I work on interactive and learning-based algorithms for simulating crowds and modeling the environments they move through. My contributions fall in two connected areas: simulation of multi-agent motion, from pedestrians to ant swarms[1,2,3,4,5,6,7,8,9,10,11,12,13,14], and the analysis and synthesis of 3D scenes, which supply the spaces those agents inhabit[15,16,17,18,19,20,21,22,23,24,25,26]. These areas are converging on a single goal: realistic, data-grounded virtual environments in which populations of simulated agents can be studied at scale. Such environments allow planners, designers, engineers, and researchers to explore how people move through a space and how changes to it affect outcomes.

My research has been recognized by the *ACM SIGGRAPH Conference on Motion in Games*, where I received the **Best Paper Award**[10], by the ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games, where I received the **Best Student Paper Award**, and by the ACM SIGGRAPH/Eurographics Symposium on Computer Animation, where I received the **Best Presentation Award**[14]. In 2018, I was also an internationally selected **finalist** at both the *ACM SIGGRAPH Thesis Fast Forward* and the *ACM SIGGRAPH Asia Doctoral Consortium*.

The two thrusts below summarize my research work, followed by my future research agenda.

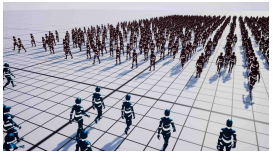
Thrust I: Crowds, Groups, and Multi-Agent Motion Dynamics

Multi-agent simulation has applications in entertainment, pedestrian analysis, urban planning, robotics, and autonomous systems. A **crowd** is a collection of independent, self-actuated agents. Each agent has individual navigational goals in this shared environment. Since agents share their environment, they can interact and collide with each other. Agent movement is controlled by a navigation algorithm, which must ensure that **agents progress towards their goals, while avoiding collisions**. However, computing collision-free agent motion is difficult due to the complexity of such dynamic interactions. Field experiments are one important avenue for testing a navigation algorithm. A more efficient and less costly approach is to conduct virtual simulation experiments, which allow quick insight into the dynamics and performance of such navigation algorithms. Furthermore, a simulation allows us to easily test "what-if" scenarios.

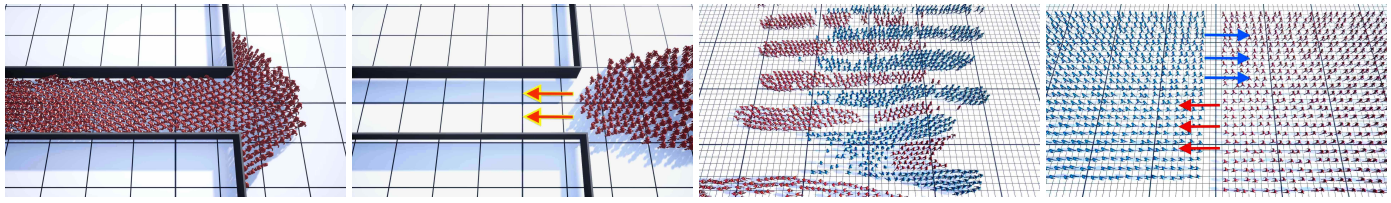


Agent

During a simulation, agents run through a continuous cycle of sensing and acting, where each cycle corresponds to a time step. At the beginning of each cycle, each agent independently computes a trajectory to its goal, while avoiding collisions with other agents or obstacles. In addition to collision-free movement, a simulation should **capture both individual and group behavior** observed in real crowds, while being computationally interactive.



Multi-Agent Crowds



(a) A group of agents passing through a narrow corridor. **(b)** Two groups exchanging positions.

Figure 1: Emergent phenomena in crowd simulation. Our results[10] recreate different phenomena observed in pedestrian crowds, including: (a) clogging and arching near bottlenecks[27], (b) groups self-organizing into lanes[28], with stable and real-time performance for +100K individuals. These results are made possible by framing the crowd motion as a constraint optimization problem, which is then efficiently solved using a GPU.

