SIGGRAPH 2022
VANCOUVER+ 8-11 AUG

CCP: CONFIGURABLE CROWD PROFILES

Andreas Panayiotou[1], Theodoros Kyriakou[1], Marilena Lemonari[2], Yiorgos Chrysanthou[1,2], and Panayiotis Charalambous[1]
[1] CYENS – Centre of Excellence, Nicosia, Cyprus, [2] University of Cyprus, Nicosia, Cyprus

Source Code
Trained Policies
Video

ABSTRACT

Diversity among agents' behaviors and heterogeneity in virtual crowds in general, is an important aspect of crowd simulation as it is crucial to the perceived realism and plausibility of the resulting simulations. Most of the existing systems optimize for specific behaviors such as goal seeking, and neglect to consider other behaviors and how these interact together to form diverse agent profiles. In this paper, we present a **Reinforcement Learning based** framework for **learning multiple agent behaviors concurrently**. We optimize the agent policy by varying the importance of the **selected behaviors** (*GOAL SEEKING*, *COLLISION AVOIDANCE*, *GROUPING*, *and INTERACTION WITH POIs*) while training; essentially, we have a **reward function that changes dynamically during training**. The importance of each separate sub-behavior **is added as input to the policy**, resulting in the development of a **single model** capable of enabling dynamic **run-time manipulation of agent profiles**; thus, allowing configurable profiles. Our system provides users with the ability to control and **mix agent behaviors** thus creating personality profiles and **assign different profiles to groups of agents**. Moreover, we demonstrate that interestingly the proposed **model generalizes to situations** not seen in the training data such as (a) higher density crowds, (b) behavior weights outside the training intervals and (c) to scenes with more intricate environment layouts.

