

When No Safety Margin Is Admissible: Obstacle Inflation, Localisation Uncertainty and the Closing of the Feasible Window in Mobile Robot Navigation

Dr. Zubair Ahmed Khan

Assistant Professor, Department of Computer Science and Engineering, RIT College, Raipur, Chhattisgarh, India.

Email: zubairahmedkhanresearch@gmail.com

ABSTRACT

Mobile robot planners inflate obstacles by a safety margin to absorb localisation error. A larger margin reduces collisions and also closes gaps, so that in a cluttered workspace it can eliminate every feasible route. The margin is therefore constrained from both sides, and it is not guaranteed that any value satisfies both. We simulated a differential-drive robot on 140×140 occupancy grids at 5 cm resolution, planning with A* on inflated maps and executing under a stationary correlated localisation error. Across 40 random maps per configuration and 12 execution trials per plan, we measured planning success, collision rate, mission success and path length, sweeping the inflation margin, the workspace clutter and the localisation error. At 20 per cent clutter and a localisation error of 1.0 cells, mission success rose from 2.9 per cent with no margin to a maximum of 71.3 per cent at 2.5 cells, then fell to 49.2 per cent at 3.5 cells and to zero at 4.0 cells, where no plan existed on any map. Defining safety as fewer than 5 per cent collisions and feasibility as a plan on at least 80 per cent of maps, an admissible window existed at 8 and 14 per cent clutter but closed at 20 per cent and above, where safety required at least 3.0 cells and feasibility permitted at most 2.5. The optimal margin sat near three times the localisation standard deviation while that error was small, falling to 1.75 times it as the error grew and feasibility became binding. A margin sized only against localisation error is under-specified. Above a clutter threshold no single inflation radius was both safe and feasible, and the remedy is a change to the system rather than to the parameter.

Keywords: mobile robots; path planning; obstacle inflation; localisation uncertainty; autonomous navigation; safety margins.

1. INTRODUCTION

A mobile robot planning a route through a mapped environment does not plan against the map it was given. It plans against a map whose obstacles have been grown by a safety margin, because the robot has physical extent and because it does not know exactly where it is [1,2,12,13]. The inflation radius is one of the most consequential parameters in a navigation stack and one of the most casually set [18,19].

The reasoning behind it is usually one-sided. Localisation error is estimated, a multiple of its standard deviation is chosen, the robot radius is added, and the sum becomes the inflation radius. Larger is safer, and the parameter is treated as a safety dial with a monotone effect.

It is not monotone in the quantity that matters. Inflation closes gaps as well as widening clearances, and a doorway, a pallet aisle or a corridor between two machines admits a path only

while the inflated obstacles on either side do not meet. Past that point the planner does not produce a risky path; it produces no path, and the robot stops [10,11].

A stopped robot is safe and useless, and the two failure modes are therefore not comparable on a single axis. What matters operationally is mission success — the probability that the robot both finds a route and traverses it without contact — and that quantity is the product of two terms moving in opposite directions as the margin grows [20,25].

The consequence we examine here is that the admissible margin is bounded from both sides, and that those bounds can cross. Where they do, no setting of the parameter is acceptable and the problem cannot be solved by tuning. We set out to find where that happens and what determines it.

1.1 Contributions

We quantify mission success as a joint function of inflation margin, workspace clutter and localisation error; we identify the clutter threshold above which the safe and the feasible margin ranges cease to overlap; we express the optimal margin in units of localisation error and show where that scaling rule breaks; and we separate the planning and execution contributions to failure so that a designer can tell which constraint is binding.

2. METHOD

2.1 Workspace and Planner

Workspaces are 140×140 occupancy grids at 5 cm per cell, populated with randomly placed rectangular obstacles until a target occupied fraction is reached, with the start and goal regions kept clear and the boundary occupied [12,13]. Clutter is defined as that occupied fraction.

Obstacles are inflated by binary dilation with a disc of the specified radius, and a path is sought with 8-connected A* using Euclidean distance as the heuristic [9]. If no path exists the trial is recorded as a planning failure and no execution is attempted. Figure 1 shows one workspace planned at three margins.

