Perception-Driven Tube Scaling for State-Dependent Dynamic Tube-Based MPC

Filip Keerberg

fkeerberg@tudelft.nl

Responsible professor: Dr. Anahita Jamshidnejad

Supervisor: Sven Schoonebeek

TU Delft

1. Background and Motivation

- Use of autonomous robotics in post-disaster search and rescue (SaR) scenarios can significantly lower the risk to human life
- Robot has **trade-off** between **efficiency** and **safety**:
  - Must efficiently locate and aid victims
  - Must avoid hazards; damage to self, or others
- Hazards are spread throughout environment; safe/hazardous zones
- **External disturbances**: robot is moved differently from its intention
- Model Predictive Control (MPC) is a common control scheme for autonomous robotics
  - Exhaustively search action sequences and predicted outcomes
  - Select action sequence which results in best predicted outcome
  - *However*, not robust to (large) external disturbances
- Tube-based MPC (TMPC) enforces that any position/state is safe iff the worst-case disturbance is still safe
  - Can cause agent to be overly cautious
- Introduce **dynamic tube scaling**, change how cautious the agent is over time
- **Perception driven:** the robot will not always have perfect information of its environment. **Goal**: create tube scaling function on the agent's perception of the environment.

Testbed / Simulation

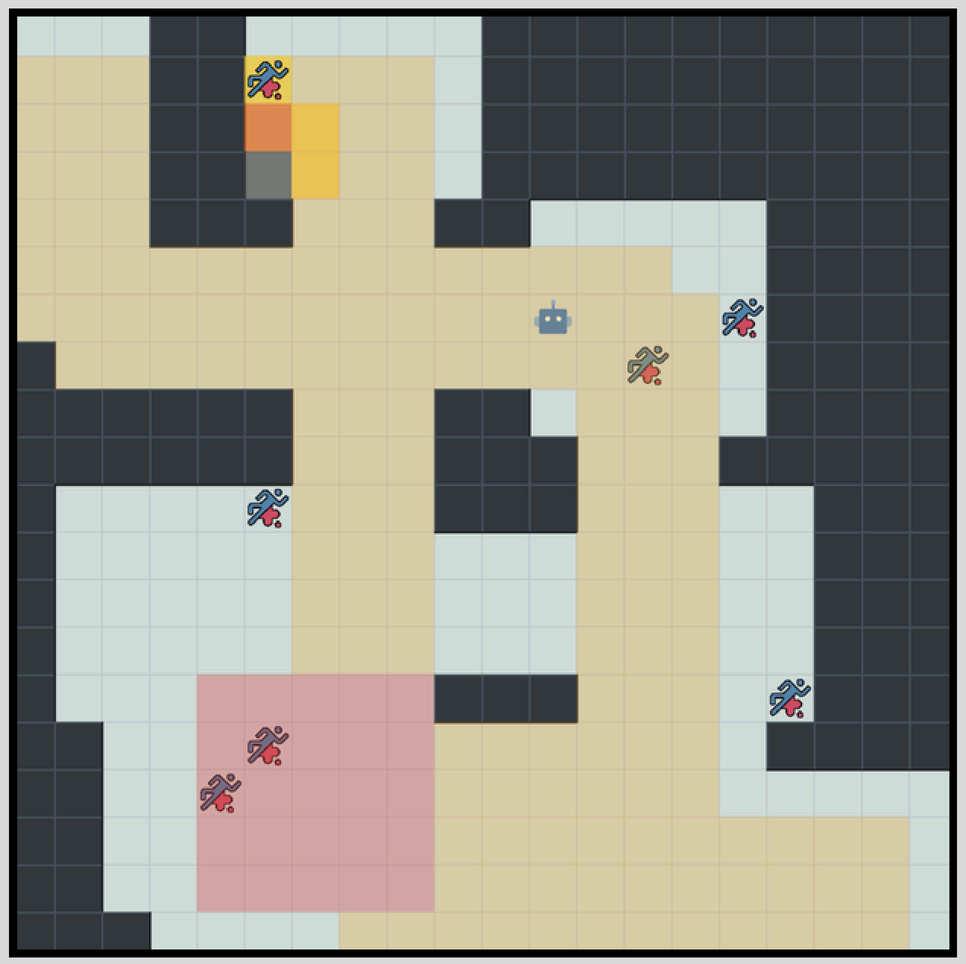

- Dark gray: Walls
- Bodies: Victims
- Blue robot: Agent
- Light gray: Burnt out fire
- Orange: Burning fire
- Yellow highlight: Moderate risk of disturbance
- Red highlight: Extreme risk of disturbance

2. Research Question

MAIN QUESTION

how can linking the tube size of a SDD-TMPC agent to said agent's perception of the environment, influence the agent's robustness, conservativeness, and performance in a SaR scenario?

SUB-QUESTIONS

- How should disturbance bounds be defined as a function of perceived confidence and disturbance risk, such that performance is improved over TMPC, and empirical robustness is maintained?
- To what extent are the performance gains and empirical robustness retained under increased perceptual uncertainty?
- How conservative is the resulting bound compared to a fixed worst-case bound used by conventional TMPC?