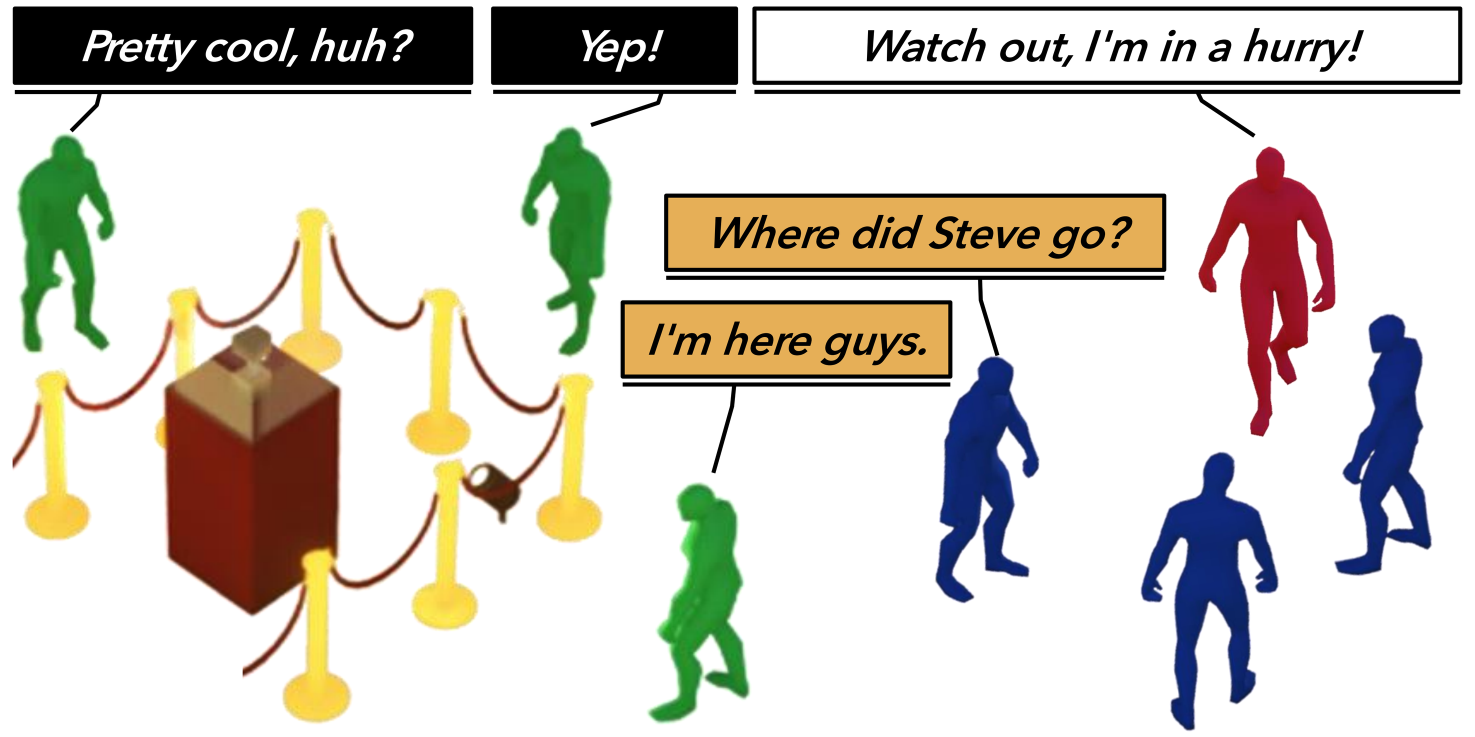

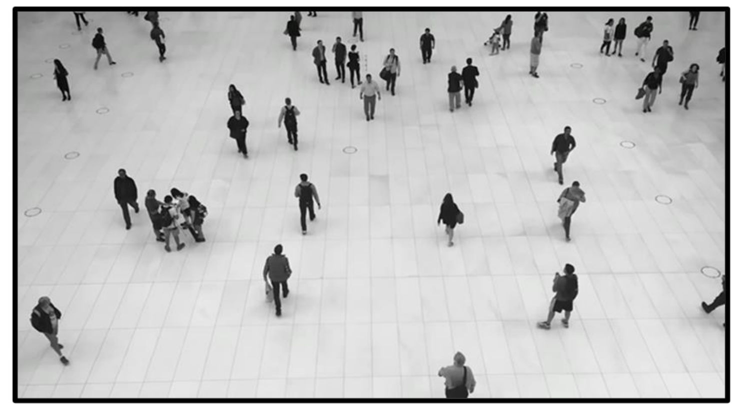

METHODOLOGY

Training Scene - Observations - Actions

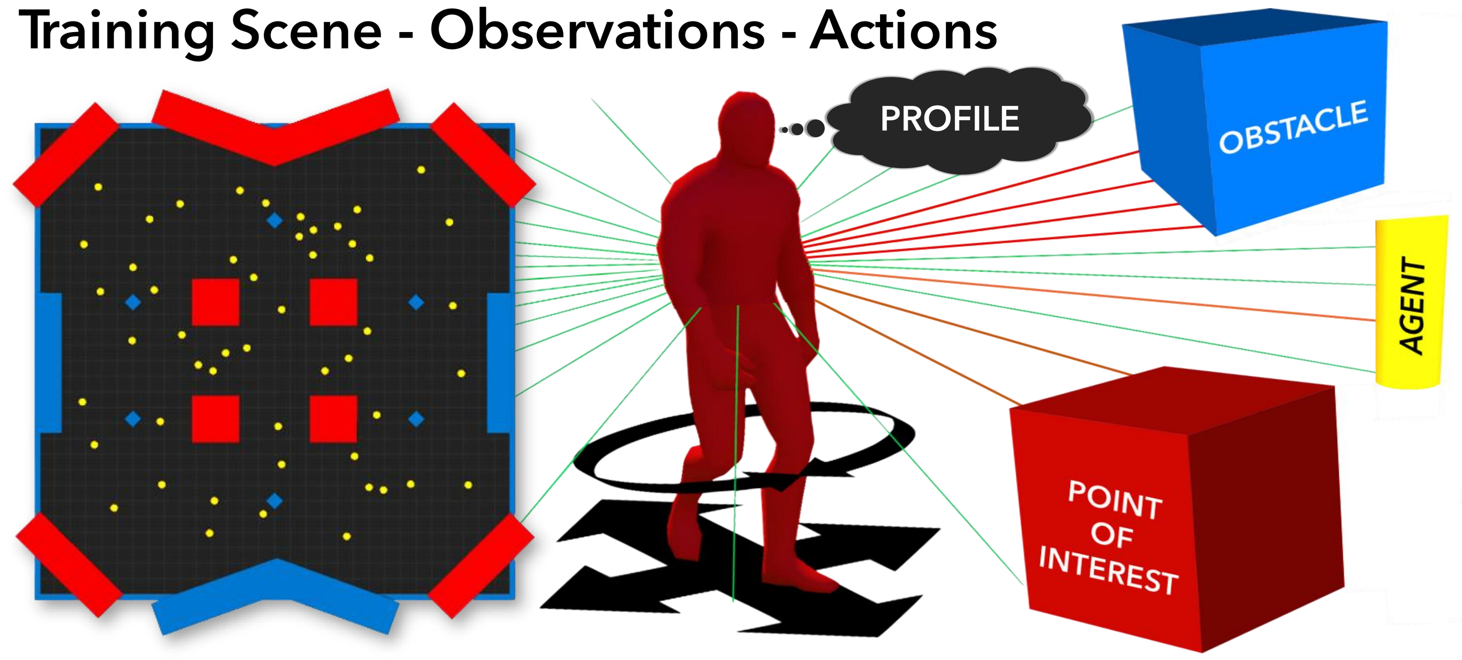

Parameterized Reward Function

A set of values $\{w_g, w_{ca}, w_{gr}, w_i\}$ define a profile of an agent.

Event	Symbol	Base Reward	Weight	Dense
Reached Goal	R_g	+1.0	w_g	N
Agent Collision	R_{ca}	-.01	w_{ca}	N
Obstacle Collision	R_{co}	-.5	w_{ca}	N
Towards Goal	R_{gt}	+.00075	w_g	Y
Away from Goal	R_{ga}	-.00025	w_g	Y
In Group	R_{gr}	+.001	w_{gr}	Y
Interacting	R_i	+.001	w_i	Y
Living Penalty	R_l	-.00015	w_g	Y

Training Strategy

We use a simple Fully Connected Neural Network (FCNN) and the Proximal Policy Optimization (PPO) algorithm. We utilize a curriculum-based method during training; we gradually increase the difficulty of the environment and behaviors the agents have to face.

$$R^t = \boxed{w_g}(R_g + R_{gt} + R_{ga} + R_l) + \boxed{w_{ca}}(R_{ca} + R_{co}) + \boxed{w_{gr}}R_{gr} + \boxed{w_i}R_i$$