2.2 Localisation Error and Execution

The robot follows the planned cell sequence while its true pose differs from its believed pose. We model that difference as a stationary first-order autoregressive process in each axis, with correlation $\rho = 0.90$ between successive steps and a stationary standard deviation σ [3]:




Fig.1: One workspace planned at three inflation margins. At the largest margin no route survives.

$$e_k = \rho e_{k-1} + w_k, w_k \sim N(0, \sigma^2(1 - \rho^2)) \quad (1)$$

The autoregressive form matters. A random walk would grow without bound and would describe dead reckoning rather than localisation; a localising robot corrects against a map, so its error is correlated in time but stationary in magnitude [14-17]. Using a random walk here inflates collision rates by an artefact of path length rather than of geometry.

A trial is recorded as a collision if the true pose at any step of the path falls in a genuinely occupied cell of the uninflated map. Mission success is the product of the planning success rate and the collision-free execution rate.

2.3 Experiments

Four sweeps were run: inflation margin from 0.0 to 4.0 cells at the base clutter and localisation error; a full grid of margin against clutter; localisation error from 0 to 2.0 cells; and path length against margin. Each configuration used 40 independent maps with 12 execution trials per successful plan.

Parameters are given in Table 1. They describe a warehouse-scale differential-drive platform and are illustrative rather than drawn from a specific machine. The simulation was implemented in Python using NumPy, SciPy and Matplotlib [28-30].

Table 1: Simulation parameters.

Parameter	Value	Note
Workspace	140 × 140 cells	7.0 m square at 5 cm resolution
Obstacles	Random rectangles, 4–14 cells	Start and goal regions kept clear
Base clutter	20%	Occupied fraction of the workspace
Localisation error σ	1.0 cells	5 cm, stationary
Error correlation ρ	0.90	Correlated between steps, bounded in magnitude
Planner	8-connected A*	Euclidean heuristic on the inflated grid
Maps per configuration	40	Independent random workspaces
Trials per plan	12	Independent execution noise realisations

3. RESULTS

3.1 An Interior Optimum, Not a Safety Dial

Figure 2 and Table 2 sweep the margin at 20 per cent clutter. The two component rates behave as expected and in opposition: planning success falls monotonically from 100.0 per cent with no inflation to 0.0 per cent at 4.0 cells, while collision-free execution rises from 2.9 to 98.3 per cent.

Their product has a clear interior maximum. Mission success rose from 2.9 per cent at zero margin to 71.3 per cent at 2.5 cells, then fell back to 49.2 per cent at 3.5 cells and to zero at 4.0 cells.

The final column is the one that resists intuition. At 4.0 cells the collision rate is undefined because no robot ever moved: the inflated obstacles merged on every map and the planner returned failure every time. A system tuned by minimising collisions alone converges on exactly this configuration, which has a perfect safety record and completes no missions.




Fig.2: Planning success, collision-free execution and their product against inflation margin.

Table 2: Margin sweep at base clutter and localisation error.

Margin (cells)	Plan exists (%)	Collision rate (%)	Mission success (%)	Mean path length (m)
0.0	100.0	97.1	2.9	7.65
0.5	100.0	97.1	2.9	7.65
1.0	95.0	77.4	21.5	8.07
1.5	92.5	59.5	37.5	8.09
2.0	87.5	27.6	63.3	8.31
2.5	85.0	16.2	71.3	8.73
3.0	57.5	2.2	56.3	9.34
3.5	50.0	1.7	49.2	9.04
4.0	0.0	—	0.0	—

Margin in cells of 5 cm. Localisation error $\sigma = 1.0$ cells.

3.2 The Window Closes

Figure 4 and Table 3 apply two operational criteria: safety, taken as a collision rate below 5 per cent, and feasibility, taken as a plan existing on at least 80 per cent of maps. The margin ranges satisfying each are plotted against clutter, and the full grid of outcomes appears in Figure 3 and Table 4.

