I work on **interactive** and **learning-based** algorithms for simulation and modeling. My research has been recognized by the *ACM SIGGRAPH Conference on Motion in Games*, where I received the **Best Paper Award**[1] in 2017, by the ACM Symposium on Interactive 3D Graphics and Games, where I received the **Best Student Paper Award** in 2025[2], and by ACM Symposium on Computer Animation, where I received the **Best Presentation Paper Award** in 2025[3]. Overall, my research contributions are in **simulation of physics-based AI agents**[4,5,6,7,8,9,10,1,11], **deep learning and reinforcement learning**[12,13,10,14,15,16,17], and **architectural scene analysis and synthesis**,[14,18,19,17,20,16,21,22,23,15]. I work on:

1. 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.

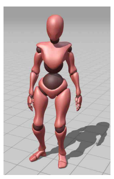

Agent

Agents in a simulation run through a continuous cycle of sensing and acting, where each cycle correlates 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.

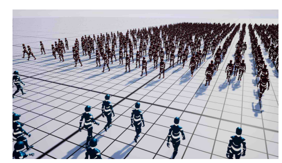

Multi-Agent Crowds

Physics-based Crowd Simulation

One core problem in multi-agent dynamics is to simulate dense, large crowds interactively. We proposed to use Position-Based Dynamics (PBD), in conjunction with novel geometry-based crowd dynamics constraints to simulate up to 100,000+ agents in real time[1]. We achieved this using GPU parallelization of the PBD solver (Fig. 1).

Later, we proposed additional novel PBD constraint for simulating crowds moving in groups and formations[6], and for simulating agents that are not disc-shaped, but with capsules, since a human figure projection on the ground plane

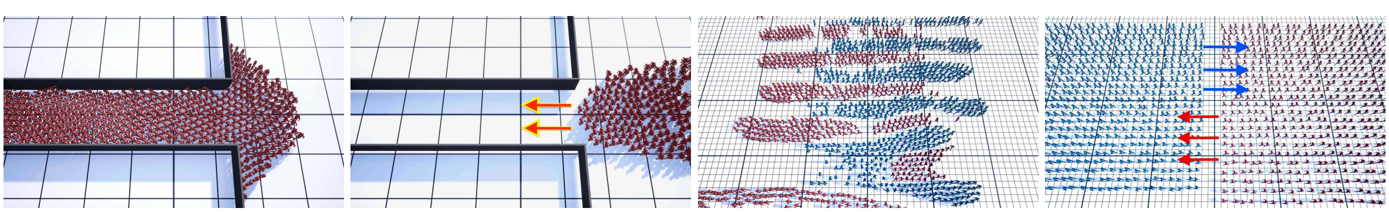

(a) A group of agents passing through a narrow corridor. **(b)** Two groups exchanging positions.

Figure 1: Emergent phenomena in crowd simulation. Our results[1] recreate different phenomena observed in pedestrian crowds, without any scripting or other user-directed control. These phenomena include: (a) clogging and arching near bottlenecks,[24] (b) groups self-organizing into lanes,[25] with stable and real-time performance for 100,000+ individuals. These results are made possible by framing the crowd motion as a constraint optimization problem, which is then efficiently solved using a GPU.

approximates a capsule more than a disc[5]. Such capsule approximation for agents allows more captures angular motion and turns more accurately. In addition to PBD, my group has also researched how other physics-based simulation approaches such as Projective Dynamics can be reformulated to simulate crowds[26]. We demonstrated that such approach leads to quick resolution of dense opposing crowd flows, which is notoriously difficult to compute with previous work.

Learning-based Crowd Simulation

Deterministic physics-based rules often fail to capture the somewhat stochastic nature of crowd dynamics. To attach this challenge, we proposed to learn the most-preferred crowd simulation behaviors from crowd-sourced votes via **Bayesian optimization**. 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 is preferred to users. We demonstrate our method's robustness in multiple scenarios and metrics, where we show it is superior compared to alternate policies and prior work[13].

Other Crowds Research

Outside of pedestrian crowds, we also simulated 3D crowds[9], including robot-like agents and ant swarm aggregations[3,27]. Understanding the dynamics of ant swarms and bacteria colonies is important for many applications from health, traffic to manufacturing. This work modeled ants as self-actuating PBD particles, which spontaneous form elastic links and also maintain fluid-like properties such as maintaining a certain density within an aggregation (Figure 3).

Other work includes innovative way finding, and accelerating the rendering a crowd simulation. To achieve higher computational rates of crowd simulation, we proposed a **Perceptual Foveated Rendering** approach, where instead of rendering each agent with the same fidelity, we render agents that a viewer focuses on with a higher resolution, and crowds not in the area of Fovea with lower resolution (Figure 3). This is done by modifying the animation refresh rate of agents limbs[2]. In regards to way finding, we proposed a **path planing** method that allows groups of agents to maintain their formation while navigating an obstructed scene. This is done by constructing a vector field encodes the supported formations, and allows switching to alternative formations requiring less clearance[7].