Physics-based Crowd Simulation

A core problem in multi-agent dynamics is to simulate dense, large crowds interactively. Our work allows simulation of +100K agents in real time[10], which we achieved using GPU parallelization, and a specialized physics-based solver that uses geometric constraints that capture agent-agent dynamics in a crowd (Fig. 1). **This result is currently the state-of-the-art in terms of crowd simulation speed**[29].

Later, we extended this framework to simulate groups moving in formations[5], including path planning that allows groups of agents to maintain their formation while navigating[6]. We also developed methods for simulating non disc-shaped agents, since a human or robot figure may be better approximated with other geometry[4], which allows capturing angular motion and turns more accurately. In addition, my group has also researched how other physics solver frameworks can be reformulated to simulate crowds[2]. Such frameworks allow quicker computation for specific crowd simulation scenarios compared to prior work.

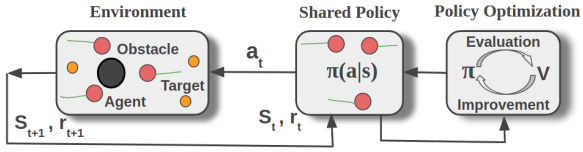


Figure 2: RL for multi-agent navigation[12].

Learning Crowd Behaviors

Physics solvers often fail to capture the somewhat stochastic nature of crowd dynamics. To address this challenge, we learned how crowds should move using crowdsourcing via **Bayesian optimization**[12]. Our method uses **deep reinforcement learning** for simulating crowds, where crowdsourcing is used to select policy hyper-parameters (Figure 2). Training agents with such parameters results in a crowd simulation that users prefer compared to state-of-the-art work.

Ongoing and future research directions

When simulating a crowd, the simulation is often only one part of the computational bottleneck, while visualizing an animated crowd is another. We developed a **Perceptual Foveated Rendering** approach[1]., where instead of animating each agent with the same fidelity, we focus on animating agents in the center of attention, and agents outside that region require fewer animation resources (Figure 3).

Outside of pedestrian crowds, we simulated 3D crowds[8], such as ant swarm aggregations[14,30]. Understanding the dynamics of ant swarms and bacterial colonies is important for applications ranging from healthcare and transportation to manufacturing. This work modeled ants as self-actuating particles, which spontaneously form elastic links and also maintain fluid-like properties such as maintaining a certain density within an aggregation (Figure 3).

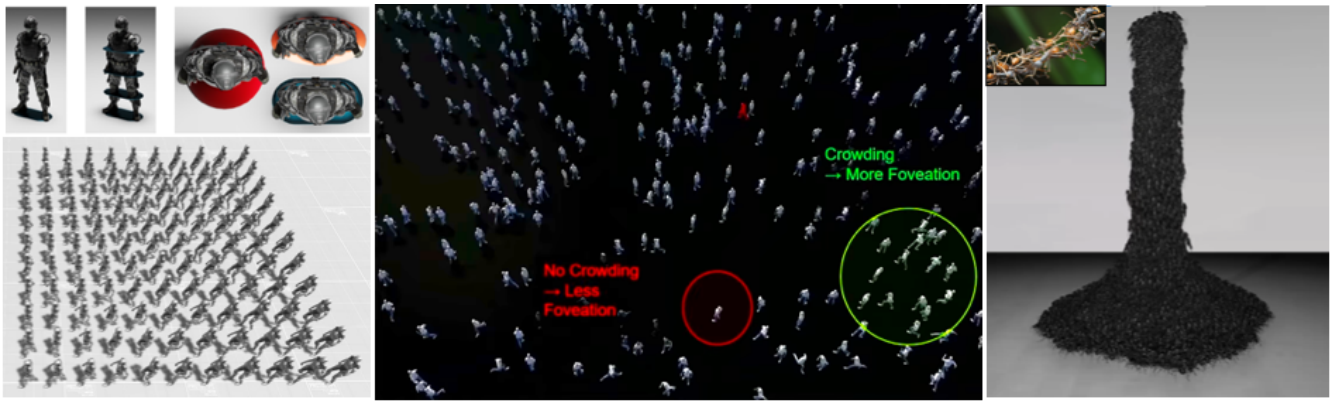


Figure 3: Recent multi-agent research of our lab includes simulating crowds as capsule-shaped particles which more accurately capture the pedestrian profile than discs, simulating agents moving in formation (left), accelerating rendering of crowd simulations using foveated rendering (middle), and simulating non-human, 3D crowds such as ants (right).