At 8 and 14 per cent clutter the two ranges overlap: safety required at least 3.0 cells and feasibility permitted up to 3.5, leaving a narrow admissible window of 0.5 cells.

At 20 per cent clutter and above the ranges no longer intersect. Safety still required at least 3.0 cells; feasibility permitted at most 2.5. No value of the parameter satisfies both criteria, and the best achievable mission success at that clutter was 75.8 per cent.

This is the paper's central result and it is a statement about the system rather than about the tuning. When the window is closed, an engineer adjusting the inflation radius is choosing which requirement to violate, not searching for a setting that meets both. Reporting the chosen value without reporting that the window was closed conceals a decision that should have been escalated [20,21].

The diagnostic is cheap. Both criteria are measurable from data a deployment already generates — planning failure rate and contact events — and plotting them against the margin on a representative set of maps takes one afternoon.




Fig.4: The safety requirement and the feasibility ceiling against clutter.

Table 3: The admissible window by clutter.

Clutter (%)	Minimum margin for safety (cells)	Maximum margin for feasibility (cells)	Window (cells)	Best mission success (%)
8	3.0	3.5	0.5	96.7
14	3.0	3.5	0.5	91.9
20	3.0	2.5	none	75.8
26	3.0	1.5	none	60.2
32	3.0	1.5	none	19.0

Safety: collision rate below 5%. Feasibility: a plan exists on at least 80% of maps.

Table 4: Mission success (%) across margin and clutter.

Clutter	0.0	0.5	1.0	1.5	2.0	2.5	3.0	3.5	4.0
8%	28.7	28.7	52.5	76.0	90.4	94.8	96.3	96.7	0.0
14%	10.0	10.0	39.2	52.3	69.8	84.8	91.9	87.1	0.0
20%	2.5	2.5	27.3	36.3	64.4	70.0	73.3	75.8	0.0
26%	0.8	0.8	11.5	25.2	42.7	60.2	17.3	15.0	0.0
32%	0.2	0.2	10.4	14.2	19.0	11.5	10.0	10.0	0.0

Columns give the inflation margin in cells.

3.3 Scaling with Localisation Quality

Figure 5 and Table 5 vary the localisation error. With perfect localisation the optimal margin is zero and mission success is 100.0 per cent, as it must be: inflation exists only to absorb uncertainty.

At $\sigma = 0.5$ cells the optimum was 1.5 cells, exactly 3.00 times σ , and at $\sigma = 1.0$ cells it was 3.0 cells, again 3.00 times σ . The familiar three-sigma heuristic is therefore correct in this regime.




Fig.3: Mission success across the margin and clutter grid.

It stops being correct as the error grows. At $\sigma = 1.5$ cells the optimum was 2.00σ and at $\sigma = 2.0$ cells it was 1.75σ . The margin the geometry would demand no longer fits in the workspace, so the optimum is set by feasibility rather than by uncertainty and the ratio falls.

The practical reading is that a margin rule expressed purely in units of σ is valid only while the resulting margin is small relative to the free space. Beyond that, the workspace is the binding constraint and the rule silently stops applying — and the sign that it has stopped is a rising planning failure rate, not anything visible in the localisation statistics.

The value of localisation quality is also visible in the last column of Table 5. Improving σ from 1.0 to 0.5 cells raised the best achievable mission success from 70.6 to 96.5 per cent. Where the window is closed, this is the lever that reopens it [14-17].

Table 5: Effect of localisation error.

σ (cells)	σ (cm)	Optimal margin (cells)	Optimal margin (σ units)	Best mission success (%)
0.0	0	0.0	—	100.0
0.5	3	1.5	3.00	96.5
1.0	5	3.0	3.00	70.6
1.5	8	3.0	2.00	43.5
2.0	10	3.5	1.75	19.2

3.4 The Distance Cost

Table 6 and Figure 6 record path length. The shortest route, planned with no inflation, averaged 7.65 m. At the optimal margin it averaged 8.73 m, 14.1 per cent longer, and at the largest feasible margin 9.04 m.