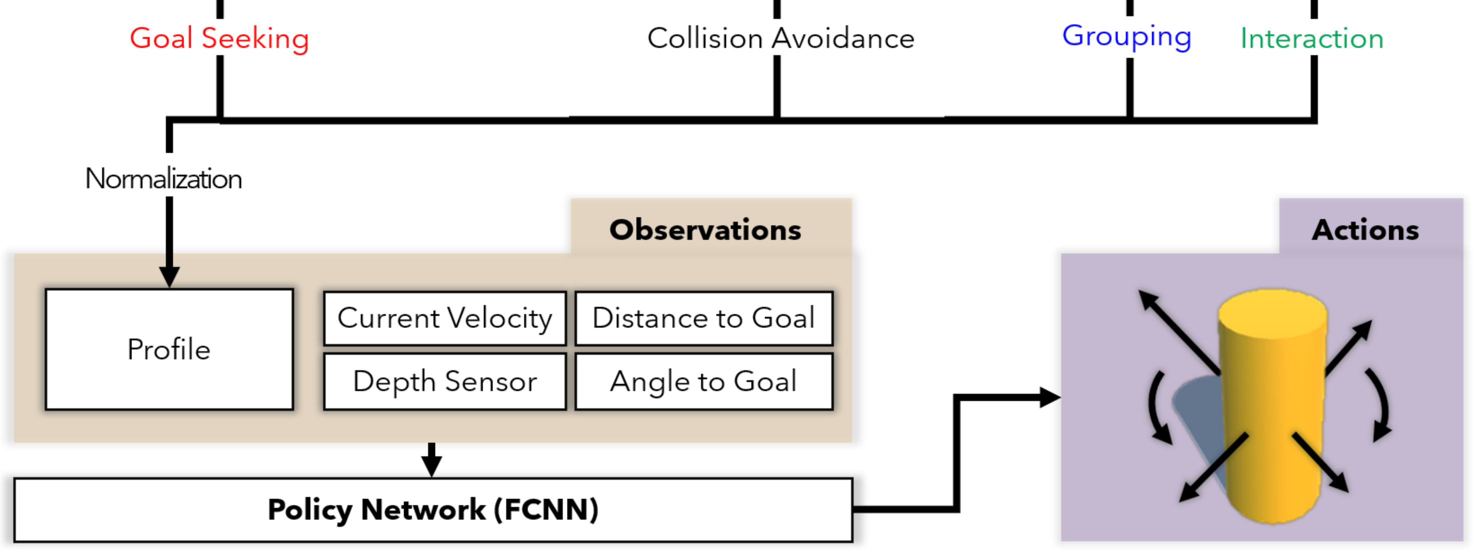

EXPERIMENTS AND EVALUATION

We demonstrate our framework's practicality when designing and authoring realistic scenes. We showcase one case study, a museum exhibition, which represents a large, abstract complex environment.