We are extending our multi-agent methodology to bacterial motion, robotic agents, and other biological entities, with learning-based approaches to assist in creating realistic dynamics. We are also developing digital twins of real environments from street imagery, integrating computer vision, pedestrian simulation, traffic simulation, and behavioral modeling to produce realistic virtual scenes populated with pedestrians and vehicles. The framework will support planners, first responders, and researchers in exploring how emergency situations unfold and how interventions affect outcomes, validated against real-world observations. These topics are integrated in pending federal grant proposals.

Thrust II: 3D Scene Synthesis and Understanding

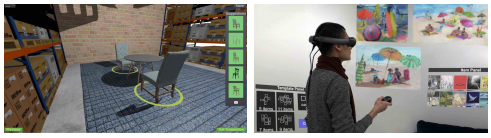
Virtual worlds are **challenging to construct**, requiring professionals with extensive training in 3D modeling software. My lab's research accelerates 3D content creation with computational methods for creating and understanding scenes.

One method we proposed uses a user-defined set of 3D objects and spatial arrangement constraints from which our method computes layout proposals (Figure 4). This is achieved by viewing layout synthesis through the lens of deformable body simulation. We formulate both problems as geometric constraint systems and solve them using continuous optimization. A layout synthesis solution is then an arrangement that satisfies such constraints[32]. The **novelty of our work**[21] is in enabling **real-time, interactive synthesis of layouts**, even for **large scale** scenes, which **were previously intractable** via traditional optimization-based approaches[33]. This work was later used to create a dataset of 3D spaces, which is **widely used today** for multiple AI tasks[34].

We also developed a **deep-learning generative transformer** approach for scene synthesis[18], by formulating the synthesis task to be similar to text sequence generation. An interior is represented by a series of tokens, where the goal of the neural network is to generate a sequence of tokens that represent a realistic space. In contrast to prior transformer-based and data-driven approaches[35], our method improves the quality of synthesized layouts beyond that of the rooms in the original dataset by incorporating additional terms in the loss function that encode ergonomic and geometric commonsense, such as keeping a table within reach of a seated person.

Ongoing and future research directions

Objects may not always fit a space, regardless of their function, due to differing visual features. For example, how well would a wheel fit into a car design? To understand this question better, we introduced an image-based classification approach for estimating 3D object and room style[15,24]. To accomplish this, we used (i) images where each object appears in multiple "styled" backgrounds, and (ii) image annotation labels that encode the style spectrum of each object in a



System for designing style-compatible spaces[17,24]

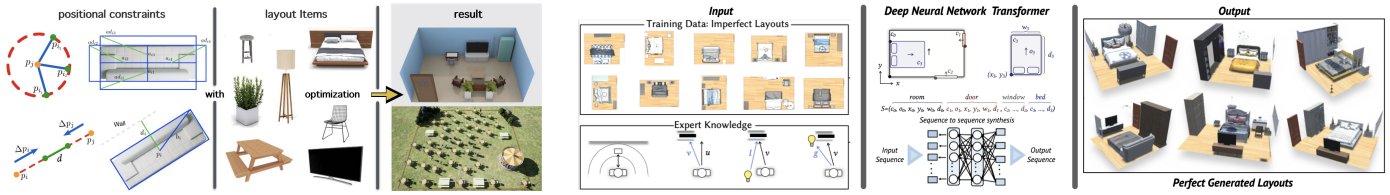


Figure 4: Left: Our results[21] in 3D scene synthesis are faster by an order of magnitude than previous probabilistic methods and allow **real-time, interactive synthesis of scenes that were previously intractable**[31]. Given input positional constraints and selected layout items, the algorithm rapidly outputs a variety of synthesized layouts. **Right**: We proposed a data-driven method for layout synthesis that combines interior expert knowledge with a data-driven generator based on a Transformer architecture[18]. The synthesized layouts integrate desirable properties not present in the original dataset.

relative manner. For example, a label indicates whether one table is more modern than another table. We have extended our style based work to Virtual Reality, and built interactive systems for quick synthesis of furniture in interior rooms based on interior design rules[17,16].

