NEURAL INFORMATION PROCESSING SYSTEMS
2026 · Evaluations & Datasets Track

SnT uni.lu Université du Luxembourg

ECLIPSE: A Spacecraft Rendezvous Trajectories Dataset with Controlled In-Orbit Lighting Conditions

Nidhal Eddine Chenni Arunkumar Rathinam Abid Ali Djamila Aouada

SnT, University of Luxembourg

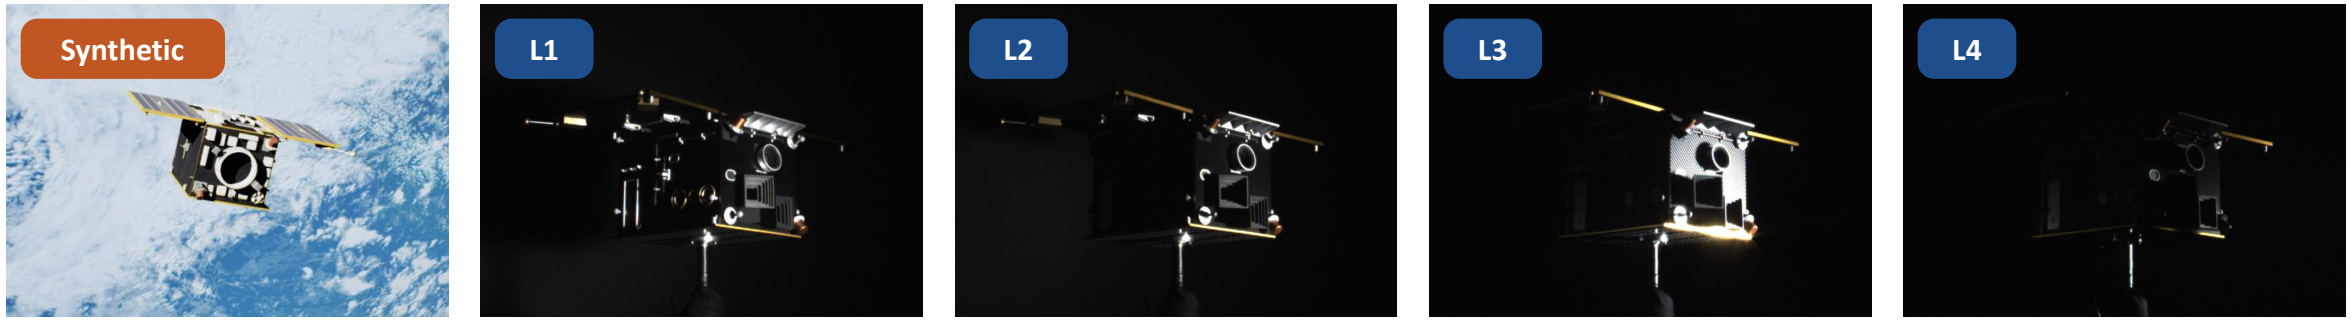

Same trajectory, same pose, four controlled lighting conditions

Why spacecraft pose estimation needs better data

Pose drives proximity operations

Rendezvous, inspection and docking need the 6-DoF pose of the target from a single camera.

Real in-orbit images are scarce

Acquiring them is expensive and risky, so models are trained mainly on synthetic renders.

The sim-to-real gap

Models trained on synthetic data degrade on real images. Domain adaptation is still an open problem.



Train: synthetic render

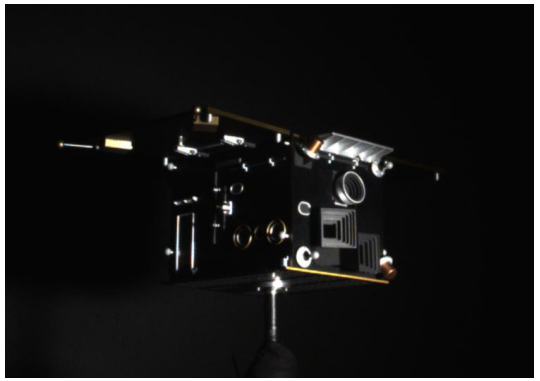

Deploy: real camera image

Rotation error on real L1 images

65.2° trained on synthetic vs **2.0°** trained on real

EfficientNet-B5 keypoint model, mean of 3 runs

In orbit, lighting changes along the trajectory

- Illumination is not random: it follows the motion of the target and its position along the orbit.
- The phase angle between the Sun and the camera decides which faces are lit, shadowed or specular.
- Sunlit and dawn-dusk parts of the orbit give very different images of the same target.

