REPORT DOCUMENTATION PAGE

Form Approved OMB NO. 0704-0188

1. REPORT DATE (DD-MM-YYYY)	2. REPORT TYPE	3. DATES COVERED (From - To)
24-07-2022	Final Report	7-Jan-2019 - 6-Mar-2022

4. TITLE AND SUBTITLE
Final Report: Efficient Computational Models for Simulating Large-Scale Heterogeneous Crowds

5a. CONTRACT NUMBER
W911NF-19-1-0069

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER
611102

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

6. AUTHORS

7. PERFORMING ORGANIZATION NAMES AND ADDRESSES
University of Maryland - College Park
The University of Maryland
Office of Research Administration
College Park, MD 20742 -5141

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS (ES)
U.S. Army Research Office
P.O. Box 12211
Research Triangle Park, NC 27709-2211

10. SPONSOR/MONITOR'S ACRONYM(S)
ARO

11. SPONSOR/MONITOR'S REPORT NUMBER(S)
73554-CS-H.8

12. DISTRIBUTION AVAILIBILITY STATEMENT
Approved for public release; distribution is unlimited.

13. SUPPLEMENTARY NOTES
The views, opinions and/or findings contained in this report are those of the author(s) and should not contrued as an official Department of the Army position, policy or decision, unless so designated by other documentation.

14. ABSTRACT

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	15. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			Ming Lin
UU	UU	UU	UU		19b. TELEPHONE NUMBER 301-405-2662

RPPR Final Report
as of 25-Jul-2022

Agency Code: 21XD

Proposal Number: 73554CSH

Agreement Number: W911NF-19-1-0069

INVESTIGATOR(S):

Name: Ph.D. Ming C. Lin
Email: lin@umd.edu
Phone Number: 3014052662
Principal: Y

Organization: **University of Maryland - College Park**
Address: The University of Maryland, College Park, MD 207425141
Country:

DUNS Number: 790934285
EIN: 526002033
Report Date: 06-Jun-2022
Date Received: 24-Jul-2022
Final Report for Period Beginning 07-Jan-2019 and Ending 06-Mar-2022
Title: Efficient Computational Models for Simulating Large-Scale Heterogeneous Crowds
Begin Performance Period: 07-Jan-2019
End Performance Period: 06-Mar-2022
Report Term: 0-Other

Submitted By: Ph.D. Ming Lin
Email: lin@umd.edu
Phone: (301) 405-2662

Distribution Statement: 1-Approved for public release; distribution is unlimited.

STEM Degrees: 5 **STEM Participants:** 1

Major Goals: - Real-time simulation and visualization of massive crowd flows using a novel hybrid representation of discrete-continuum paradigm;

- Continuum modeling of pedestrian dynamics that can be used to simulate the macroscopic and complex crowd phenomena;

- Accurate pedestrian models based on elliptical representations and local navigation algorithms;

- Modeling of intricate pedestrian dynamics for accurate movement on uneven terrains;

- Data-driven crowd modeling based on machine learning and Bayesian inference; - Applications to battlefield simulation,emergency response and training, engineering design evaluation, and crowd monitoring;

-Statistical techniques for validation and verification of crowd simulation algorithms.

Accomplishments: DenseCAvoid: Real-time Navigation in Dense Crowds using Anticipatory Behaviors:
We present DenseCAvoid, a novel algorithm for navigating a robot through dense crowds and avoiding collisions by anticipating pedestrian behaviors. Our formulation uses visual sensors and a pedestrian trajectory prediction algorithm to track pedestrians in a set of input frames and compute bounding boxes that extrapolate to the pedestrian positions ina future time. Our hybrid approach combines this trajectory prediction with a Deep Reinforcement Learning-based collision avoidance method to train a policy to generate smoother, safer, and more robust trajectories during runtime. We train our policy in realistic 3-D simulations of static and dynamic scenarios with multiple pedestrians. In practice, our hybrid approach generalizes well to unseen, real-world scenarios and can navigate a robot through dense crowds (～1-2 humans per square meter) in indoor scenarios, including narrow corridors and lobbies.

