



Robocup 2025

Salvador, Bahia, Brazil

Mapping Challenge Details


The Labyrinth Ground Truth - Prelims & Semi Finals

Generated via Leica BLK Handheld Lidar



The Labyrinth Ground Truth - Finals

Generated via Leica BLK Handheld Lidar



The Labyrinth Ground Truth

- Generated via Leica BLK Handheld Lidar
- Identified objects were manually located and labelled
- For finals, we changed the maze entrance, which added an internal loop.
- When mapping for finals, a flashlight was used to assist with the color data in the resulting point cloud (it was very dark otherwise, as seen in the first ground truth). The color data made it easier to label identification object locations.
- The labyrinth was covered with tarps to prevent the mapping of the venue ceiling, which was considered to be outside the mapping area.

Mapping Scores

There were two scores which equally contribute to overall mapping score

- 3D Point Cloud Score (3D)
- Identification Score (ID)

Differing from the German Open, the ID score was normalized based on the total available detection points, so both 3D and ID scored range from 0-100, and they can be added together for a total maximum of 200 points per mapping run (without multipliers)

Mapping Scores

- Multipliers
 - Score multipliers were implemented but not used by any participating team
 - Autonomous map (4x) - This should be defined more concretely in the next rules iteration. Only one team attempted, but they were unable to produce a map.
 - 3D Color Data (1.25x) - No teams provided color information.
 - 3D Thermal Data (1.6x) - No teams provided heat information.

Mapping Scores

- 3D Point Cloud Score

- Average Cloud Distance (Error / E) - Evaluation map to Ground Truth map. For each point in the evaluation map, it finds the closest distance to a point in the ground truth map, and then averages them all together. (In Meter units)
- Completeness (C) - Looks at distance the other way, from ground truth to evaluation map, and calculates the percentage of ground truth points which are within a threshold distance to a point in the evaluation map. A 5cm threshold was used. This represents the percentage of the ground truth map
- Clouds must be properly aligned and cropped to only include the mapping area of interest to be scored appropriately. This an origin point and reference axis was provided, but several teams did not use it properly to orient their map so these corrections were made to submissions by the judge. It's not too much extra work, but as teams get more comfortable we should make it the team's responsibility to provide the initial map orientation before we apply ICP.

Score Calculation: $C / (1+E)$

Mapping Scores

- Identification Score

- Total points (P) - Each type of object detection was worth a specified point amount. If the detection was present in the submitted csv file AND below a determined error threshold, then the points were awarded and summed.
- Available Points (A) - Total points available from identifiable objects
- Average Error (E) - The average error of all submitted detections. This influences the score calculations so that a team with equal detection points can pull ahead if their detections are more accurate. (In Meter units)
- For this evaluation to be as smooth as possible, label formatting should be strictly enforced. This year every team required minor formatting adjustments to be scored correctly.
- If everything is formatted correctly, the scoring is instant and requires no extra work.
- Different from the German Open, points were normalizing based on the total available points.

Score Calculation: $(P / (1+E)) / A$

Basic Rules

- No victim box
- No dexterity
- 3 kinds of identification targets (Apriltags = 1pt, Hazmat = 2pts, Real Objects = 10 pts)
- 20 identification objects for prelims-semifinals (10 Apriltag, 5 Hazmat, 5 Real Objects)
- 30 identification objects for finals (20 Apriltag, 5 Hazmat, 5 Real Objects)
- 20 minutes for mapping
 - Teams may reset as many times as they want, and create as many maps as they want
 - Teams were not required to start mapping at any particular point in the labyrinth
- 5 minutes to submit their files for scoring
- The rules specify up to 4 map submissions for semi finals and finals. I restricted it to 3, and think it should stay this way to encourage autonomy.

Submissions

During Preliminary rounds, teams were allowed to create as many maps as they wanted within the 20 minutes mapping period, but they were only allowed to submit files for one, which they believed to be their best map. This encourages cautious driving and thorough mapping in earlier stages.

During semi finals and finals, teams were allowed to submit up to 3 maps within the 20 minute period, which rewards faster mapping, but only if the maps don't become significantly worse as a result. The scores for each mapping run are then summed together.

Observations

In the preliminary rounds, very few teams were prepared for 3D mapping, but by finals we got submissions from almost every team. Most teams submitted only a 3D map, and a few teams submitted Identifications as well to varying degrees of success.

In finals, there were 16 3D map submissions from 6 teams, and 4 Identification list submissions from 2 teams.