We are currently investigating how clutter and a room's organization affect the perception of space[25], as well as how suitable the room is for certain work tasks[26]. This work, done jointly with perception researchers, compares computational clutter measures against human judgments, so that synthesized scenes can be evaluated by how people actually experience a space and not only by geometric plausibility. Our goal is to bring these perceptual measures back into layout synthesis, so that generated rooms are not only realistic but also comfortable and usable for the people and agents that occupy them.

Next Steps

The following agenda is imbued in pending grant proposals:

Generative AI for agent-based simulations. Off-the-shelf AI diffusion models can produce visually plausible scenes, but often struggle to maintain consistent bodies, interactions, and motion across many agents. We will adapt generative models to produce structured, controllable representations of complex multi-agent behavior by training on curated video data with annotations that identify individual agents and their movements. These observations can be projected into a common spatial representation and converted into virtual agents that physics-based simulation can refine and extend. We will evaluate the resulting behaviors using visual metrics, trajectory statistics, and human studies. The resulting framework will provide a way to generate and study complex biological, robotic, and engineered systems—including rare or difficult-to-observe scenarios—without requiring them to be directly captured or reproduced physically.

Collective self-assembly across biological and engineered systems. Biological organisms and engineered agents can self-organize into structures and collective behaviors that are difficult to design explicitly, from insect aggregations and animal groups to robot swarms and adaptive materials. Yet we have limited understanding of the mechanics that govern these systems and how local interactions can produce robust, large-scale structures. Building on our prior work, we will develop models of agents that interact through simple local rules, allowing them to assemble, deform, separate, and reorganize in response to their environment. We will validate these models against observations from biology and engineering. Beyond graphics, these models can generate predictions that are difficult to test physically and provide a sandbox for studying collective behavior, swarm robotics, self-assembly, adaptive materials, and resilient infrastructure.

Simulation for real-world emergency preparedness. Emergencies change how people and vehicles move, but they are hard to simulate and ethically difficult to stage. In a pending grant proposal, we reconstruct animation-ready street scenes from Street View imagery, adding various hazards and mixed pedestrian and vehicle traffic. We plan to validate this against real-world observations. We also plan a fined-tuned language model that will let practitioners without simulation expertise build scenarios from plain-language descriptions, supporting preparedness in emergency management and transportation.

Computational bottlenecks in simulation. Interactive simulation of large collectives is computationally expensive, mainly due to computing agent-agent interactions, which is tied to nearest neighbor search, whose naive cost grows quadratically with agent count. No single spatial data structure performs best, since performance depends on density and interaction complexity. We plan to benchmark various spatial data structures, including recent neural acceleration approaches to identify optimal computational regimes, test whether they can be predicted online from local statistics such as density, and release the results as open-source infrastructure, with adaptive data structures as a stretch goal.

References

[1] Florin-Vladimir Stancu, Tomer Weiss, and Rafael Kuffner dos Anjos. Foveated animations for efficient crowd simulation. *Proceedings of the ACM on Computer Graphics and Interactive Techniques*, 8(1):1–14, 2025.

[2] Tomer Weiss. Projective multi-agent dynamics. In *Proceedings MIG 2025 - 18th ACM Conference on Motion, Interaction and Games*, Proceedings MIG 2025 - 18th ACM Conference on Motion, Interaction and Games. Association for Computing Machinery, Inc, December 2025. Publisher Copyright: © 2025 Copyright held by the owner/author(s).; 18th Annual ACM SIGGRAPH Conference on Motion, Interaction, and Games, MIG 2025 ; Conference date: 03-12-2025 Through 05-12-2025.