Frozone: Freezing-Free, Pedestrian-Friendly Navigation in Human Crowds:
We present Frozone, a novel algorithm to deal with the Freezing Robot Problem (FRP) that arises when a robot navigates through dense scenarios and crowds. Our method senses and explicitly predicts the trajectories of pedestrians and constructs a Potential Freezing Zone (PFZ); a spatial zone where the robot could freeze or be obtrusive to humans. Our formulation computes a deviation velocity to avoid the PFZ, which also accounts for social constraints. Furthermore, Frozone is designed for robots equipped with sensors with a limited sensing range and field of view. We ensure that the robot's deviation is bounded, thus avoiding sudden angular motion which

could lead to the loss of perception data of the surrounding obstacles. We have combined Frozone with a Deep Reinforcement Learning-based (DRL) collision avoidance method and use our hybrid approach to handle crowds of varying densities. Our overall approach results in smooth and collision-free navigation in dense environments. We have evaluated our method's performance in simulation and on real differential drive robots in challenging indoor scenarios.

Realtime Collision Avoidance for Mobile Robots in Dense Crowds using Implicit Multi-sensor Fusion and Deep Reinforcement Learning:
We present a novel high fidelity 3-D simulator that significantly reduces the sim-to-real gap for collision avoidance in dense crowds using Deep Reinforcement Learning (DRL). Our simulator models realistic crowd and pedestrian behaviors, along with friction, sensor noise and delays in the simulated robot model. We also describe a technique to incrementally control the randomness and complexity of training scenarios to achieve better convergence and generalization capabilities. We demonstrate the effectiveness of our simulator by training a policy that fuses data from multiple perception sensors such as a 2-D lidar and a depth camera to detect pedestrians and computes smooth, collision-free velocities. Our novel reward function and multi-sensor formulation results in smooth and unobtrusive navigation. We observe that our algorithm outperforms prior dynamic navigation techniques in terms of metrics such as success rate, trajectory length, mean time to goal, and smoothness.

ACSEE: Antagonistic Crowd Simulation Model with Emotional Contagion and Evolutionary Game Theory:
Antagonistic crowd behaviors are often observed in cases of serious conflict. Antagonistic emotions between different opposing groups and the way they spread through contagion in a crowd are important causes of such behaviors. Moreover, games between different opposing groups determine the level of crowd violence. We present an antagonistic crowd simulation model, ACSEE, which is integrated with antagonistic emotional contagion and evolutionary game theories. Our approach models the antagonistic emotions between agents in different roles using two components: mental emotion and external emotion. We combine enhanced susceptible-infectious-susceptible (SIS) and game approaches to evaluate the role of antagonistic emotional contagion in crowd violence. Our evolutionary game theoretic approach incorporates antagonistic emotional contagion through deterrent force, which is modelled by a mixture of emotional and physical forces defeating the opponents. We evaluate our approach on real-world scenarios consisting of different kinds of agents. We also compare the simulated crowd behaviors with real-world crowd videos and use our approach to predict the trends of crowd movements.

GamePlan: Game-Theoretic Multi-Agent Planning with Human Drivers at Intersections, Roundabouts, and Merging (Published in RAL/ICRA 2022)
We present a new method for multi-agent planning involving human drivers and autonomous vehicles (AVs) in unsignaled intersections, roundabouts, and during merging. In multi-agent planning, the main challenge is to predict the actions of other agents, especially human drivers, as their intentions are hidden from other agents. Our algorithm uses game theory to develop a new auction, called GamePlan, that directly determines the optimal action for each agent based on their driving style (which is observable via commonly available sensors). GamePlan assigns a higher priority to more aggressive or impatient drivers and a lower priority to more conservative or patient drivers; we theoretically prove that such an approach is game-theoretically optimal prevents collisions and deadlocks. We compare our approach with prior state-of-the-art auction techniques including economic auctions, time-based auctions (first-in first-out), and random bidding and show that each of these methods result in collisions among agents when taking into account driver behavior. We compare with methods based on DRL, deep learning, and game theory and present our benefits over these approaches. Finally, we show that our approach can be implemented in the real-world with human drivers.

STCrowd: A Multimodal Dataset for Pedestrian Perception in Crowded Scenes (IEEE CVPR 2022)
Accurately detecting and tracking pedestrians in 3D space is challenging due to large variations in rotations, poses and scales. The situation becomes even worse for dense crowds with severe occlusions. However, existing benchmarks either only provide 2D annotations, or have limited 3D annotations with low-density pedestrian distribution, making it difficult to build a reliable pedestrian perception system especially in crowded scenes. To better evaluate pedestrian perception algorithms in crowded scenarios, we introduce a large-scale multimodal dataset, STCrowd. Specifically, in STCrowd, there are a total of 219K pedestrian instances and 20 persons per frame on average, with various levels of occlusion. We provide synchronized LiDAR point clouds and camera images as well as their corresponding 3D labels and joint IDs. STCrowd can be used for various tasks, including LiDAR-only, image-only, and sensor-fusion based pedestrian detection and tracking. We provide baselines for most of the tasks. In addition, considering the property of sparse global distribution and density-varying local distribution of pedestrians, we further propose a novel method, Density-aware Hierarchical heatmap Aggregation (DHA), to enhance pedestrian perception in crowded scenes. Extensive experiments show that our new method achieves

state-of-the-art performance on the STCrowd dataset, especially on cases with severe occlusion. The dataset and code will be released to facilitate related research when the paper is published.