Fig.5: Optimal margin against localisation error, in absolute terms and in units of σ .

The increase is modest and it compounds. A fleet operating continuously pays it on every mission in time, energy and throughput, so a margin chosen larger than necessary is a permanent tax rather than a one-off cost. This is a further reason to identify the optimum rather than to err upward on the assumption that extra margin is free.

Table 6: Path length against margin.

Margin (cells)	Mean path length (m)	Increase over shortest (%)	Mission success (%)
0.0	7.65	0.0	2.9
0.5	7.65	0.0	2.9
1.0	8.07	5.5	21.5
1.5	8.09	5.8	37.5
2.0	8.31	8.6	63.3
2.5	8.73	14.1	71.3
3.0	9.34	22.1	56.3
3.5	9.04	18.2	49.2




Fig.6: Mean executed path length against inflation margin.

4. DISCUSSION

4.1 Principal Findings

Four findings follow. Mission success has an interior optimum in the inflation margin, at 2.5 cells in the base configuration, with zero mission success at 4.0 cells where no plan exists at all. Above 20 per cent clutter the safe and feasible margin ranges ceased to overlap, so no parameter value satisfied both criteria. The three-sigma heuristic held while localisation was good and degraded to 1.75σ once feasibility became binding. And the optimal margin cost 14 per cent in path length.

The second is the one we would press. Parameter tuning presupposes that an acceptable value exists, and in cluttered workspaces that presupposition fails. The failure is invisible if only one criterion is monitored, because each criterion alone can always be satisfied by moving the parameter in its own direction.

4.2 Relation to Existing Work

Obstacle inflation and costmap construction are standard components of navigation stacks and are documented as such [1,2,18,19]. Planning under localisation and motion uncertainty has a substantial literature, including belief-space planning, chance-constrained formulations and probabilistic collision checking, which reason about uncertainty explicitly rather than through a fixed margin [3-6,26,27]. Reactive and predictive local avoidance layers exist precisely to decouple long-horizon planning from short-horizon safety [7,8,23,24].

Our contribution is narrower and, we think, complementary: a quantification of when the conventional fixed-margin approach has no admissible setting, expressed in quantities a deployment already measures. Where the window is open, a fixed margin remains a reasonable and cheap engineering choice, and our results give the value. Where it is closed, the result is an argument for adopting one of the methods above rather than for continuing to tune [22,25].

4.3 Design Guidance

Figure 7 and Table 7 set out four recommendations. Two deserve comment.

Table 7: Which constraint binds, and what to do about it.

Symptom	Binding constraint	Ineffective response	Effective response
Contacts with obstacles	Margin too small for σ	Retuning the planner cost function	Increase margin, or improve localisation
Planner returns no path	Margin too large for the gaps	Increasing margin further	Reduce margin, or widen the route
Both, at different times	Window closed	Any single margin value	Improve localisation or add reactive avoidance
Long routes, few failures	Margin above the optimum	Accepting it as safe	Reduce toward the optimum and measure

Measure both failure modes, always together. Planning failures and contacts are recorded by different subsystems and are often reviewed by different people, which is how a closed window goes unnoticed. A single plot of both against the margin, on a representative map set, is the whole diagnostic.

Treat localisation quality as the lever of first resort. Improving σ lowers the safety requirement without touching the feasibility ceiling, and is the only change in our study that reopens a closed window without altering the workspace or the robot. It is also usually cheaper than re-laying an aisle.




Fig.7: Proposed practice for setting and auditing the inflation margin.

5. LIMITATIONS

The collision criterion is a union over every step of the path, so a longer route has more opportunities to make contact and the absolute collision rates depend on path length as well as on geometry. This is physically reasonable but it means our rates are specific to the workspace scale simulated and should not be read as expected rates for a real platform.

There is no reactive avoidance layer. A real robot detects an imminent obstacle and stops or deviates, which converts many of our collisions into delays. Including such a layer would lower the safety requirement and widen the admissible window, so our window estimates are conservative in that respect and the qualitative closing result would occur at higher clutter than we report.

