



# DEEP REINFORCEMENT LEARNING PATH PLANNER FOR SELF-DRIVING WHEELCHAIRS



STUDENTS: CHRIS CHEN, THOMAS LEE, SIRAAJ RIZVI, NIXON SARDJONO, MAAHNUM ZAHEER

## Motivation and Goal

- Self-driving wheelchairs must handle tight, crowded indoor spaces.
- Navigation should feel safe, smooth, and predictable for riders.
- Planner failures can cause unsafe stops, discomfort, or contact.
- We benchmark SAC, TEB, and Hybrid in repeatable Gazebo scenarios.
- Goal: select the strongest software setup before wheelchair trials.

8

scenarios

240


planner trials

10


runs each

## Scenario Design

### Reference cases



### Extension cases



Repeatable evaluation protocol    Shared simulation testbed    Demo: one Gazebo session

- C1-C4: reference navigation cases.
- C5-C8: added stress cases covering crowds, sudden obstacles, degraded sensing, and drift.
- Scenarios are shown in compact views for readability.
- All planners use the same scenario definitions.

## Shared ROS/Gazebo Testbed

### Fair comparison pipeline



### Recorded for each trial



Scenario effects reviewed consistently    Runtime uses one measurement pipeline    Figures from one benchmark summary

- Benchmark mode uses repeatable, isolated simulation trials.
- Demo mode can advance through C1-C8 in one Gazebo session.
- Timeouts, contact events, and recovery behavior remain visible.
- Scenario-specific effects are validated during playback and review.

## Benchmark Protocol

### What is recorded

- Success or timeout for every staged goal.
- Contact events and safety-relevant outcomes.
- Completion time, path length, and recovery behavior.
- Runtime resource usage from the same logging pipeline.
- Updated figures are generated from the current benchmark summary.