3. Methodology and Implementation

- **Develop search and rescue simulator**
  - Regions with disturbance probability
  - Traversable/untraversable terrain
  - Dynamic fires
  - Mission: locate all victims without colliding with anything
- **Agents**
  - Local scanning
  - Confidence values
  - Line-of-sight
  - BFS path-finding to closest unexplored tile
- **Agent Implementations**
  - MPC: Solely explore, don't consider disturbances
  - TMPC: Assume that a disturbance will occur at every step
    - Always keep at least 1 tile distance from walls, fires, and victims
    - Cannot traverse 1 wide safe corridors
  - SDD-TMPC: Scale cautiousness on perceived disturbance risk and confidence
    - If perceived probability of disturbance > 0, act like TMPC
    - If perceived probability of disturbance = 0, act like MPC
- **Perception/Scan Noise**
  - Disturbance risk in the range 0 to 1
  - Introduce uniformly sampled noise from -x to x into scanned risk value
  - Experiment with different magnitudes of noise: 0.1, 0.2, 0.4

5. Conclusion and Future Work

- Sub-question 1: scaling tube size on perceived environment is an effective method of retaining safety of TMPC, while increasing performance
- Sub-question 2: A boundary between noise magnitudes 0.2, and 0.4 is identified where safety and performance advantage are no longer present
- Sub-question 3: The mapping proposed in the research is significantly less conservative than worst-case TMPC

Future work: Continuous model (not a grid) to investigate continuously varying tube sizes. Introduce further uncertainty into scanning (i.e. a wall can be mistaken for a victim).

4. Results

- Performed 6 sets of 250 seeded simulations
  - 3 sets without noise: MPC, TMPC, SDD-TMPC
  - 3 sets with noise: 3x SDD-TMPC

- Disturbances may 'push' the agents into obstacles
- MPC has most collisions
- TMPC, SDD-TMPC have **0 wall/victim collisions**
- Fire collisions occur due to partial observability of a dynamic variable, not a fault of the controllers
  - Unfair to assess safety on fire collisions

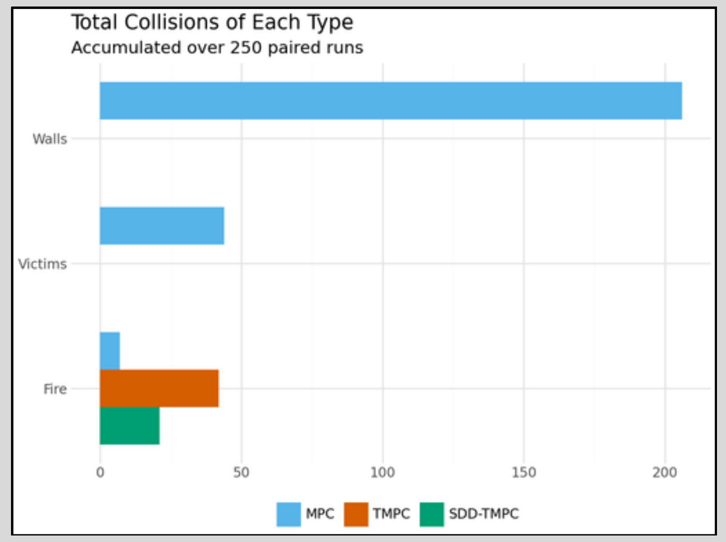

- Due to frequent collisions, MPC success-rate drops to 48.8%
- Due to over-conservativeness, TMPC success-rate: 39.8%
- Mix of safety and efficiency results in SDD-TMPC rate: 82.4%

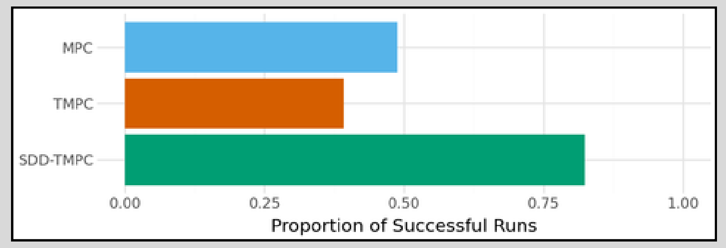

McNemar test between TMPC and SDD-TMPC, p = 1.71 x 10^{-31}, significant statistical improvement in success-rate

- As perceptual noise increases, success rate tends to decrease.
- Interestingly, rate from 0.0 → 0.1 increases. Small noise relaxes conservativeness while maintaining caution. With more significant noise, conservativeness becomes too lax
- At noise magnitude 0.4, odds-ratio (OR) of successes compared to TMPC is no longer statistically significant (OR = 1.35 [0.74, 2.47])

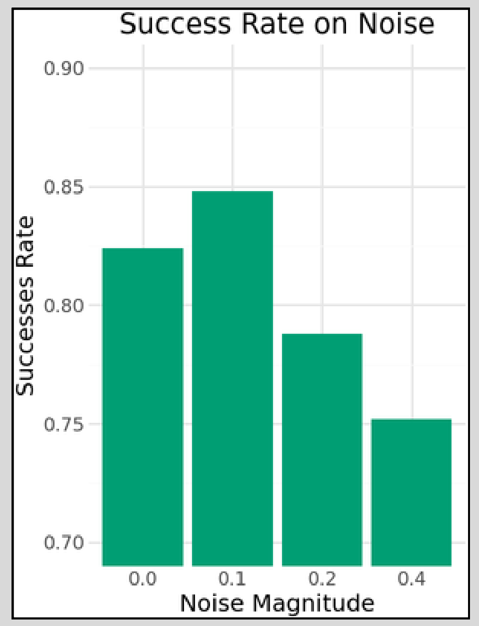

- Similar to success rate, safety also degrades at noise magnitude 0.4
- Some limit between noise magnitudes 0.2 and 0.4, where performance and safety of SDD-TMPC is lost

Noise	Wall	Victim
0.00	0	0
0.10	0	0
0.20	0	0
0.40	5	4

- Assign TMPC tube size of 1
- Assign MPC tube size of 0
- Record SDD-TMPC average tube size over all runs
- SDD-TMPC in 'TMPC mode' only 32.3% of the time (95% CI [29.9%, 34.7%])