Key question

How much of the domain gap is caused by lighting alone, when the pose stays the same?

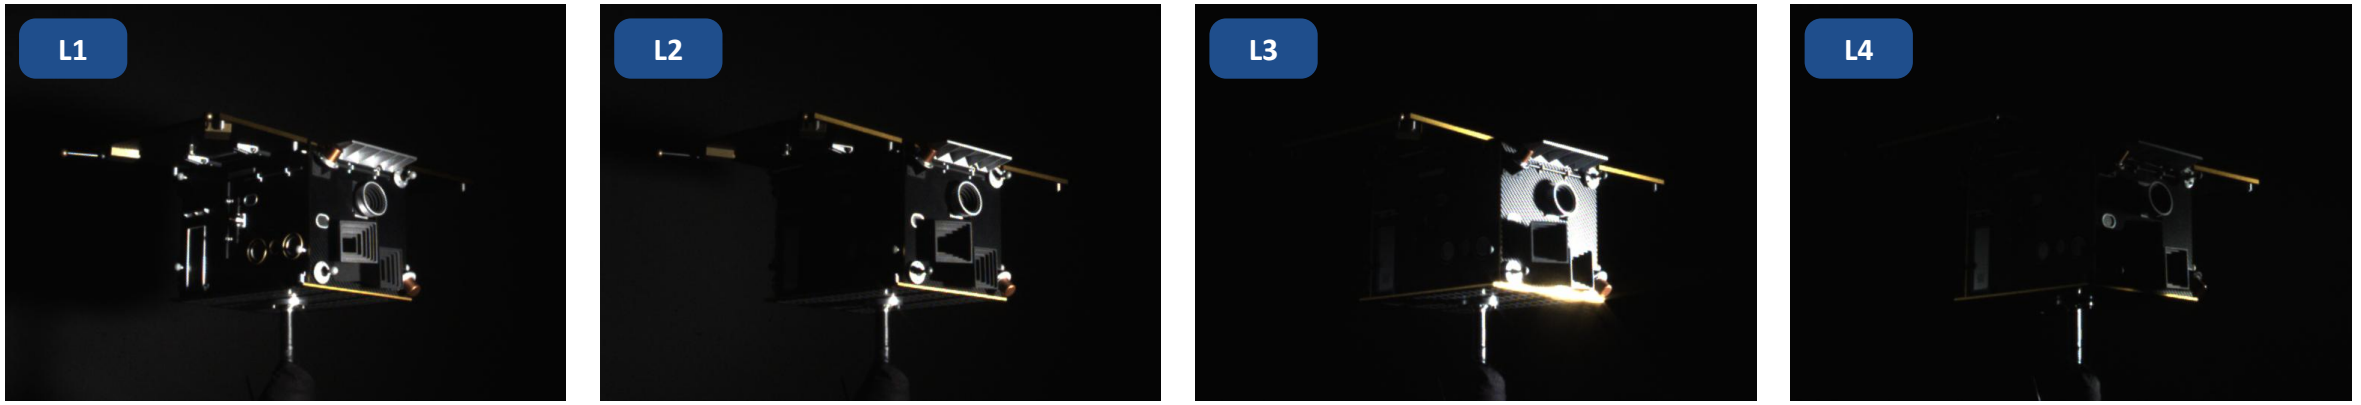

Light moves away from the camera axis in 30° steps: less visible surface, longer shadows

Existing benchmarks do not isolate lighting

Property	SPEED / SPEED+	SHIRT	SPADES	SPARK	DLVS3	ECLIPSE
Synthetic data	✓	✓	✓	✓	✓	✓
Robotic testbed (HIL)	✓	✓	✓	✓	-	✓
Trajectories	-	✓	✓	✓	-	✓
Real lighting setups	2	1	Fixed	Not controlled	-	**4, controlled**
Lighting-paired trajectories	-	-	-	-	-	✓
Resolution	1920×1200	1920×1200	1440×1080	1440×1080	1024×1024	**1440×1080**
Images	15K / 70K	4.7K	179.4K	120K	1M	**~270K**

SPEED+

2 lighting setups (Lightbox, Sunlamp), poses sampled at random, no trajectories.