Training Opportunities: Training of undergraduate and graduate students.

Results Dissemination: Research results are published in top international scientific venues. Please see the project summary under "Accomplished" and "Products".

Honors and Awards: Dinesh Manocha was inducted into ACM SIGGRAPH Academy in 2019.

Ming Lin was selected to be UMD College Park Distinguished University Professor in 2019.

Dinesh Manocha was selected to be UMD College Park Distinguished University Professor in 2020.

Ming Lin was awarded Washington Academy of Science Distinguished Career Award in Computer Science in 2020.

Ming Lin was inducted into ACM SIGGRAPH Academy in 2020.

Ming Lin received the Spirit of IEEE Computer Society Award in 2020.

Ming Lin was named Dr. Barry Mersky and Capital One E-Nnovate Endowed Professor in 2021.

Ming Lin was selected to be among 50 Women in Robotics You Need to Know About in 2021.

Ming Lin was inducted as the Fellow of National Academy of Inventors in 2022.

Ming Lin was inducted as the Fellow of IEEE VR Academy in 2022.

ARO funding was acknowledged.

Protocol Activity Status:

Technology Transfer: The resulting software systems are made publicly available for download and have been integrated into various open-source software systems, including ROS.

PARTICIPANTS:

Participant Type: PD/PI
Participant: Ming C Lin
Person Months Worked: 6.00 **Funding Support:**
Project Contribution:
National Academy Member: N

Participant Type: Co PD/PI
Participant: Dinesh Manocha
Person Months Worked: 6.00 **Funding Support:**
Project Contribution:
National Academy Member: N

Participant Type: Graduate Student (research assistant)
Participant: Rohan Chandra
Person Months Worked: 3.00 **Funding Support:**
Project Contribution:
National Academy Member: N

Participant Type: Graduate Student (research assistant)
Participant: Uttaran Bhattacharya
Person Months Worked: 12.00 **Funding Support:**
Project Contribution:
National Academy Member: N

Participant Type: Graduate Student (research assistant)
Participant: Trisha Mittal
Person Months Worked: 12.00 **Funding Support:**
Project Contribution:
National Academy Member: N

Participant Type: Graduate Student (research assistant)
Participant: Adarsh Sathyamoorthy
Person Months Worked: 12.00 **Funding Support:**
Project Contribution:
National Academy Member: N

Participant Type: Graduate Student (research assistant)
Participant: Sang Son
Person Months Worked: 6.00 **Funding Support:**
Project Contribution:
National Academy Member: N

Participant Type: Undergraduate Student
Participant: Laura Zheng
Person Months Worked: 12.00 **Funding Support:**
Project Contribution:
National Academy Member: N

ARTICLES:

Publication Type: Journal Article Peer Reviewed: Y **Publication Status:** 1-Published
Journal: Proceedings of 2019 International Conference on Robotics and Automation (ICRA)
Publication Identifier Type: DOI Publication Identifier: 10.1109/ICRA.2019.8794239
Volume: Issue: First Page #:
Date Submitted: 9/30/19 12:00AM Date Published: 5/20/19 4:00AM
Publication Location:
Article Title: ADAPS: Autonomous Driving Via Principled Simulations
Authors: Weizi Li ; David Wolinski ; Ming C. Lin
Keywords: Autonomous Driving
Abstract: Autonomous driving has gained significant advancements in recent years. However, obtaining a robust control policy for driving remains challenging as it requires training data from a variety of scenarios, including rare situations (e.g., accidents), an effective policy architecture, and an efficient learning mechanism. We propose ADAPS for producing robust control policies for autonomous vehicles. ADAPS consists of two simulation platforms in generating and analyzing accidents to automatically produce labeled training data, and a memoryenabled hierarchical control policy. Additionally, ADAPS offers a more efficient online learning mechanism that reduces the number of iterations required in learning compared to existing methods such as DAGGER [1]. We present both theoretical and experimental results. The latter are produced in simulated environments, where qualitative and quantitative results are generated to demonstrate the benefits of ADAPS.
Distribution Statement: 3-Distribution authorized to U.S. Government Agencies and their contractors
Acknowledged Federal Support: **Y**

