ICRA 2026 · SPACE ROBOTICS WORKSHOP · SPOTLIGHT

From SLAM drift to goal drift.

Measuring coordination sensitivity on lunar analogue terrain.

Andreja Malus · Tena Žužek · Jure Dvoršak · Rok Vrabič
University of Ljubljana · Faculty of Mechanical Engineering

UNIVERSITY OF LJUBLJANA Faculty of Mechanical Engineering

ROVERS at LUNA · ESA/DLR Cologne

THREE FINDINGS

Coordination under SLAM drift is structurally tolerant—within limits.

1. Tolerance is real, but bounded. Distance ranks can change without changing the chosen frontier, because the rule only commits to the closest uncontested cell. Once drift reaches the same physical scale as the grid, that silent margin starts to disappear.

2. Cell resolution sets the SLAM requirement. Localisation accuracy must stay within the occupancy-cell scale throughout the run. This is a coordination requirement, not a number that can be read from the SLAM specification alone.

3. Pose accuracy and map consistency are independent. Accurate poses still do not rescue a corrupted frontier set: if each occupancy grid drifts in its own frame, robots deconflict over different physical cells.

PAPER · FIGURES · DATA
andreja.malus@fs.uni-lj.si
openreview.net/forum?id=Jkx3OydlWX

PROBLEM

Multi-robot coordination breaks silently when SLAM drifts. **We quantify how much drift it can absorb.**

QUESTION

How much SLAM drift can a decentralised frontier rule absorb before robots claim **different goals**? And, turning the question around, what accuracy must a SLAM backend actually deliver to keep coordination intact?

AUDIT

For every coordination tick we replay the deconfliction rule twice: once with **live poses** as the robot saw them, once with **RTAB-Map graph-optimised poses** as the post-hoc reference. Both are SE(2)-aligned and interpolated to the tick timestamp.

METRICS

Rank changed — when *any* frontier in the ordering shifted (sensitivity of the distance computation). **Goal changed** — when the *claimed* cell itself shifted (impact on the actual assignment).

A · ALGORITHM

Distance-rank deconfliction

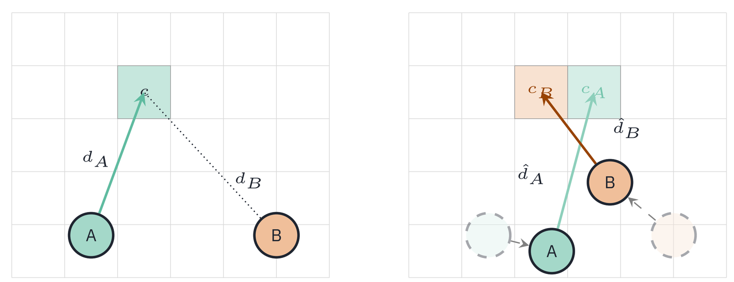

Left: shared inputs make every robot compute the same assignment.
Right: SLAM drift shifts both peer poses and the grid, so identical algorithms now compute incompatible answers.

Each robot publishes its **live pose**, its currently selected **frontier cell**, and occupancy-grid updates; it receives the same state fr om peers and merges peer cells that are still unknown in its own grid. Given that shared state, every robot runs the same deterministic assignment locally:

1. Build the candidate **frontier set** from the current shared occupancy grid.
2. For every robot in the shared state, compute its **Euclidean distance** to every candidate cell.
3. Assign the **closest** available robot–cell pair, remove that cell and its 8-cell neighbourhood, and repeat until every robot has one goal.

B · PIPELINE PRUNING

Two stages already prune the frontier set

Cells are filtered twice before the distance-rank step the audit measures. **Cluster selection** keeps only the best-scoring connected frontier region (**44 – 84 %** of cells removed across all runs). **Peer deconfliction** then drops cells already claimed by a closer peer (**6 – 44 %** outdoors). The audited assignment runs on what survives.

At LUNA the peer stage was offline (see H), so the FR1–FR3 audit captures pose-error sensitivity only.

C · PLATFORM

Leo Rover fleet

Five **Leo Rovers** (~8.5 kg each, differential drive) running ROS2 on **NVIDIA Jetson Orin Nano**. Primary sensor: **Intel RealSense D456** RGB-D + IMU. **RTAB-Map** supplies visual-inertial SLAM and the 2D occupancy grid; **vox_nav** handles path planning. Peers exchange pose, selected frontier cell, and grid updates over WiFi at ~**1 Hz**.