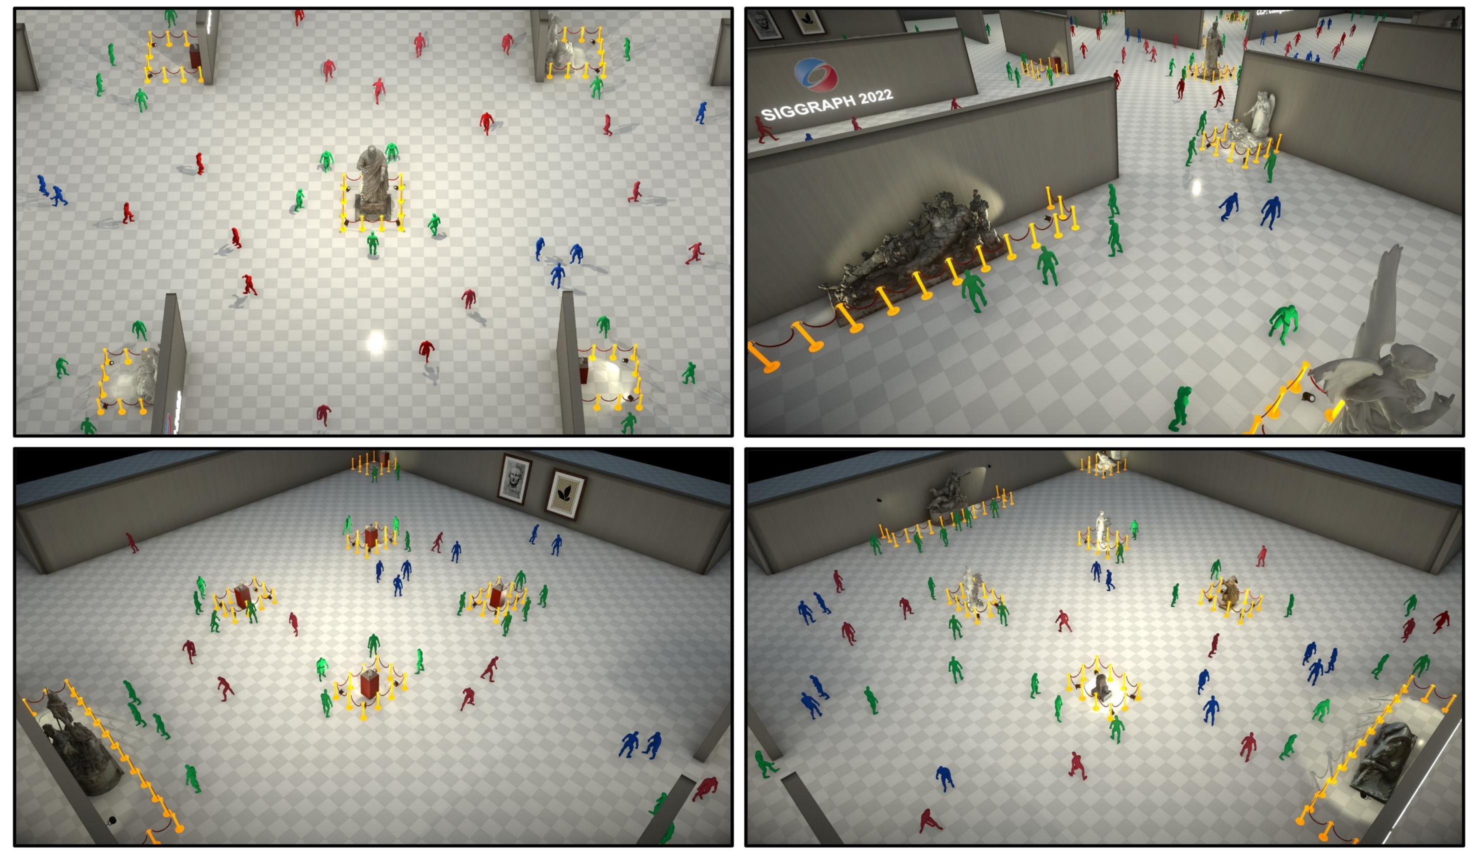

Comparison with Power Law and Real-World Data

We generate the Fundamental Diagram of our method as compared to simulation data from Power Law and real-world data. In all cases, we observe that higher densities lead to slower agents.

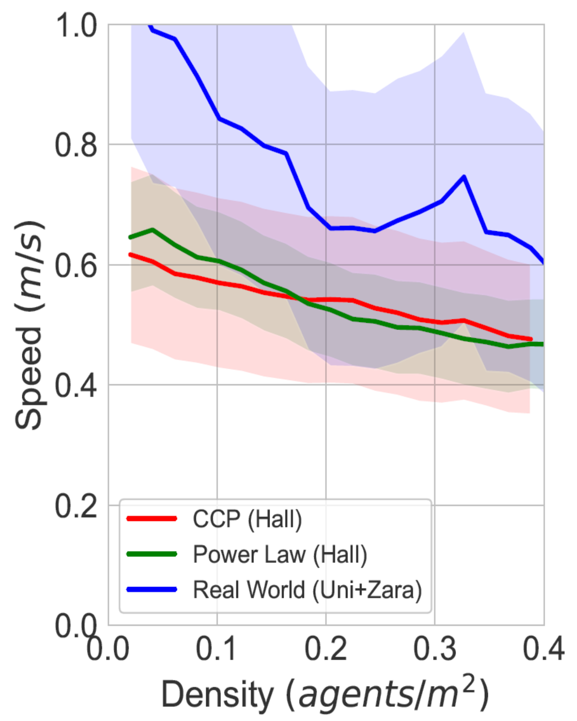

Simulation Environments

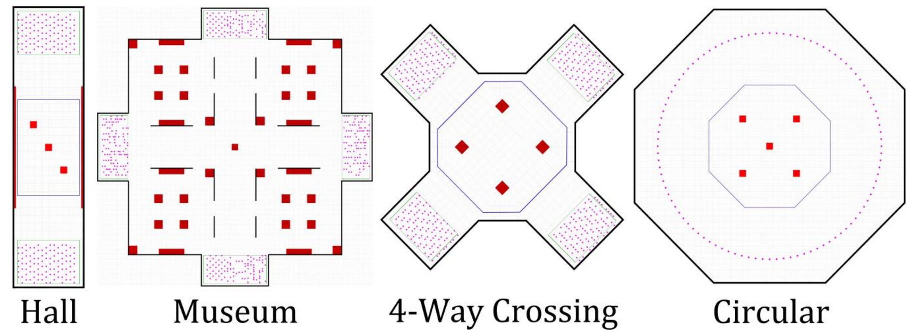

Effect of weight changes for the Interaction Behavior

As w_i increases, agents move slower and keep short distances to POIs.