Publication Type: Journal Article Peer Reviewed: Y **Publication Status:** 1-Published
Journal: IEEE Robotics and Automation Letters
Publication Identifier Type: DOI Publication Identifier: 10.1109/LRA.2020.2996593
Volume: 5 Issue: 3 First Page #: 4352
Date Submitted: 10/11/20 12:00AM Date Published: 5/22/20 4:00AM
Publication Location:
Article Title: Frozone: Freezing-Free, Pedestrian-Friendly Navigation in Human Crowds
Authors: Adarsh Jagan Sathyamoorthy, Utsav Patel, Tianrui Guan, Dinesh Manocha
Keywords: Avoidance
Abstract: We present Frozone, a novel algorithm to deal with the Freezing Robot Problem (FRP) that arises when a robot navigates through dense scenarios and crowds. Our method senses and explicitly predicts the trajectories of pedestrians and constructs a Potential Freezing Zone (PFZ); a spatial zone where the robot could freeze or be obtrusive to humans. Our formulation computes a deviation velocity to avoid the PFZ, which also accounts for social constraints. Furthermore, Frozone is designed for robots equipped with sensors with a limited sensing range and field of view. We ensure that the robot's deviation is bounded, thus avoiding sudden angular motion which could lead to the loss of perception data of the surrounding obstacles. We have combined Frozone with a Deep Reinforcement Learning-based (DRL) collision avoidance method and use our hybrid approach to handle crowds of varying densities. Our overall approach results in smooth and collision-free navigation in dense environments.
Distribution Statement: 2-Distribution Limited to U.S. Government agencies only; report contains proprietary info
Acknowledged Federal Support: **Y**

Publication Type: Journal Article Peer Reviewed: Y **Publication Status:** 1-Published
Journal: IEEE International Conference on Robotics and Automation
Publication Identifier Type: Other Publication Identifier: https://doi.org/10.48550/arXiv.2002.03038
Volume: 1 Issue: 1 First Page #: 1
Date Submitted: 7/24/22 12:00AM Date Published: 5/31/20 4:00PM
Publication Location:
Article Title: DenseCAvoid: Real-time Navigation in Dense Crowds using Anticipatory Behaviors
Authors: Adarsh Jagan Sathyamoorthy, Jing Liang, Utsav Patel, Tianrui Guan, Rohan Chandra, Dinesh Manocha
Keywords: Navigation
Abstract: We present DenseCAvoid, a novel algorithm for navigating a robot through dense crowds and avoiding collisions by anticipating pedestrian behaviors. Our formulation uses visual sensors and a pedestrian trajectory prediction algorithm to track pedestrians in a set of input frames and compute bounding boxes that extrapolate to the pedestrian positions ina future time. Our hybrid approach combines this trajectory prediction with a Deep Reinforcement Learning-based collision avoidance method to train a policy to generate smoother, safer, and more robust trajectories during runtime. We train our policy in realistic 3-D simulations of static and dynamic scenarios with multiple pedestrians. In practice, our hybrid approach generalizes well to unseen, real-world scenarios and can navigate a robot through dense crowds (?1-2 humans per square meter) in indoor scenarios, including narrow corridors and lobbies.
Distribution Statement: 2-Distribution Limited to U.S. Government agencies only; report contains proprietary info
Acknowledged Federal Support: **Y**

Publication Type: Journal Article Peer Reviewed: Y **Publication Status:** 1-Published
Journal: IEEE/RSJ International Conference on Intelligent Robots and Systems
Publication Identifier Type: Other Publication Identifier:
Volume: Issue: First Page #: 1
Date Submitted: 7/24/22 12:00AM Date Published: 10/25/20 8:00AM
Publication Location:
Article Title: Enhanced Transfer Learning for Autonomous Driving with Systematic Accident Simulation
Authors: Shivam Akhauri, Laura Zheng, Ming Lin.
Keywords: Self-driving car
Abstract: We present a new method for multi-agent planning involving human drivers and autonomous vehicles (AVs) in unsignaled intersections, roundabouts, and during merging. In multi-agent planning, the main challenge is to predict the actions of other agents, especially human drivers, as their intentions are hidden from other agents. Our algorithm uses game theory to develop a new auction, called GamePlan, that directly determines the optimal action for each agent based on their driving style (which is observable via commonly available sensors). GamePlan assigns a higher priority to more aggressive or impatient drivers and a lower priority to more conservative or patient drivers; we theoretically prove that such an approach is game-theoretically optimal prevents collisions and deadlocks. We compare our approach with prior state-of-the-art auction techniques including economic auctions, time-based auctions (first-in first-out), and random bidding and show that each of these methods result in co
Distribution Statement: 2-Distribution Limited to U.S. Government agencies only; report contains proprietary info
Acknowledged Federal Support: **Y**