0%

Hybrid collisions

same

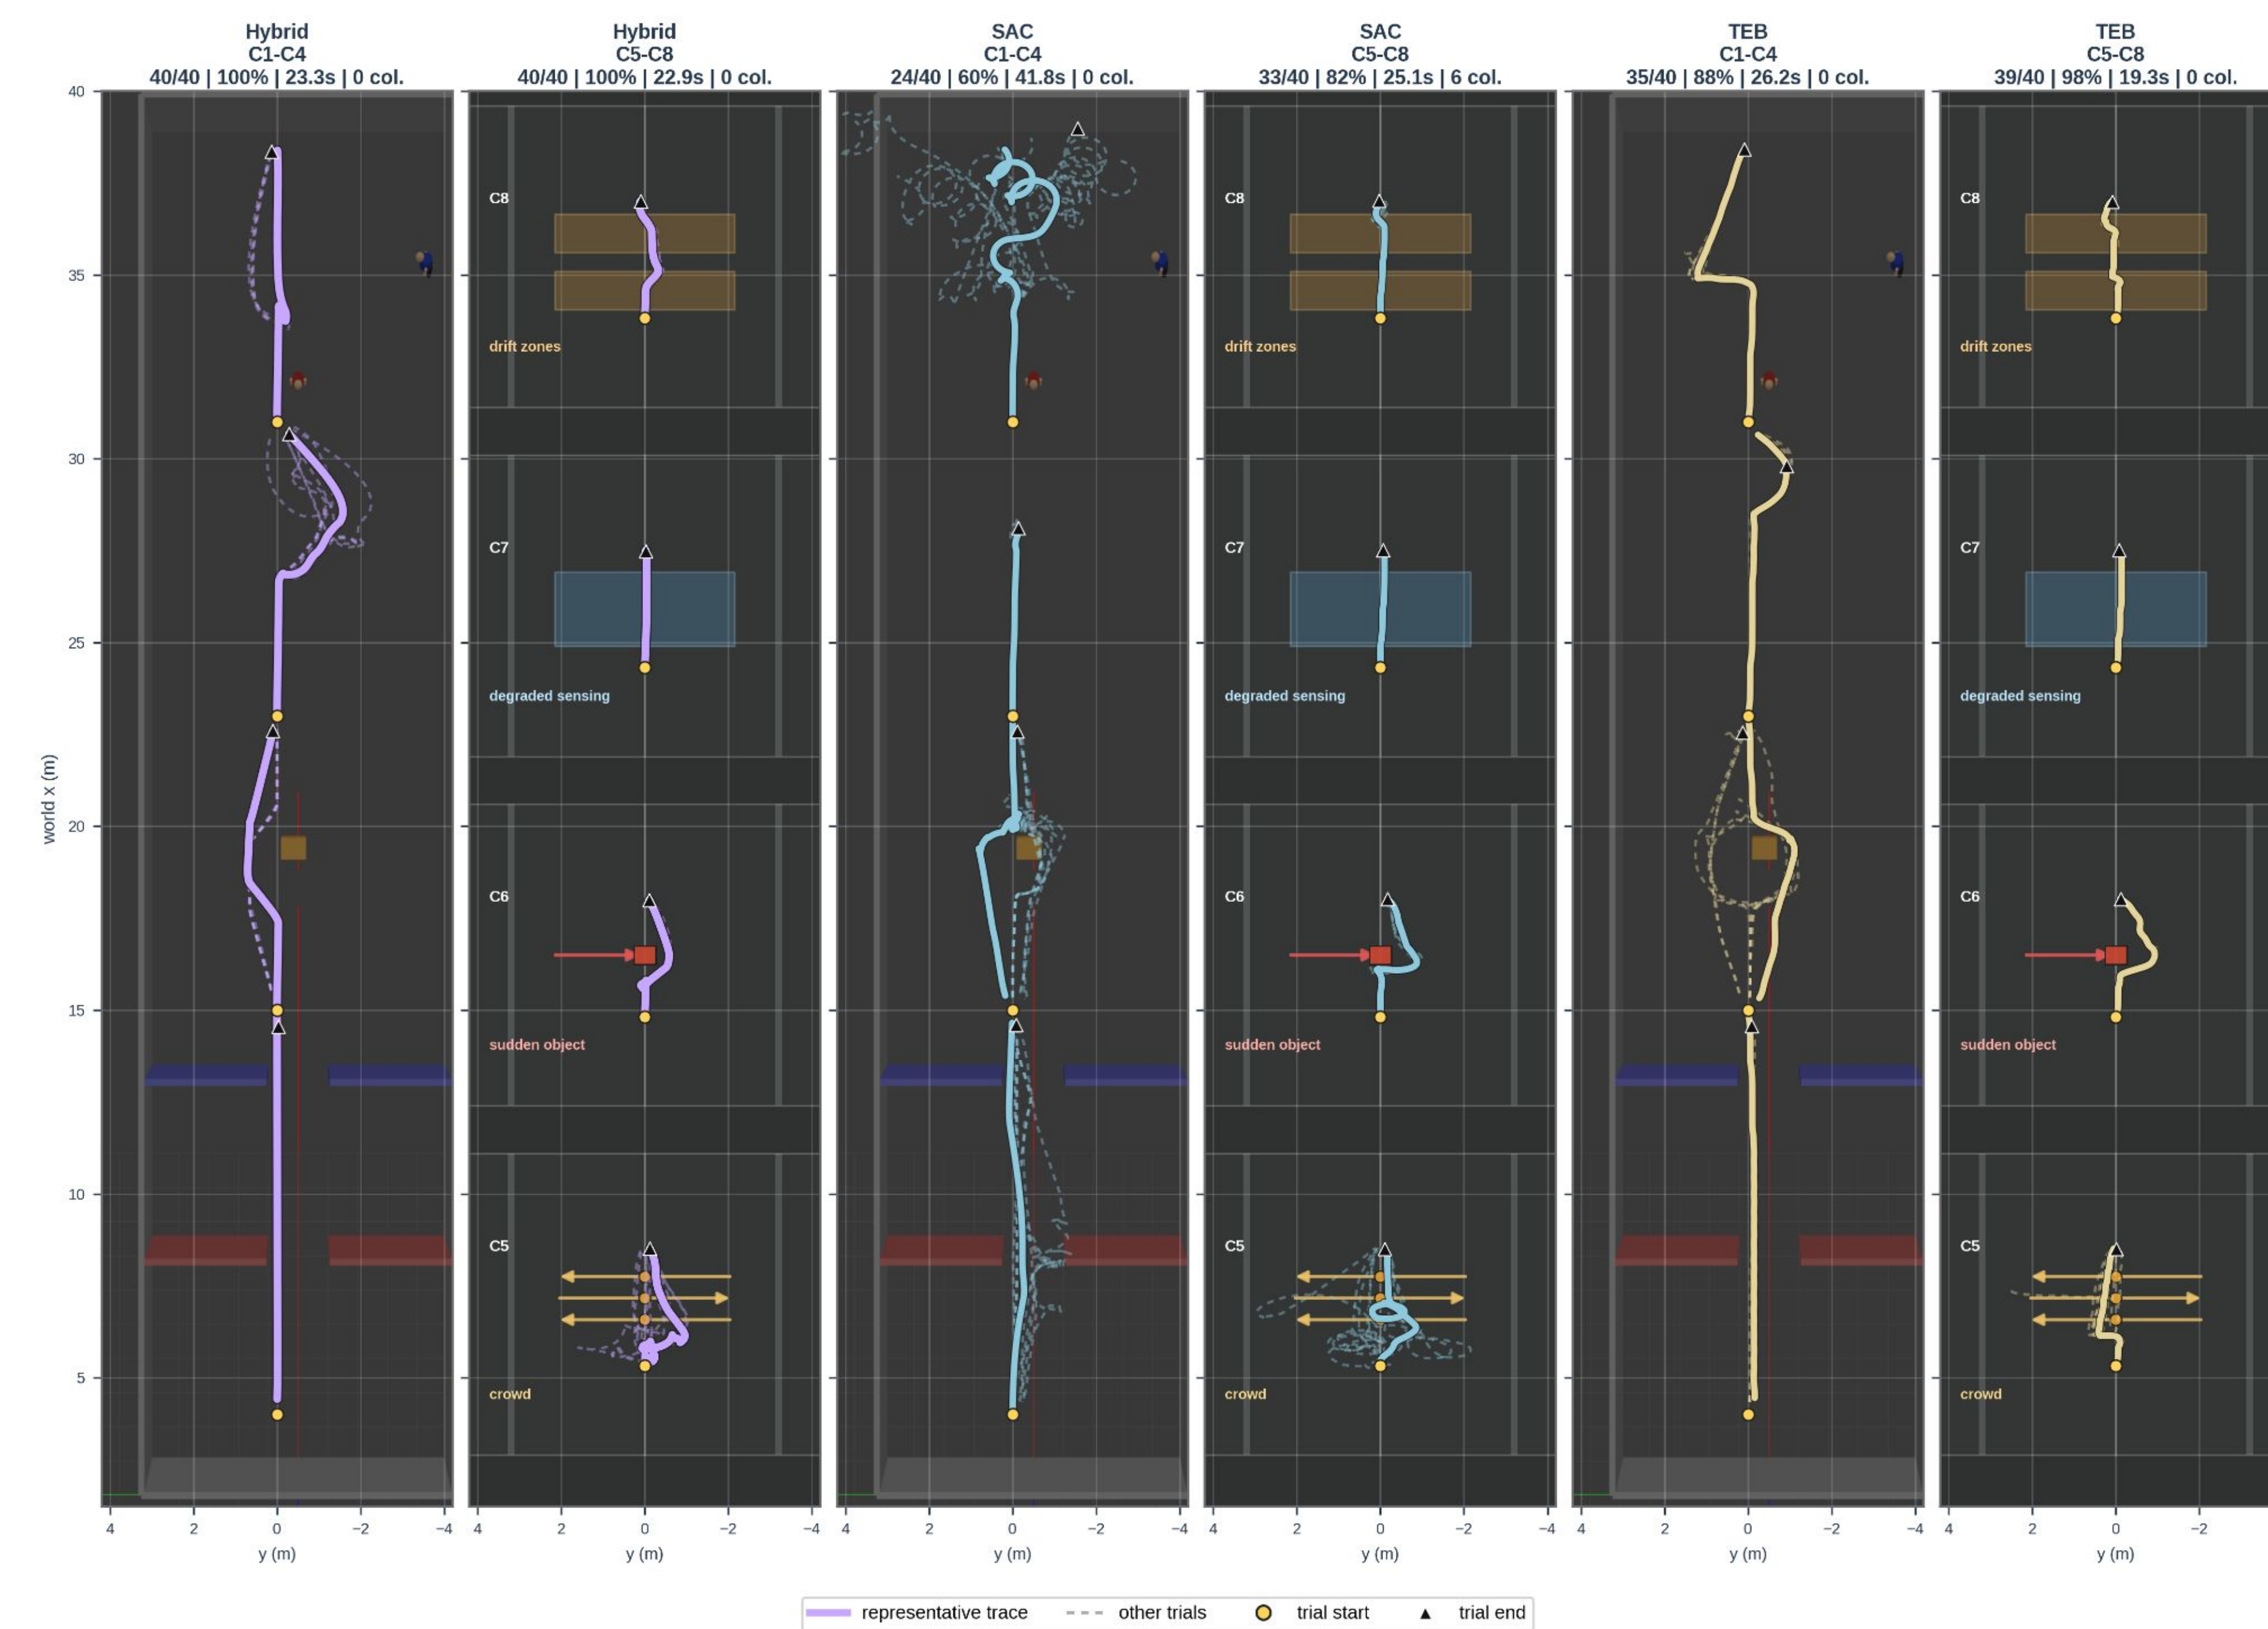
map + robot

same

repeatable logs

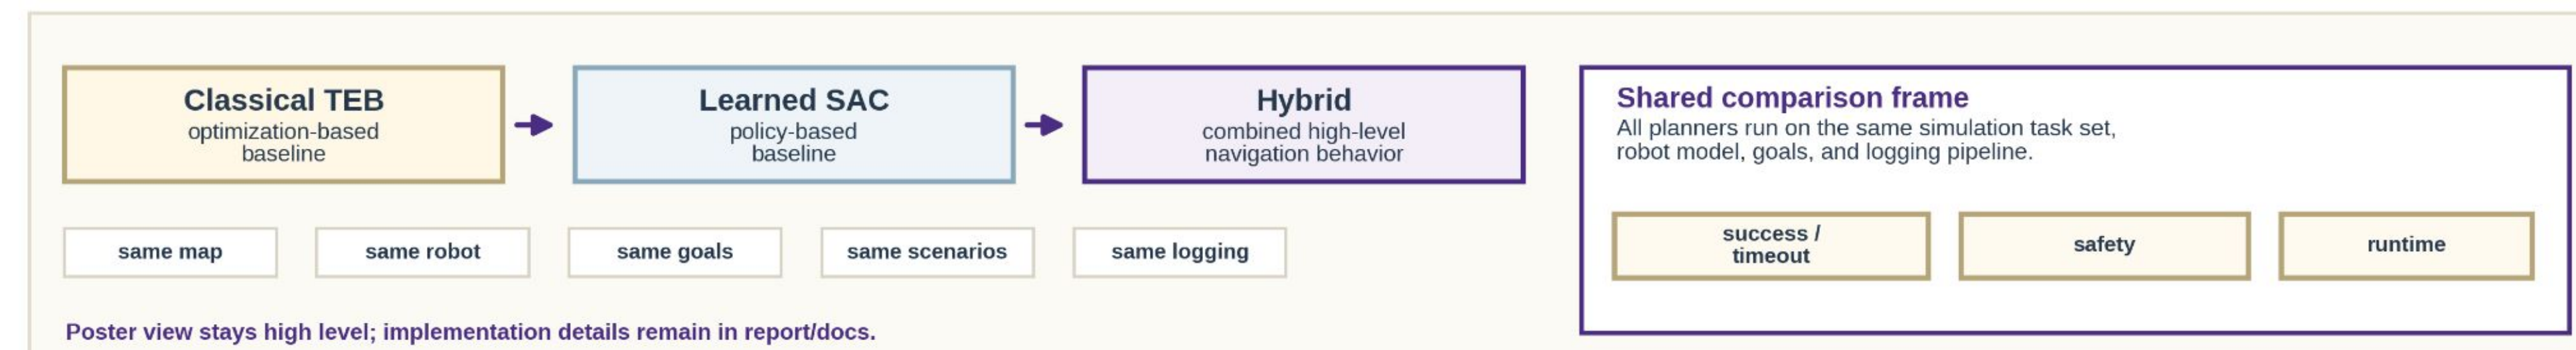
## Benchmark Trajectories

Six Gazebo top-down strips: C1-C4 use the full course view; C5-C8 use focused close-up summaries of the local interaction region.



Order: Hybrid C1-C4, Hybrid C5-C8, SAC C1-C4, SAC C5-C8, TEB C1-C4, TEB C5-C8. Solid lines show representative traces; faint dashed lines show other trials.

## Planning Approaches




## Conclusion

Hybrid completes all 80 simulated C1-C8 trials with zero logged collisions under the same testbed, scenarios, and logging pipeline, making it the strongest current software-side configuration before physical wheelchair validation.