Localisation error is modelled as isotropic, stationary and independent of position, whereas real error grows in feature-poor regions, is often anisotropic along a corridor, and is correlated with exactly the narrow passages where margin matters most. That correlation would make the problem worse than modelled.

We used a grid planner with a fixed resolution and a holonomic 8-connected motion model, so kinematic constraints, path smoothing and controller tracking error are all absent. A non-holonomic platform following a smoothed path has a different effective footprint, and the margin that suits it would differ.

Finally, obstacles are random rectangles with no structure. Real workspaces have corridors, doorways and aisles whose widths are designed, and in such environments the feasibility ceiling is set by a few critical passages rather than by an average clutter figure. The clutter parameter used here is a coarse proxy for that geometry and we would not expect the threshold value to transfer.

6. CONCLUSION

Inflating obstacles to absorb localisation error trades collision risk against the existence of a route. In the workspaces simulated, mission success peaked at 2.5 cells — 2.5 times the localisation standard deviation — and fell to zero at 4.0 cells, where the planner found no path on any map.

Above 20 per cent clutter the two requirements stopped overlapping: safety demanded at least 3.0 cells and feasibility permitted at most 2.5. In that regime the parameter has no acceptable value, and an engineer adjusting it is selecting which requirement to break.

The three-sigma rule was accurate while localisation was good and misleading once the workspace became the binding constraint, with no signal of the transition in the localisation statistics themselves. Monitoring planning failures alongside contacts is what makes the transition visible, costs nothing, and distinguishes a parameter that needs tuning from a system that needs changing.

DECLARATIONS

Ethics approval and consent to participate. Not applicable. This is a simulation study using synthetic data generated by the author. It did not involve human participants, animals or personal data, and institutional review board approval and informed consent were therefore not required.

Consent for publication. Not applicable.

Data availability statement. No external data were used. All results were generated by the simulation script described in Section 2, developed by the author in Python using NumPy, SciPy and Matplotlib [28-30], with all random seeds fixed. Every parameter is stated in full in Table 1, and every figure and every value in Tables 2 to 6 can be regenerated exactly by executing that script. The script is available from the corresponding author upon reasonable request and will be deposited in a public repository with a permanent identifier on acceptance.

Conflict of interest. The author declares no competing interests, financial or non-financial, in relation to the subject matter of this manuscript. The author has no financial or advisory relationship with any manufacturer of mobile robots, navigation software or localisation hardware.

Funding. This research received no specific grant from any funding agency in the public, commercial or not-for-profit sectors.

Author's contribution. Dr. Zubair Ahmed Khan is the sole author of this work and was responsible for the conception and design of the study, the design and implementation of the simulation experiments, the analysis and verification of all reported values, the preparation of the figures and tables, the formulation of the design guidance, and the drafting, revision and final approval of the manuscript. The author is accountable for all aspects of the work.

Acknowledgements. None declared.

REFERENCES

1. Lu DV, Hershberger D, Smart WD. Layered costmaps for context-sensitive navigation. Proc IEEE/RSJ Int Conf Intell Robots Syst. 2014;709–15.
2. Marder-Eppstein E, Berger E, Foote T, Gerkey B, Konolige K. The office marathon: robust navigation in an indoor office environment. Proc IEEE Int Conf Robot Autom. 2010;300–07.
3. Thrun S, Burgard W, Fox D. Probabilistic robotics. Cambridge: MIT Press; 2005.
4. Van Den Berg J, Abbeel P, Goldberg K. LQG-MP: optimized path planning for robots with motion uncertainty and imperfect state information. Int J Robot Res. 2011;30(7):895–913.
5. Platt R, Tedrake R, Kaelbling L, Lozano-Perez T. Belief space planning assuming maximum likelihood observations. Proc Robot Sci Syst. 2010.
6. Blackmore L, Ono M, Williams BC. Chance-constrained optimal path planning with obstacles. IEEE Trans Robot. 2011;27(6):1080–94.