My group has been working on extending these results to multiple crowd benchmarks,[11] including agent navigation in 3D spaces,[9], formation control, and robotics applications.[13] In our latest results, we show how **deep reinforcement learning** is valuable as a motion suggestion engine where the simulation mechanics are controlled by our method.[1]

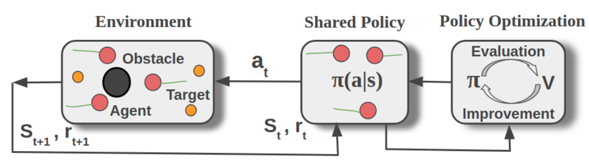

Figure 2: RL for multi-agent navigation[13].

2. Scene Synthesis and Style Classification

Virtual worlds are **challenging to construct**, requiring professionals with extensive training in 3D modeling software. My lab's research accelerates 3D content creation with contributions in the following:

Layout Synthesis

We have developed computational methods for creating spaces using **fast optimization** and **generative AI**[21,17]

In the **optimization** approach, users provide a 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

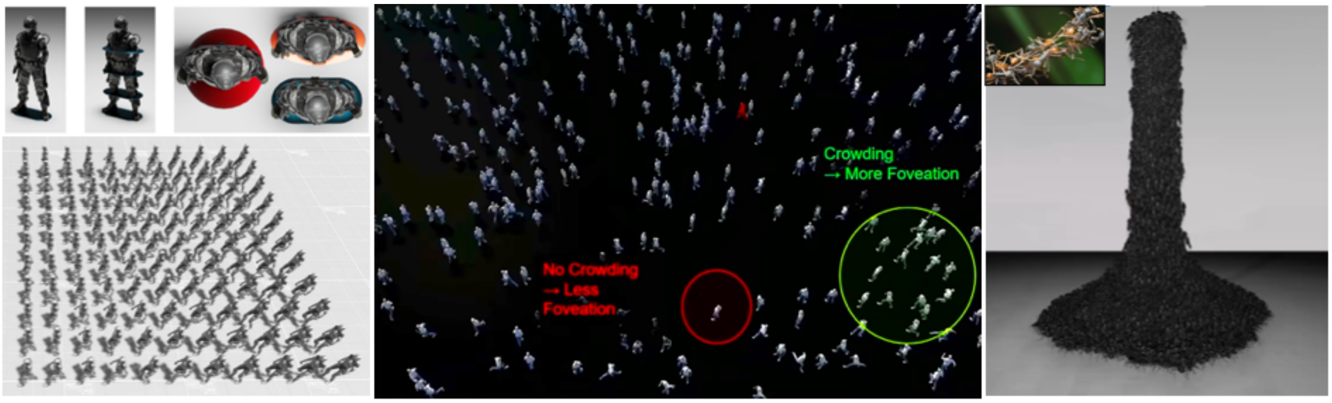

Figure 3: Recent multi-agent research of our lab includes simulating crowds as capsule-shaped particles which more accurately capture the pedestrian profile than discs, simulating agents moving in formation (left), accelerating rendering of crowd simulations using foveated rendering (middle), and simulating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

deformable body simulation. Both layouts and deformable bodies can be described by geometric constraints, which should be satisfied for a layout to be realistic, or for the movement of a body to be plausible. Both cases can be tackled using continuous optimization procedures. A layout synthesis solution is then an arrangement that satisfies such constraints[28]. 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.[29] This work was later used to create a dataset of 3D spaces, which is **widely used today** for multiple deep learning tasks.[30]

We also developed a **deep-learning generative transformer** approach for scene synthesis[17], 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. Our main innovation contrary to other transformer-based scene synthesis (and data-driven) work is that our synthesis **improves on the quality of rooms available in the original dataset**. This is done by using additional terms in the loss function, which encode ergonomic and geometric common-sense principles, such as the a table should be within reach of a person sitting.

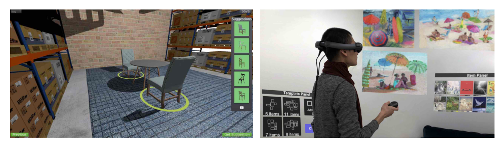

System for designing style-compatible spaces.[19,15]

Most recently, 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[19,18].

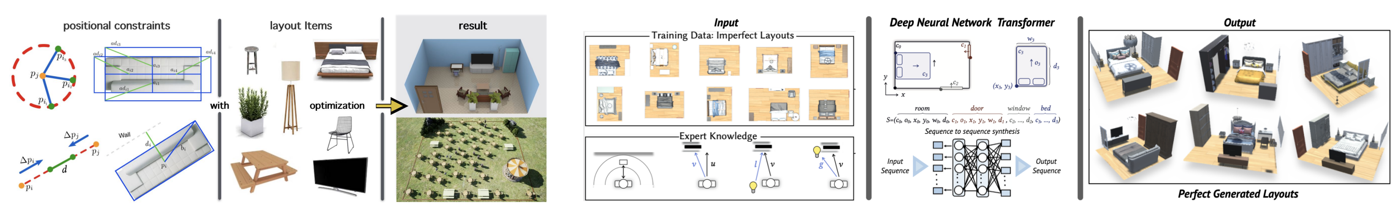

Figure 4: Left: Our results[21] in 3D scene synthesis are faster by an order of magnitude than previous probabilistic methods and allow **real-time, interactive synthesis of scenes that were previously intractable**. Given input positional constraints and selected layout items, the algorithm rapidly 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 Deep Neural Network Transformer architecture.[17] The synthesized layouts integrate desirable properties not present in the original dataset.