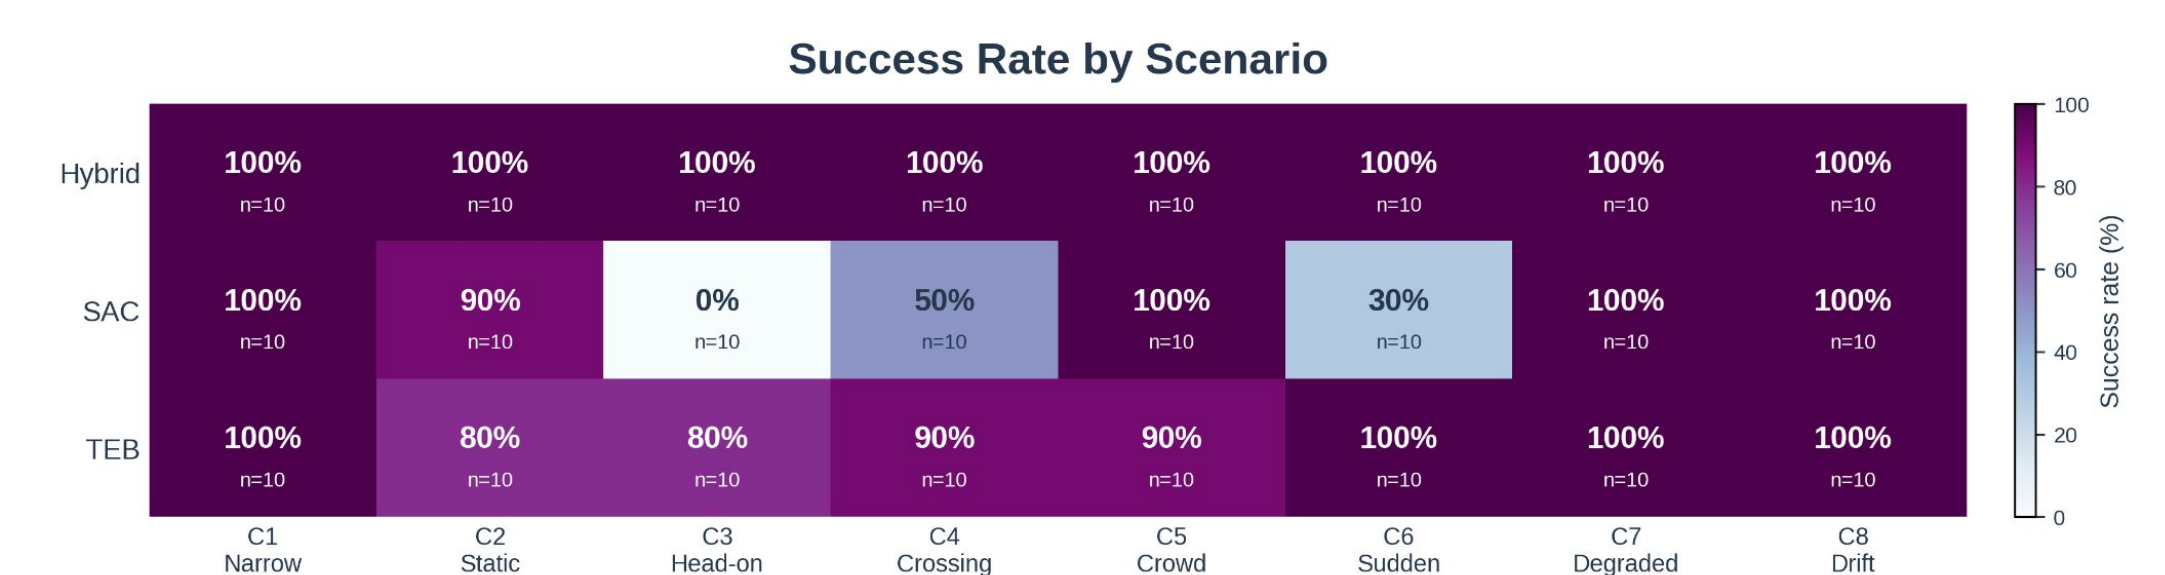
## Results at a Glance

### Full C1-C8 Benchmark Summary




Bars show successful runs out of 80; badges show missed runs, mean time, and logged contacts.

## Scenario-Level Results



- Hybrid reaches 100% in every simulated scenario.
- TEB remains competitive but drops in C2-C5.
- SAC drops most in C3, C4, and C6.

## Time and Resource Metrics



## Future Work, Acknowledgements, and References

### Future Work

- Run the full suite on the physical wheelchair.
- Add onboard safety limits and emergency-stop checks.
- Use real logs to refine recovery behavior.
- Expand rooms, pedestrians, and obstacle layouts.

### References

[1] K. Nakhleh, M. Raza, M. Tang, M. Andrews, R. Boney, I. Hadzic, J. Lee, A. Mohajeri, and K. Palyutina, "SACPlanner: Real-World Collision Avoidance with a Soft Actor Critic Local Planner and Polar State Representations," arXiv:2303.11801, 2023.  
 [2] V. D. Sharma, J. Lee, M. Andrews, and I. Hadzic, "Hybrid Classical/RL Local Planner for Ground Robot Navigation," arXiv:2410.03066, 2024.  
 [3] C. Roesmann, W. Feiten, T. Woesch, F. Hoffmann, and T. Bertram, "Efficient Trajectory Optimization Using a Sparse Model," European Conference on Mobile Robots, 2013.

### Acknowledgements

Cyberworks Robotics; Vivek Burhanpurkar, Nick Kourtzanidis, Payman Arabshahi, Sam Burden; UW ENGINE.