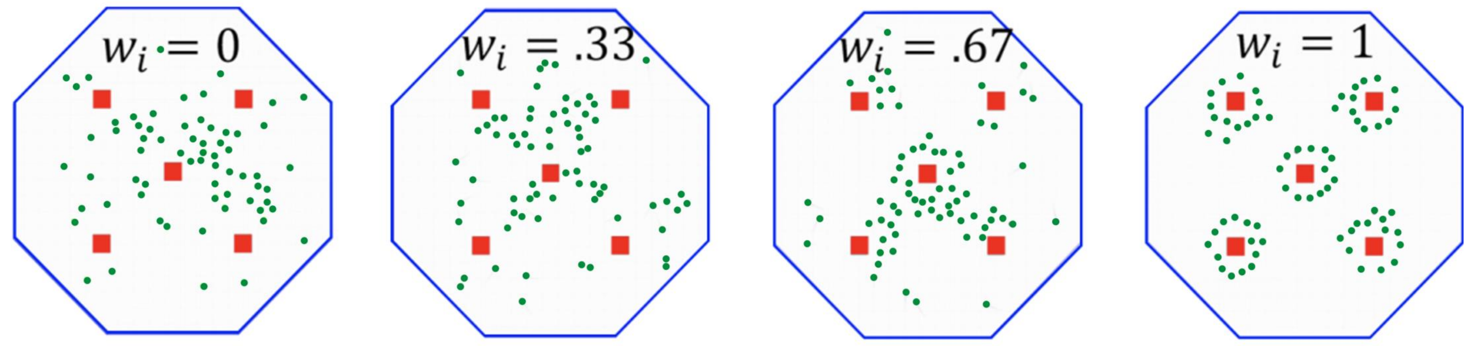

Weight Sensitivity in the Hall environment

To test the effect of weights, we vary the value of a single weight in the range [0-1] while keeping the values of the other weights at 0.5.

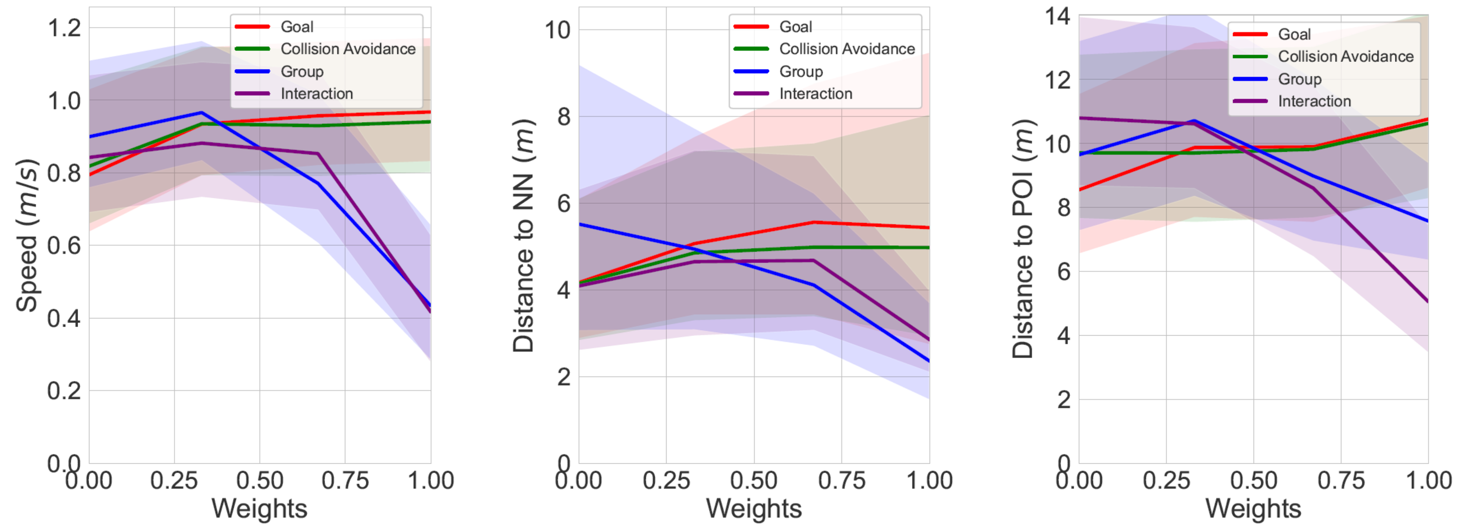

Density Sensitivity in the Crossing environment

We select set of weights reflecting behavior's effect profoundly.