Publication Type: Journal Article Peer Reviewed: Y **Publication Status:** 1-Published
Journal: IEEE Robotics and Automation Letters
Publication Identifier Type: DOI Publication Identifier: 10.1109/LRA.2022.3144516
Volume: 7 Issue: 2 First Page #: 2676
Date Submitted: 7/24/22 12:00AM Date Published: 4/1/22 4:00AM
Publication Location:
Article Title: GamePlan: Game-Theoretic Multi-Agent Planning With Human Drivers at Intersections, Roundabouts, and Merging
Authors: Rohan Chandra, Dinesh Manocha
Keywords: Multi-agent, planning
Abstract: We present a new method for multi-agent planning involving human drivers and autonomous vehicles (AVs) in unsignaled intersections, roundabouts, and during merging. In multi-agent planning, the main challenge is to predict the actions of other agents, especially human drivers, as their intentions are hidden from other agents. Our algorithm uses game theory to develop a new auction, called GamePlan, that directly determines the optimal action for each agent based on their driving style (which is observable via commonly available sensors). GamePlan assigns a higher priority to more aggressive or impatient drivers and a lower priority to more conservative or patient drivers; we theoretically prove that such an approach is game-theoretically optimal prevents collisions and deadlocks. We compare our approach with prior state-of-the-art auction techniques including economic auctions, time-based auctions (first-in first-out), and random bidding and show that each of these methods result in co
Distribution Statement: 1-Approved for public release; distribution is unlimited.
Acknowledged Federal Support: **Y**

Publication Type: Journal Article Peer Reviewed: Y **Publication Status:** 1-Published
Journal: IEEE CVPR Proceedings
Publication Identifier Type: DOI Publication Identifier: https://doi.org/10.48550/arXiv.2204.01026
Volume: 1 Issue: 1 First Page #: 1
Date Submitted: 7/24/22 12:00AM Date Published: 6/22/22 4:00AM
Publication Location: New Orlean, LA
Article Title: STCrowd: A Multimodal Dataset for Pedestrian Perception in Crowded Scenes
Authors: Peishan Cong, Xinge Zhu, Feng Qiao, Yiming Ren, Xidong Peng, Yuenan Hou, Lan Xu, Ruigang Yang,
Keywords: Pedestrian, perception
Abstract: We present a new method for multi-agent planning involving human drivers and autonomous vehicles (AVs) in unsignaled intersections, roundabouts, and during merging. In multi-agent planning, the main challenge is to predict the actions of other agents, especially human drivers, as their intentions are hidden from other agents. Our algorithm uses game theory to develop a new auction, called GamePlan, that directly determines the optimal action for each agent based on their driving style (which is observable via commonly available sensors). GamePlan assigns a higher priority to more aggressive or impatient drivers and a lower priority to more conservative or patient drivers; we theoretically prove that such an approach is game-theoretically optimal prevents collisions and deadlocks. We compare our approach with prior state-of-the-art auction techniques including economic auctions, time-based auctions (first-in first-out), and random bidding and show that each of these methods result in co
Distribution Statement: 1-Approved for public release; distribution is unlimited.
Acknowledged Federal Support: **Y**

CONFERENCE PAPERS:

RPPR Final Report
as of 25-Jul-2022

Publication Type: Conference Paper or Presentation **Publication Status:** 1-Published
Conference Name: 2019 International Conference on Robotics and Automation (ICRA)
Date Received: 24-Jul-2022 Conference Date: 20-May-2019 Date Published:
Conference Location: Montreal, QC, Canada
Paper Title: ADAPS: Autonomous Driving Via Principled Simulations
Authors: Weizi Li, David Wolinski, Ming C. Lin
Acknowledged Federal Support: **Y**

Partners

,

I certify that the information in the report is complete and accurate:
Signature: Ming C. Lin
Signature Date: 7/24/22 6:53PM

Please see previously uploaded papers during the project duration 2019-2022.