SHIRT

Rendezvous trajectories, but a single Lightbox condition.

SPARK'24 / SPADES

Real trajectories, but lighting fixed or not controlled across them.

ECLIPSE at a glance

240,000
synthetic images
400 trajectories × 600 frames

30,800
real HIL images
14 trajectories × 4 lightings × 550 frames

4
lighting setups
L1 to L4, 30° steps

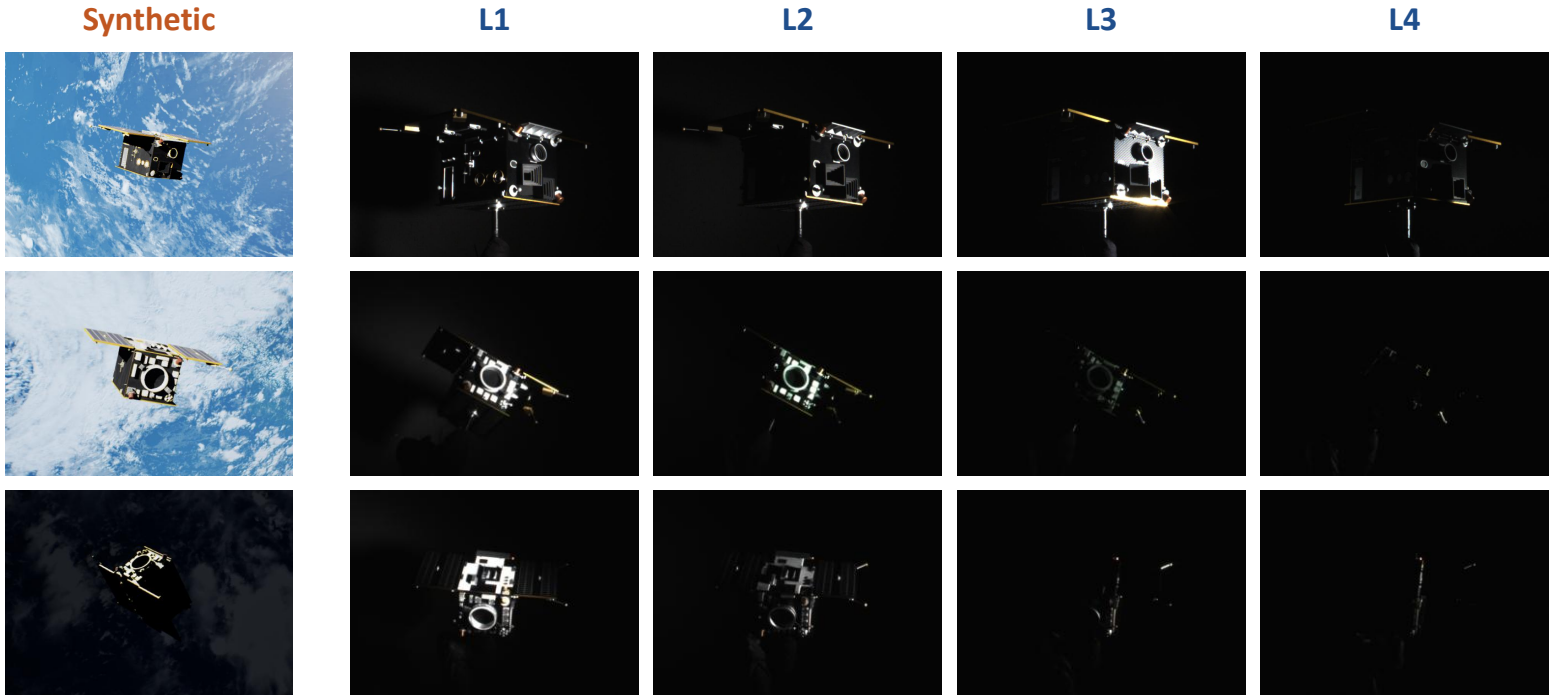

Key idea

Every real trajectory is replayed with the same robot motion under each lighting. Pose labels stay fixed, so lighting is the only variable.