D · TEST ENVIRONMENTS

Outdoor field vs. LUNA facility

OUTDOOR · OR1, OR2

Open ground with rich visual features and varied texture — a favourable environment for visual SLAM. This is the **baseline** regime: where coordination is *supposed* to work, and where we expect tolerance.

LUNA · FR1, FR2, FR3

A lunar surface analogue at ESA/DLR's LUNA facility in Cologne, with low-texture sand zones and simulated lunar illumination. Featureless terrain **stresses** visual SLAM and starves loop closures — exposing what happens to coordination once drift accumulates. In these runs, the audit measures pose-error sensitivity only; missing peer deconfliction would compound map-frame errors.

E · CASE STUDY FR2

What drift looks like in a single run

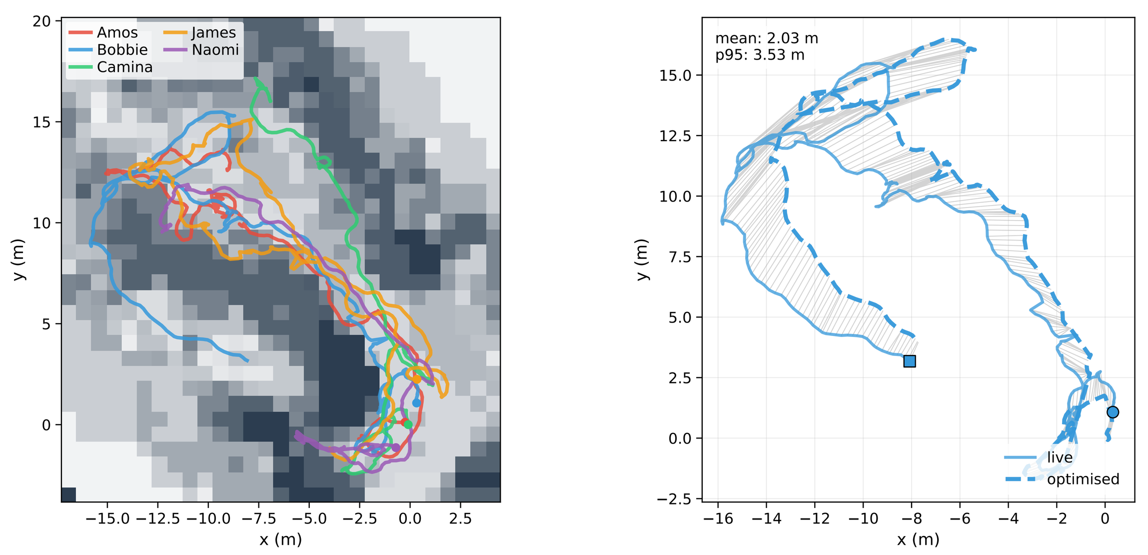

FR2 at LUNA. Left: live trajectories of all five robots on the shared occupancy grid. Right: Bobbie's live vs. graph-optimised trajectory, with grey per-tick displacement lines showing where the two diverged.

Bobbie's mean error reached **2.03 m** (p95 **3.53 m**) — about **2.5× the cell size**. The graph-optimised reference itself is bounded at ~**0.26 m** (p95, over 140 loop closures), so the audit baseline is trustworthy.