[3] Alexander Chen and Tomer Weiss. How to break crowd simulation algorithms. In *Proceedings of the Special Interest Group on Computer Graphics and Interactive Techniques Conference Talks*, SIGGRAPH Talks '25, New York, NY, USA, 2025. Association for Computing Machinery.

[4] Bilas Talukdar and Tomer Weiss. Generalized, dynamic multi-agent torso crowds. In *ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games*, I3D '25, 2025.

[5] Tomer Weiss. Fast position-based multi-agent group dynamics. *Proceedings of the ACM on Computer Graphics and Interactive Techniques*, 6(1):1–15, 2023.

[6] Ritesh Sharma, Tomer Weiss, and Marcelo Kallmann. Formation-aware planning and navigation with corridor shortest path maps. In *Computer Graphics Forum*, volume 43, page e14995. Wiley Online Library, 2024.

[7] Tomer Weiss. Implementing position-based real-time simulation of large crowds. In *2019 IEEE International Conference on Artificial Intelligence and Virtual Reality (AIVR)*. IEEE, 2019.

[8] Ritesh Sharma, Tomer Weiss, and Marcelo Kallmann. Plane-based local behaviors for multi-agent 3d simulations with position-based dynamics. In *2020 IEEE International Conference on Artificial Intelligence and Virtual Reality (AIVR)*, pages 214–217. IEEE, 2020.

[9] Masaki Nakada, Tao Zhou, Honglin Chen, Tomer Weiss, and Demetri Terzopoulos. Deep learning of biomimetic sensorimotor control for biomechanical human animation. *ACM Transactions on Graphics (TOG)*, 37(4):1–15, 2018.

[10] Tomer Weiss, Alan Litteneker, Chenfanfu Jiang, and Demetri Terzopoulos. Position-based multi-agent dynamics for real-time crowd simulation. In *ACM SIGGRAPH Conference on Motion in Games*, MIG '17, pages 10:1–10:8, 2017. **Best Paper Award**.

[11] Tomer Weiss, Alan Litteneker, Chenfanfu Jiang, and Demetri Terzopoulos. Position-based real-time simulation of large crowds. *Computers & Graphics*, 2018.

[12] Bilas Talukdar, Yunhao Zhang, and Tomer Weiss. Learning crowd motion dynamics with crowds. *Proceedings of the ACM on Computer Graphics and Interactive Techniques*, 7(1):1–17, 2024.

[13] Bilas Talukdar, Yunhao Zhang, and Tomer Weiss. Learning to simulate crowds with crowds. In *ACM SIGGRAPH 2023 Posters*, pages 1–2. ACM New York, NY, USA, 2023.

[14] Matthew Loges and Tomer Weiss. Simulating ant swarm aggregations dynamics. *Proceedings of the ACM on Computer Graphics and Interactive Techniques*, 8(4):1–16, 2025. **Best Presentation Award**.

[15] Esra Ataer-Cansizoglu, Hantian Liu, Tomer Weiss, Archi Mitra, Dhaval Dholakia, Jae-Woo Choi, and Dan Wulin. Room style estimation for style-aware recommendation. In *2019 IEEE International Conference on Artificial Intelligence and Virtual Reality (AIVR)*. IEEE, 2019.

[16] Hameedullah Farooki, Esra Ataer-Cansizoglu, Jae-Woo Choi, and Tomer Weiss. Interactive and scalable layout synthesis with design templates. In *2020 IEEE International Conference on Artificial Intelligence and Virtual Reality (AIVR)*, pages 235–238. IEEE, 2020.

[17] Haikun Huang, Yuxuan Zhang, Tomer Weiss, Rebecca W Perry, and Lap-Fai Yu. Interactive design of gallery walls via mixed reality. In *2020 IEEE International Conference on Artificial Intelligence and Virtual Reality (AIVR)*, pages 17–26. IEEE, 2020.

[18] Kurt Leimer, Paul Guerrero, Tomer Weiss, and Przemyslaw Musialski. Layoutenhancer: Generating good indoor layouts from imperfect data. In *SIGGRAPH Asia Conference Proceedings*, 2022.