Two axes of domain shift

- Sim-to-real: synthetic vs all real images
- Within-real: L1 vs L2 vs L3 vs L4

Synthetic split: Unreal Engine 5 pipeline

1

Orbital viewpoints

2,000 camera positions on a sphere at 800 to 1,000 km altitude, grouped into 4 lighting regimes by phase angle (full to near-eclipse).

2

Trajectories

5 motion types: inspection spheres and cylinders, lemniscate fly-arounds, Clohessy-Wiltshire drift ellipses, approach-retreat passes.

3

Rendering and labels

Proba-2 CAD model, 16k NASA Blue Marble Earth, atmospheric scattering. Exact pose exported for every frame.

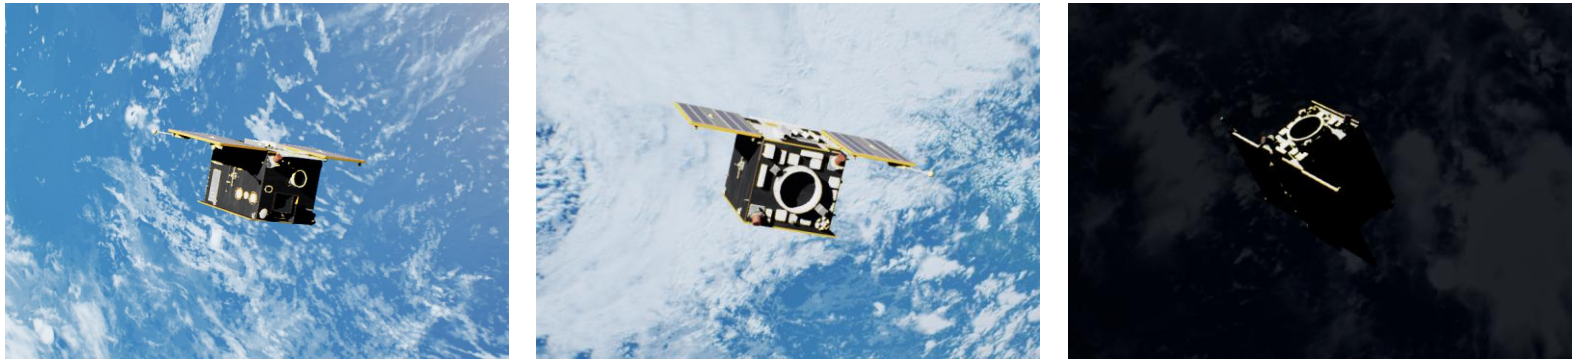

Example renders under different Sun phase angles

400 × 600 sequences × frames

11.8° to 178.6° Sun phase angles

320 / 80 train / val trajectories

Real split: hardware-in-the-loop in the Zero-G Lab



Proba-2 mockup, sun lamp and robotic arms in the SnT Zero-G Lab

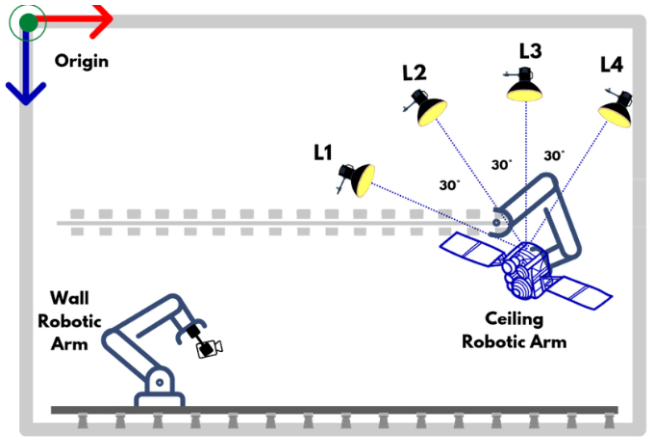

Four lamp positions

L1 lights the target from the camera side. Each next position adds 30°, until L4 is near opposition and most of the visible face is in shadow.