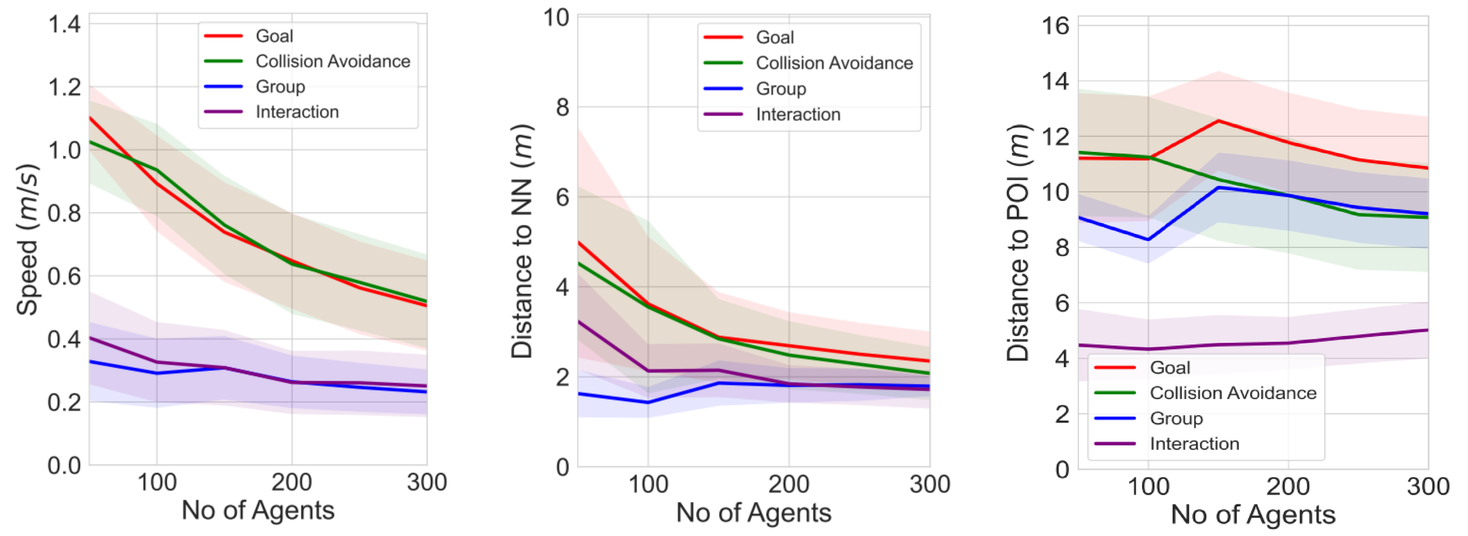

CONCLUSION

We present a Reinforcement Learning based training method, capable of successfully learning multiple behaviors concurrently. Our system is useful for creating heterogeneous crowds, enabling users to assign behaviors and have intuitive control at run-time. The generalizable qualities of our approach, both in terms of crowd size and environmental layout, further solidify our system as an efficient and practical learning-based tool for simulating crowds.

ACKNOWLEDGEMENTS

This project has received funding from the European Union's Horizon 2020 Research and Innovation Programme under Grant Agreement No 739578 and the Government of the Republic of Cyprus through the Deputy Ministry of Research, Innovation and Digital Policy. This work has received funding from the European Union's Horizon2020 research and innovation programme under the Marie Skłodowska Curie grant agreement No 860768 (CLIPE project).

CYENS CENTRE OF EXCELLENCE · University of Cyprus · clipe · Horizon 2020 European Union Funding for Research & Innovation

CONTACT US: A.PANAYIOTOU@CYENS.ORG.CY | P.CHARALAMBOUS@CYENS.ORG.CY