. Shackelford (Eds.), Encyclopedia of Personality and Individual Differences, Springer International Publishing, Cham, 2020, pp. 1932–1936. doi:10.1007/978-3-319-24612-3_1227.
[17] S. B. G. Eysenck, H. J. Eysenck, The place of impulsiveness in a dimensional system of personality description, British Journal of Social and Clinical Psychology 16 (1977) 57–68. doi:10.1111/j.2044-8260.1977.tb01003.x.
[18] A. Mehrabian, Pleasure-arousal-dominance: A general framework for describing and measuring individual differences in Temperament, Current Psychology 14 (1996) 261–292. doi:10.1007/BF02686918.
[19] L. Fu, W. Song, W. Lv, S. Lo, Simulation of emotional contagion using modified SIR model: A cellular automaton approach, Physica A: Statistical Mechanics and its Applications 405 (2014) 380–391. doi:10.1016/j.physa.2014.03.043.
[20] F. Berton, L. Hoyet, A.-H. Olivier, J. Bruneau, O. Le Meur, J. Pettre, Eye-Gaze Activity in Crowds: Impact of Virtual Reality and Density, in: 2020 IEEE Conference on Virtual Reality and 3D User Interfaces (VR), 2020, pp. 322–331. doi:10.1109/VR46266.2020.00052.
[21] A. Koilias, C. Mousas, C.-N. Anagnostopoulos, I feel a moving crowd surrounds me: Exploring

tactile feedback during immersive walking in a virtual crowd, Computer Animation and Virtual Worlds 31 (2020) e1963. doi:10.1002/cav.1963.

[22] M. Reda, A. Onsy, A. Y. Haikal, A. Ghanbari, Path planning algorithms in the autonomous driving system: A comprehensive review, Robotics and Autonomous Systems 174 (2024) 104630. doi:10.1016/j.robot.2024.104630.

[23] L. Liu, X. Wang, X. Yang, H. Liu, J. Li, P. Wang, Path planning techniques for mobile robots: Review and prospect, Expert Systems with Applications 227 (2023) 120254. doi:10.1016/j.eswa.2023.120254.

[24] L. Kotthoff, Ranking Algorithms by Performance, in: P. M. Pardalos, M. G. Resende, C. Vogiatzis, J. L. Walteros (Eds.), Learning and Intelligent Optimization, Springer International Publishing, Cham, 2014, pp. 16–20. doi:10.1007/978-3-319-09584-4_2.

[25] A. Kostovska, G. Cenikj, D. Vermetten, A. Jankovic, A. Nikolikj, U. Skvorc, P. Korosec, C. Doerr, T. Eftimov, PS-AAS: Portfolio Selection for Automated Algorithm Selection in Black-Box Optimization, in: Proceedings of the Second International Conference on Automated Machine Learning, PMLR, 2023, pp. 11/1–17.

[26] A. Kherrour, M. Robol, M. Roveri, P. Giorgini, Evaluating Heuristic Search Algorithms in Pathfinding: A Comprehensive Study on Performance Metrics and Domain Parameters, Electronic Proceedings in Theoretical Computer Science 391 (2023) 102–112. doi:10.4204/EPTCS.391.12. arXiv:2310.02346.

[27] V. Morina, R. Rafuna, A comparative analysis of pathfinding algorithms in npc movement systems for computer games, in: 12th UBT Annual International Conference On Computer Science And Engineering, 2023.

[28] T. Hwu, A. Y. Wang, N. Oros, J. L. Krichmar, Adaptive robot path planning using a spiking neuron algorithm with axonal delays, IEEE Transactions on Cognitive and Developmental Systems 10 (2017) 126–137.

[29] D. Brewer, N. Sturtevant, Benchmarks for Pathfinding in 3D Voxel Space, Proceedings of the International Symposium on Combinatorial Search 9 (2021) 143–147. doi:10.1609/socs.v9i1.18464.

[30] Z. Zhang, R. Wu, Y. Pan, Y. Wang, Y. Wang, X. Guan, J. Hao, J. Zhang, G. Li, A Robust Reference Path Selection Method for Path Planning Algorithm, IEEE Robotics and Automation Letters 7 (2022) 4837–4844. doi:10.1109/LRA.2022.3152687.

[31] Z. Meng, Optimization Path Planning Algorithm Based on Deep Reinforcement Learning, in: 2024 Asia-Pacific Conference on Software Engineering, Social Network Analysis and Intelligent Computing (SSAIC), 2024, pp. 872–876. doi:10.1109/SSAIC61213.2024.00176.

[32] M. Pascoli, F. Buttussi, K. Schekotihin, L. Chittaro, Introducing Agent Personality in Crowd Simulation Improves Social Presence and Experienced Realism in Immersive VR, IEEE Transactions on Visualization and Computer Graphics (2025) 1–15. doi:10.1109/TVCG.2025.3543740.