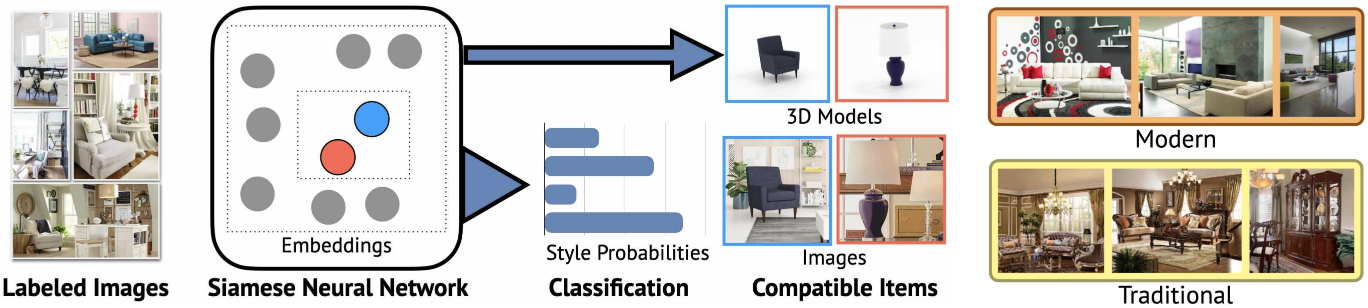

Figure 5: Left: Given style-labeled images of 3D objects and rooms, we trained a Siamese network to predict image styles.[15] After the network is trained, the stylistic compatibility between 3D objects is measured by selecting the nearest neighbors within the network's embeddings for each image displaying the 3D object. **Right**: Examples of style labels. Our work can predict the predominant style of a space, which is useful for content creation, design, and retail.[14]

Virtual Content Compatibility

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 used a siamese[31], imaged-based classification deep-learning approach for estimating 3D object style.[15] To accomplish this, we used (i) images where each object appears in multiple "styled" backgrounds (Figure 5), which is annotated with multiple style votes by experts, in contrast to other work that relies on images of objects presented in an isolated context[32], or solely on object geometry[33], (ii) comparison labels that encode the style spectrum of each object in a relative manner. For example, a label indicates whether one table is more modern than another table.

Along this line of work, we also classified room styles with a similar neural network[14], furnishing rooms with stylistically similar objects in VR[19], and how room clutter affects the perception of space[34].

References

[1] 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**.

[2] 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.

[3] 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.

[4] Florin-Vladimir Stancu, Tomer Weiss, and Rafael Kuffner Dos Anjos. Foveated animations for efficient crowd simulation. In *ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games*, I3D '25, 2025. **Best Student Paper Award**.

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

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

[7] 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.

[8] Tomer Weiss. Implementing position-based real-time simulation of large crowds. In *2019 IEEE International Conference on Artificial Intelligence and Virtual Reality (AIVR)*, pages 306–3061. IEEE, 2019.

[9] 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.

[10] M.Nakada, H.Chen, T.Zhao, T.Weiss, and D.Terzopoulos. Deep learning of neuromuscular control for pointing tasks in a biomechanical human model. In *ACM Transactions on Graphics (SIGGRAPH)*, 2018.

[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 to simulate crowds with crowds. In *ACM SIGGRAPH 2023 Posters*, pages 1–2. ACM New York, NY, USA, 2023.

[13] 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.

[14] 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)*, pages 267–2673. IEEE, 2019.

[15] 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.

[16] 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.

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

[18] 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.

[19] 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.

[20] 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.

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

[25] 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.

[26] Tomer Weiss. Projective multi-agent dynamics. In *Proceedings of the 2025 18th ACM SIGGRAPH Conference on Motion, Interaction, and Games*, pages 1–8, 2025.

[27] 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.

[28] Matthias Müller, Bruno Heidelberger, Marcus Hennix, and John Ratcliff. Position based dynamics. *Virtual Reality Interactions and Physical Simulations (VRIPHYS)*, 18(2):109–118, 2007.

[29] 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.

[30] 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.

[31] Sean Bell and Kavita Bala. Learning visual similarity for product design with convolutional neural networks. *ACM transactions on graphics (TOG)*, 34(4):1–10, 2015.

[32] Isaak Lim, Anne Gehre, and Leif Kobbelt. Identifying style of 3d shapes using deep metric learning. In *Computer graphics forum*, volume 35, pages 207–215. Wiley Online Library, 2016.

[33] Tianqiang Liu, Aaron Hertzmann, Wilmot Li, and Thomas Funkhouser. Style compatibility for 3d furniture models. *ACM Transactions on Graphics (TOG)*, 34(4):1–9, 2015.

[34] 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.