Preemptive measures were taken to secure the structure of the maze as much as possible. Securing K-rails throughout the labyrinth helped a lot with is, and prevented the structure from shifting as robots bumped into things, which helped us reuse the ground truth for several days.

Currently, no team is able to fully “solve” the mapping challenge, and there is still lots of room for growth in the scores. Still, there are several ways we could continue to increase the difficulty, such as terrain complexity, 3D/vertical navigation elements, increasing labyrinth size and complexity, tightening tolerances for scoring, and shifting towards more object detections and less AR tag markers.

Initially some teams were confused about the scoring process, which is somewhat expected for a new system, so maybe the rules can be defined a bit better, but each team seemed to understand the rules, expectations, and scoring after I explained it.

Objects Used

- Baby Doll
- Gas Tank
- Fire Extinguisher
- Hard Hat
- Backpack

Map Slices Gallery

Round: Finals

Team: Dynamics

Map 1: 92/100


Map 2: 95/100

Map 3: 95/100


Total 3D Score: 282

Also submitted object identifications for AR tags and hazmats in each map, but no real objects


Collected in the 20 minute mapping period.




[Finals] Dynamics - Map 1 [Sliced at 1.7m +/- 0.1][HIGH]




[Finals] Dynamics - Map 1 [Sliced at 1.0m +/- 0.2][LOW]




[Finals] Dynamics - Map 2 [Sliced at 1.7m +/- 0.1][HIGH]



[Finals] Dynamics - Map 2 [Sliced at 1.0m +/- 0.2][LOW]



[Finals] Dynamics - Map 3 [Sliced at 1.7m +/- 0.1][HIGH]



[Finals] Dynamics - Map 3 [Sliced at 1.0m +/- 0.2][LOW]

Map Slices Gallery

Round: Finals

Team: AutonOhm

Map 1: 82/100


Map 2: 89/100

Map 3: 85/100

Total 3D Score: 256

No identified objects submitted

Collected in the 20 minute
mapping period.



Map Slices Gallery

Rounds: Finals

Team: Solidus

Map 1: 70/100


Map 2: 75/100

Map 3: 83/100


Total 3D Score: 228

No identified objects submitted


Collected in the 20 minute
mapping period.




[Finals] SOLIDUS - Map 1 [Sliced at 1.7m +/- 0.1][HIGH]




[Finals] SOLIDUS - Map 1 [Sliced at 1.0m +/- 0.2][LOW]




[Finals] SOLIDUS - Map 2 [Sliced at 1.7m +/- 0.1][HIGH]



[Finals] SOLIDUS - Map 2 [Sliced at 1.0m +/- 0.2][LOW]



[Finals] SOLIDUS - Map 3 [Sliced at 1.7m +/- 0.1][HIGH]




[Finals] SOLIDUS - Map 3 [Sliced at 1.0m +/- 0.2][LOW]

Map Slices Gallery


Finals
Team NiTRO
Map 1

Submitted one map, but had the best object identification score by a large margin, including real objects. (Locations shown overlaid on ground truth map below)


Collected in the 20 minute mapping period.



[Finals] NiTRO - Map 1 [Sliced at 1.7m +/- 0.1][HIGH]



[Finals] NiTRO - Map 1 [Sliced at 1.0m +/- 0.2][LOW]




Additional Experiments

SkyeBrowse was used with Insta360 Video to reconstruct the mapping area.


Supplemental lighting (via small touch lights) needed to be dispersed in the mapping area for this method to perform best (official instructions state the importance of good lighting).

Three videos were recorded, walking at different speeds through the area (1, 3, and 5 minutes). Metrics clearly showed longer videos led to more accurate reconstructions, but they also took longer to process (30+ minutes to upload and process for the 5 minute video).


The resulting maps were not accurate enough to use as a competition ground truth, but would absolutely be of use to end users and first responders.




[SkyeBrowse] 1 Minute Video[Sliced at 1.7m +/- 0.1][HIGH]




[SkyeBrowse] 1 Minute Video[Sliced at 1.0m +/- 0.2][LOW]




[SkyeBrowse] 3 Minute Video[Sliced at 1.7m +/- 0.1][HIGH]



[SkyeBrowse] 3 Minute Video[Sliced at 1.0m +/- 0.2][LOW]



[SkyeBrowse] 5 Minute Video[Sliced at 1.7m +/- 0.1][HIGH]



[SkyeBrowse] 5 Minute Video[Sliced at 1.0m +/- 0.2][LOW]