[19] Benjamin S Liang, Andrew S Liang, Iran Roman, Tomer Weiss, Budmonde Duinkharjav, Juan Pablo Bello, and Qi Sun. Reconstructing room scales with a single sound for augmented reality displays. *Journal of Information Display*, pages 1–12, 2022.

[20] Mathew Schwartz, Tomer Weiss, Esra Ataer-Cansizoglu, and Jae-Woo Choi. Style similarity as feedback for product design. In *A New Perspective of Cultural DNA*, pages 27–42. Springer, 2021.

[21] Tomer Weiss, Alan Litteneker, Noah Duncan, Chenfanfu Jiang, Lap-Fai Yu, and Demetri Terzopoulos. Fast and scalable position-based layout synthesis. *IEEE Transactions on Visualization and Computer Graphics (TVCG)*, 2018.

[22] Tomer Weiss, Masaki Nakada, and Demetri Terzopoulos. Automated layout synthesis and visualization from images of interior or exterior spaces. In *IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 41–47. IEEE, 2017.

[23] Tomer Weiss, Alan Litteneker, Noah Duncan, Chenfanfu Jiang, Lap-Fai Yu, and Demetri Terzopoulos. Fast, scalable layout synthesis. In *IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*. IEEE, 2018. Extended abstract.

[24] Tomer Weiss, Ilkay Yildiz, Nitin Agarwal, Esra Ataer-Cansizoglu, and Jae-Woo Choi. Image-driven furniture style for interactive 3d scene modeling. In *Computer Graphics Forum*, volume 39, pages 57–68. Wiley Online Library, 2020.

[25] Yelda Semizer and Tomer Weiss. Visual clutter: The role of background texture, set size, and item organization. *Journal of Vision*, 23(9):4940–4940, 2023.

[26] Yelda Semizer, Sathvik Gopu, and Tomer Weiss. Visual clutter: The role of item quantity, spatial organization, and surface complexity in model and human estimates. *Journal of Vision*, 26(7):11–11, 2026.

[27] Dirk Helbing, Illés Farkas, and Tamas Vicsek. Simulating dynamical features of escape panic. *Nature*, 407(6803):487, 2000.

[28] Dirk Helbing, Lubos Buzna, Anders Johansson, and Torsten Werner. Self-organized pedestrian crowd dynamics: Experiments, simulations, and design solutions. *Transportation science*, 39(1):1–24, 2005.

[29] Wouter Van Toll and Julien Pettré. Algorithms for microscopic crowd simulation: Advancements in the 2010s. In *Computer Graphics Forum*, volume 40, pages 731–754. Wiley Online Library, 2021.

[30] Matthew Loges and Tomer Weiss. Simulating the mechanics of ant swarm aggregations. In *Proceedings of the Special Interest Group on Computer Graphics and Interactive Techniques Conference Posters*, pages 1–2, 2025.

[31] Lap-Fai Yu, Sai Kit Yeung, Chi-Keung Tang, Demetri Terzopoulos, Tony F. Chan, and Stanley Osher. Make it home: automatic optimization of furniture arrangement. *ACM Trans. Graphics*, 30(4):86, 2011.

[32] Matthias Müller, Bruno Heidelberger, Marcus Hennix, and John Ratcliff. Position based dynamics. *Journal of Visual Communication and Image Representation*, 18(2):109–118, 2007.

[33] Song-Hai Zhang, Shao-Kui Zhang, Yuan Liang, and Peter Hall. A survey of 3d indoor scene synthesis. *Journal of Computer Science and Technology*, 34(3):594–608, 2019.

[34] Huan Fu, Bowen Cai, Lin Gao, Ling-Xiao Zhang, Jiaming Wang, Cao Li, Qixun Zeng, Chengyue Sun, Rongfei Jia, Binqiang Zhao, et al. 3d-front: 3d furnished rooms with layouts and semantics. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10933–10942, 2021.

[35] Jiachen Liu, Yuan Xue, Haomiao Ni, Rui Yu, Zihan Zhou, and Sharon X Huang. Computer-aided layout generation for building design: A review. *Computational Visual Media*, 2025.