1:2.5

Proba-2 mockup

0.64 × 0.24 × 0.415 m, about 7 kg, materials matched to the CAD model

1440 × 1080

FLIR Blackfly S camera

6 mm lens, about 44.8° horizontal field of view

120,000 lux

Sun-like LED lamp at 5800 K

Close to LEO sunlight (about 163,000 lux)

Real split: trajectories and lighting shift

14 trajectories × 4 lightings = 56 sequences of 550 frames (+ 6 unpaired)

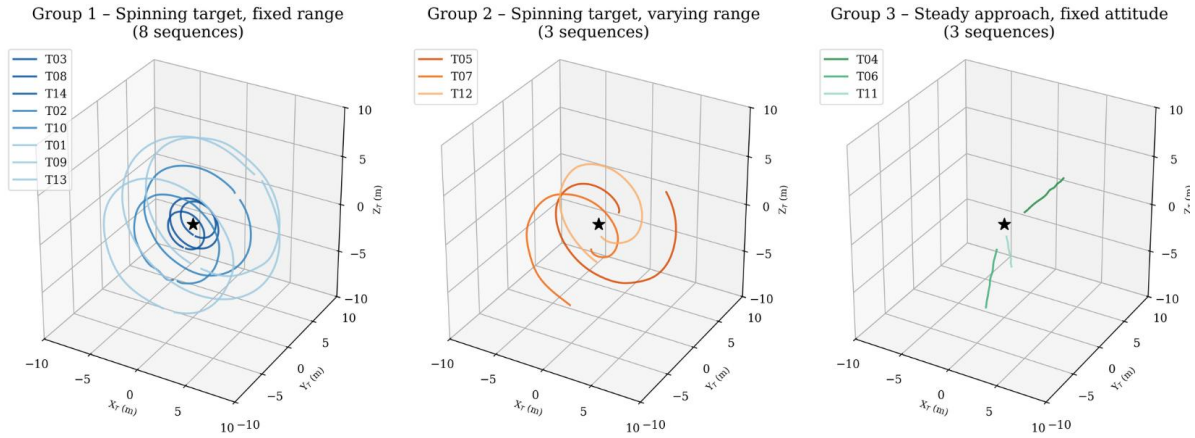

Real camera paths around the target (star) for the three motion groups.

G1 Spinning target, fixed range
8 seq. at about 2.1, 5.5 and 8.2 m, one full revolution

G2 Spin + approach
3 seq., 8 m to 2 m while spinning: all 6-DoF change

G3 Steady approach
3 seq., attitude almost fixed (≤ 3°): isolates range

Same poses, so the image change is purely lighting

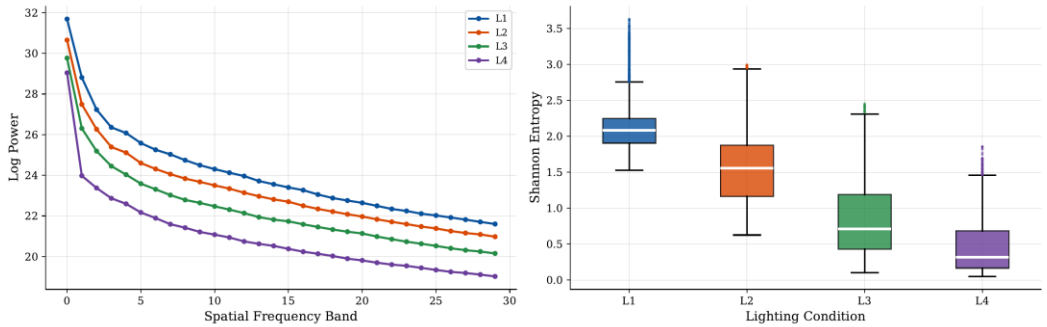

Frequency spectrum (left) and Shannon entropy (right) of target crops, over the 14 trajectories

Ordered shift: Both measures follow L1 > L2 > L3 > L4.