7. Fox D, Burgard W, Thrun S. The dynamic window approach to collision avoidance. *IEEE Robot Autom Mag.* 1997;4(1):23–33.
8. Rosmann C, Hoffmann F, Bertram T. Integrated online trajectory planning and optimization in distinctive topologies. *Robot Auton Syst.* 2017;88:142–53.
9. Hart PE, Nilsson NJ, Raphael B. A formal basis for the heuristic determination of minimum cost paths. *IEEE Trans Syst Sci Cybern.* 1968;4(2):100–07.
10. LaValle SM. *Planning algorithms.* Cambridge: Cambridge University Press; 2006.
11. Karaman S, Frazzoli E. Sampling-based algorithms for optimal motion planning. *Int J Robot Res.* 2011;30(7):846–94.
12. Elfes A. Using occupancy grids for mobile robot perception and navigation. *Computer.* 1989;22(6):46–57.
13. Moravec H, Elfes A. High resolution maps from wide angle sonar. *Proc IEEE Int Conf Robot Autom.* 1985;116–21.
14. Grisetti G, Stachniss C, Burgard W. Improved techniques for grid mapping with Rao-Blackwellized particle filters. *IEEE Trans Robot.* 2007;23(1):34–46.
15. Kohlbrecher S, von Stryk O, Meyer J, Klingauf U. A flexible and scalable SLAM system with full 3D motion estimation. *Proc IEEE Int Symp Safety Secur Rescue Robot.* 2011;155–60.
16. Hess W, Kohler D, Rapp H, Andor D. Real-time loop closure in 2D LIDAR SLAM. *Proc IEEE Int Conf Robot Autom.* 2016;1271–78.
17. Cadena C, Carlone L, Carrillo H, Latif Y, Scaramuzza D, Neira J, et al. Past, present, and future of simultaneous localization and mapping: toward the robust-perception age. *IEEE Trans Robot.* 2016;32(6):1309–32.
18. Macenski S, Martin F, White R, Clavero JG. The Marathon 2: a navigation system. *Proc IEEE/RSJ Int Conf Intell Robots Syst.* 2020;2718–25.
19. Quigley M, Conley K, Gerkey B, Faust J, Foote T, Leibs J, et al. ROS: an open-source Robot Operating System. *ICRA Workshop Open Source Softw.* 2009;3(3.2):5.
20. Guiochet J, Machin M, Waeselynck H. Safety-critical advanced robots: a survey. *Robot Auton Syst.* 2017;94:43–52.
21. International Organization for Standardization. ISO 3691-4:2023 Industrial trucks — Safety requirements and verification — Part 4: Driverless industrial trucks and their systems. Geneva: ISO; 2023.
22. Lasota PA, Fong T, Shah JA. A survey of methods for safe human-robot interaction. *Found Trends Robot.* 2017;5(4):261–349.
23. Fraichard T, Asama H. Inevitable collision states: a step towards safer robots? *Adv Robot.* 2004;18(10):1001–24.
24. Althoff M, Dolan JM. Online verification of automated road vehicles using reachability analysis. *IEEE Trans Robot.* 2014;30(4):903–18.
25. Majumdar A, Pavone M. How should a robot assess risk? Towards an axiomatic theory of risk in robotics. *Robot Res.* 2020;10:75–84.
26. Bry A, Roy N. Rapidly-exploring random belief trees for motion planning under uncertainty. *Proc IEEE Int Conf Robot Autom.* 2011;723–30.
27. Pereira AA, Binney J, Hollinger GA, Sukhatme GS. Risk-aware path planning for autonomous underwater vehicles using predictive ocean models. *J Field Robot.* 2013;30(5):741–62.
28. Harris CR, Millman KJ, van der Walt SJ, Gommers R, Virtanen P, Cournapeau D, et al. Array programming with NumPy. *Nature.* 2020;585(7825):357–62.
29. Virtanen P, Gommers R, Oliphant TE, Haberland M, Reddy T, Cournapeau D, et al. SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nat Methods.* 2020;17(3):261–72.

30. Hunter JD. Matplotlib: a 2D graphics environment. Comput Sci Eng. 2007;9(3):90–95.