F · ERROR VS. ASSIGNMENT

Coordination error scales with cell-relative drift

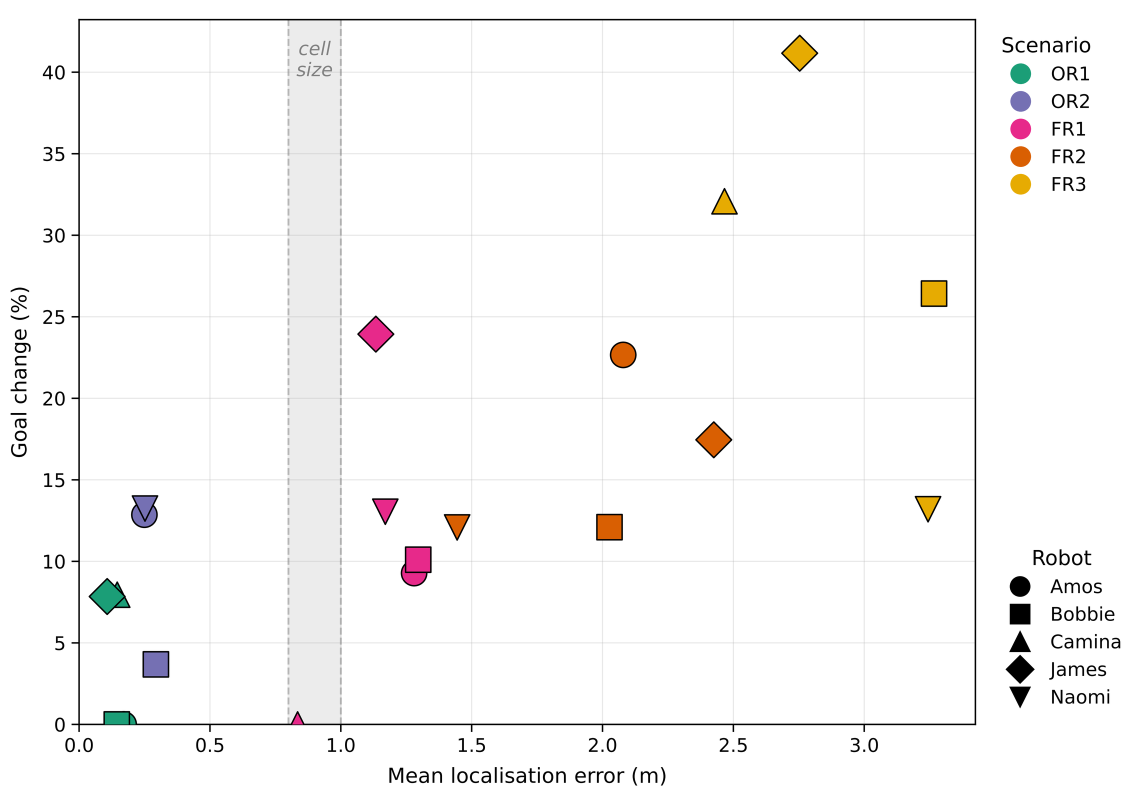

The grey band marks the **0.80 – 1.0 m** occupancy cell size — the same scale at which the algorithm reasons. Points left of this band mostly preserve the assigned cell. Points far to the right show the failure regime: when mean error grows to several cell widths, the selected frontier itself changes.

In **FR1**, a single robot operated with mean error **0.84 m** on a **0.80 m** cell grid — accuracy in proportion with cell resolution. It recorded **zero goal changes**: the only run at LUNA where the audited assignment matched the optimised reference perfectly, and the only robot whose localisation stayed at the cell scale throughout. The boundary marked on the chart is where the data actually crosses over.

G · SENSITIVITY VS. IMPACT

Most rank changes never become goal changes

Rank changed: sensitivity of the distance computation.

Goal changed: whether that sensitivity propagated to an actual assignment.

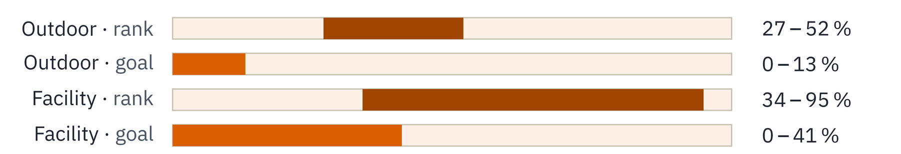

Outdoors, **70 – 100 %** of rank reorderings never flip the claimed goal — only the **single closest uncontested cell** matters. At LUNA, that silent margin shrinks: once drift exceeds the cell-claim distance, even the top-ranked cell starts to flip.

H · SUMMARY

Run summary & design requirements

Run	Robots	Mean error	Rank chg.	Goal chg.
Outdoor field				
OR1	4	0.11 – 0.17 m	27 – 40 %	0 – 8 %
OR2	3	0.25 – 0.29 m	31 – 52 %	1 – 13 %
LUNA facility · 5 robots				
FR1	5	0.84 – 1.30 m	34 – 86 %	0 – 24 %
FR2	5	1.45 – 2.43 m	40 – 77 %	12 – 24 %
FR3	5	**2.47 – 3.27 m**	**54 – 95 %**	**13 – 41 %**

Reading the LUNA numbers. At LUNA, peer deconfliction was offline — so what we measured is pose-error sensitivity alone. The missing map-frame errors would only compound those rates, making the reported FR1–FR3 numbers conservative.

1. Pose accuracy must stay within cell resolution.

The coordination layer inherits its tolerance from the occupancy-grid cell size — not from the SLAM specification alone — and that tolerance must hold *throughout* the run, not just at start-up.

2. Occupancy grids must share a common reference frame.

Accurate poses alone don't fix a misaligned frontier set. As each robot's local frame slides off the shared map, the same grid cell ends up referring to different physical locations.

FUNDING European Space Agency **TerraScout** (grant 4000143611/24/NL/MH/mp) · Slovenian Research and Innovation Agency (grant P2-0270). · The authors thank the ESA **LUNA** team for facility access and experimental support.