Fine detail disappears: High frequencies drop step by step. L4 keeps mostly coarse shadow.

Distinct regimes: L1 and L4 entropy ranges do not overlap.

Within the real domain, lighting alone breaks the model

Rotation error E_R (°), train on one lighting, test on the others

Train \ Test	L1	L2	L3	L4
L1	-	5.7°	67.3°	100.1°
L2	2.2°	-	33.7°	80.6°
L3	34.4°	19.8°	-	19.4°
L4	59.6°	41.2°	4.4°	-

EfficientNet-B5 keypoint heatmaps + PnP, mean of 3 runs

Same poses, only the light moved

5.7° L1 → L2

100.1° L1 → L4

- Error grows with the angular distance between training and test lighting.
- Neighbouring setups transfer well (L4 → L3: 4.4°).
- Translation follows the same trend (L1 → L4: E_T 0.606 vs 0.021 for L1 → L2).

Sim-to-real: adaptation helps, but the gap remains large

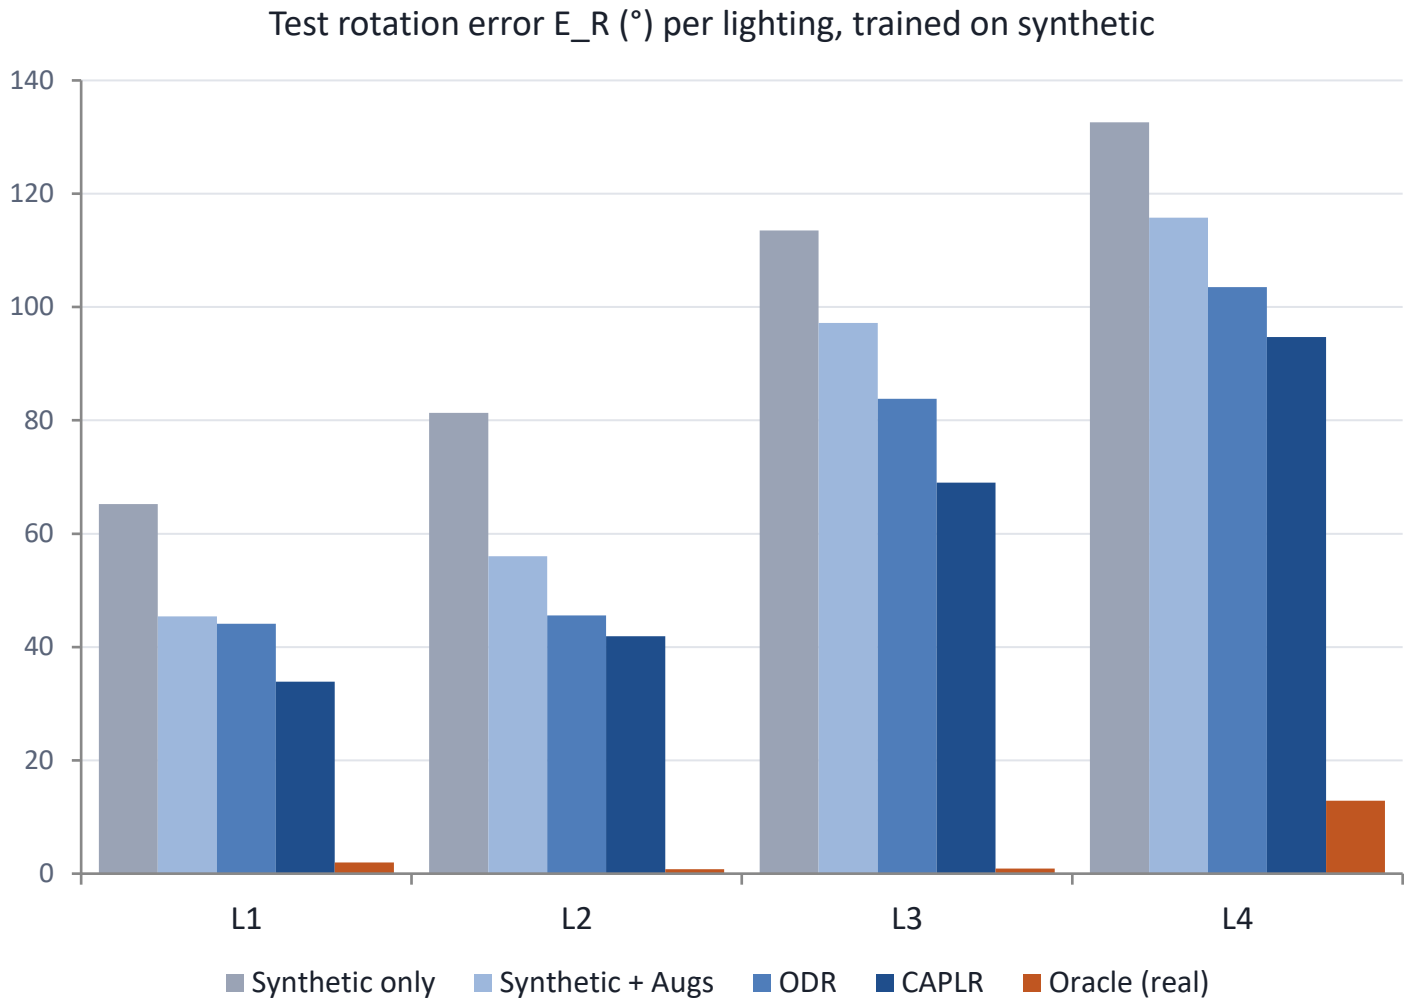

Harder light, larger error

Every method degrades from L1 to L4.

Augmentation and DA help

CAPLR gives the lowest rotation error on all lightings, ODR the lowest translation error on L2 to L4.

Large headroom

The oracle stays below 13°, so the real split is learnable and the gap is open.

Not fixed by bigger models

A DINOv3 temporal model cuts L1 error to 12.2°, but L4 stays above 100°. Larger EfficientNets give no consistent gain.

What can ECLIPSE be used for?

Lighting-robust pose estimation

Train on some lightings and test on unseen ones, with poses held fixed.

Domain adaptation

Unsupervised, source-free or test-time adaptation, with one target domain per lighting.

Temporal models and tracking

550-frame sequences for video pose estimation, filtering and tracking.

Uncertainty and failure detection

Learn when a prediction cannot be trusted, for example under backlight.

Lighting-aware guidance

Study which approach geometry gives the most informative views for navigation.

Beyond pose

Released CAD model and tools give masks, depth and surface maps for real images.

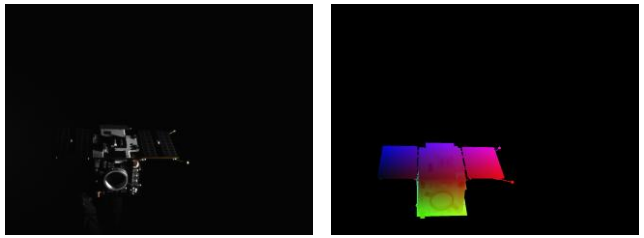

Real image and its rendered surface map

Limitations and next steps

Laboratory conditions do not fully reproduce space. Lighting is sampled at 4 discrete positions on a single target (Proba-2). Finer lighting steps and more targets are natural extensions.

Takeaways

1

Lighting is a controlled variable

14 real trajectories, each replayed under 4 lightings, plus 240,000 synthetic images.

2

Lighting is a major bottleneck

With poses fixed, errors grow by up to an order of magnitude from adjacent to opposing light, for both CNN and ViT models.

3

The sim-to-real gap is far from closed

Augmentation and domain adaptation help, but a large margin to the oracle remains.

Dataset, code and weights

Dataset **Project page**

chenni-nidhaleddine.github.io/ECLIPSE

Thank you!

nidhal.chenni@uni.lu

Funded by the Luxembourg National Research Fund (BRIDGES/2025/IS/19855391/U-ADAPT) and Infinite Orbits. Experiments on MeluXina (LuxProvide).