IRIS-VIS: A New Dataset for Visibility Estimation in an Industrial Environment

WACV 2025
TUCSON, ARIZONA • FEB 28 - MAR 4

Flavien Armangeon[1,2] Thibaud Ehret[1] Enric Meinhardt-Llopis [1] Rafael Grompone von Gioi [1] Guillaume Thibault [2] Marc Petit[2] Gabriele Facciolo[1]

[1]Université Paris-Saclay, CNRS, ENS Paris-Saclay, Centre Borelli, France [2] EDF R&D (Electricité de France), France

https://centreborelli.github.io/iris-vis

Context

- **Point visibility estimation:** Useful for point cloud rendering, surface reconstruction and 3D segmentation from images.
- **Previous work:** Focus on objects, outdoor scenes or synthetic indoor scenes. No quantitative evaluation is given on real indoor data.
- **This work:** Introduces a new dataset, Industrial Room In Saclay (IRIS), useful for many tasks including VISibility estimation (IRIS-VIS) and new metrics.
- **Advantages of IRIS-VIS:** Complex indoor scene, unlimited ground truth quantity, real LiDAR point cloud.

Visibility estimation problem

All points can be seen from a given point of view. The objective is to remove the points that would not be visible in the real scene.

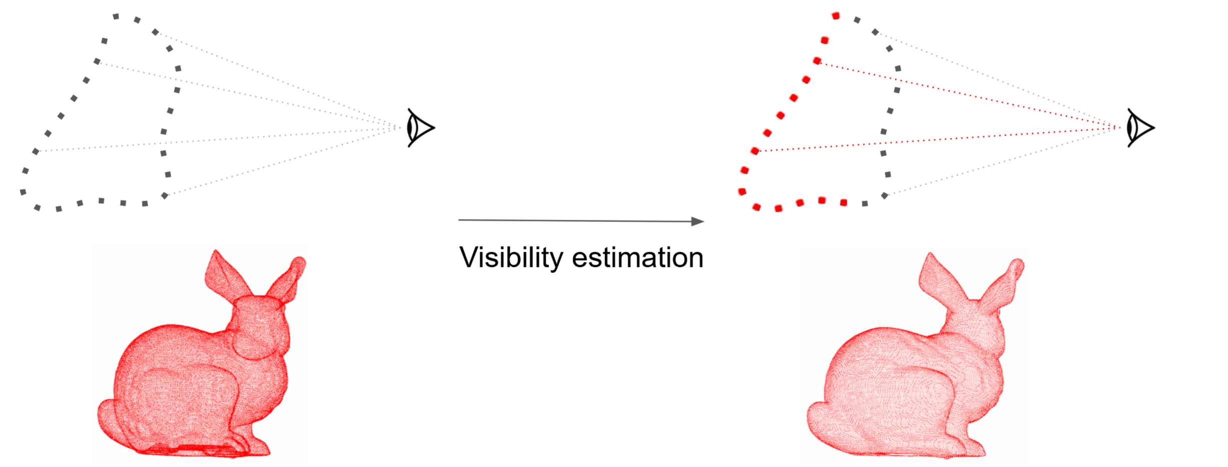

IRIS dataset

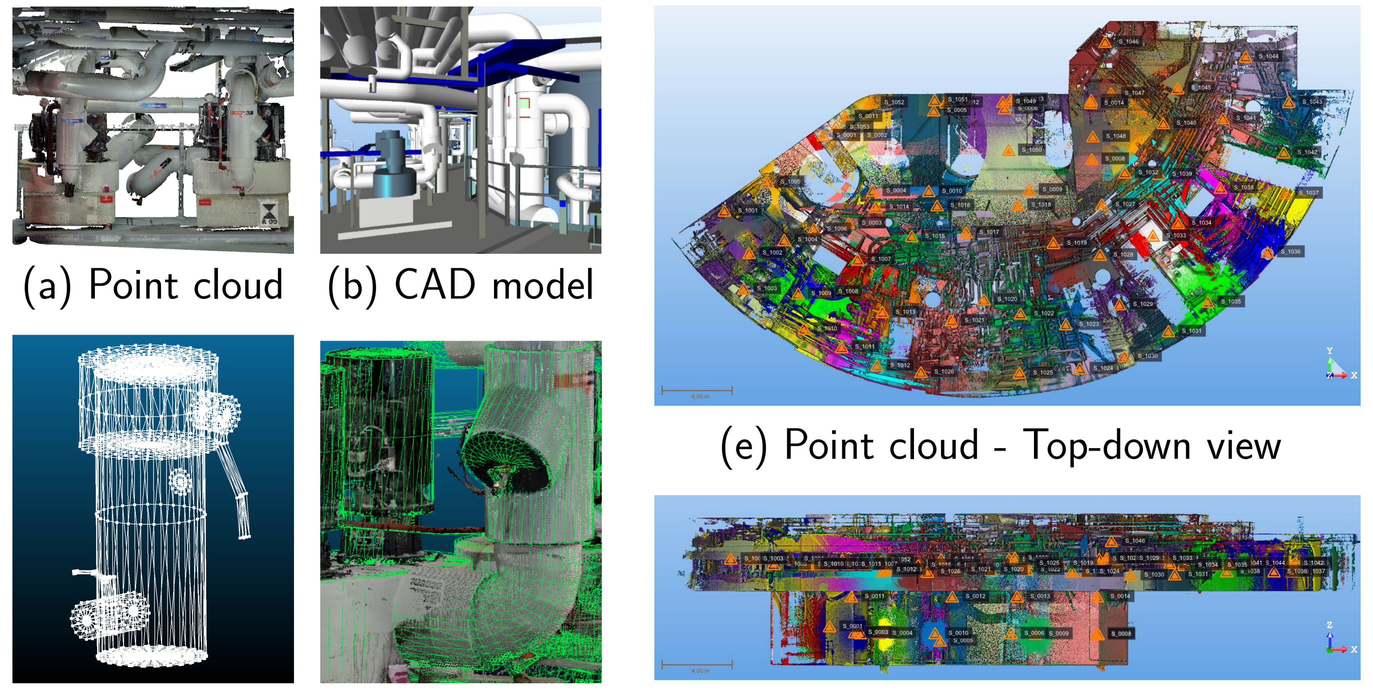

(a) Point cloud (b) CAD model (c) Mesh (d) Mesh & cloud (e) Point cloud - Top-down view (f) Point cloud - Front view

- **Scene:** A large industrial room ($530\mathrm{m}^2$) with complex objects.
- **Point cloud:** High density (2.1 billion points). Merged from 67 LiDAR acquisitions at fixed locations represented by triangles in (e) and (f).
- **CAD model:** Reconstructed close to the points with high precision. The mesh induced by the model (c) is well-paired with the cloud (d).

IRIS-VIS dataset

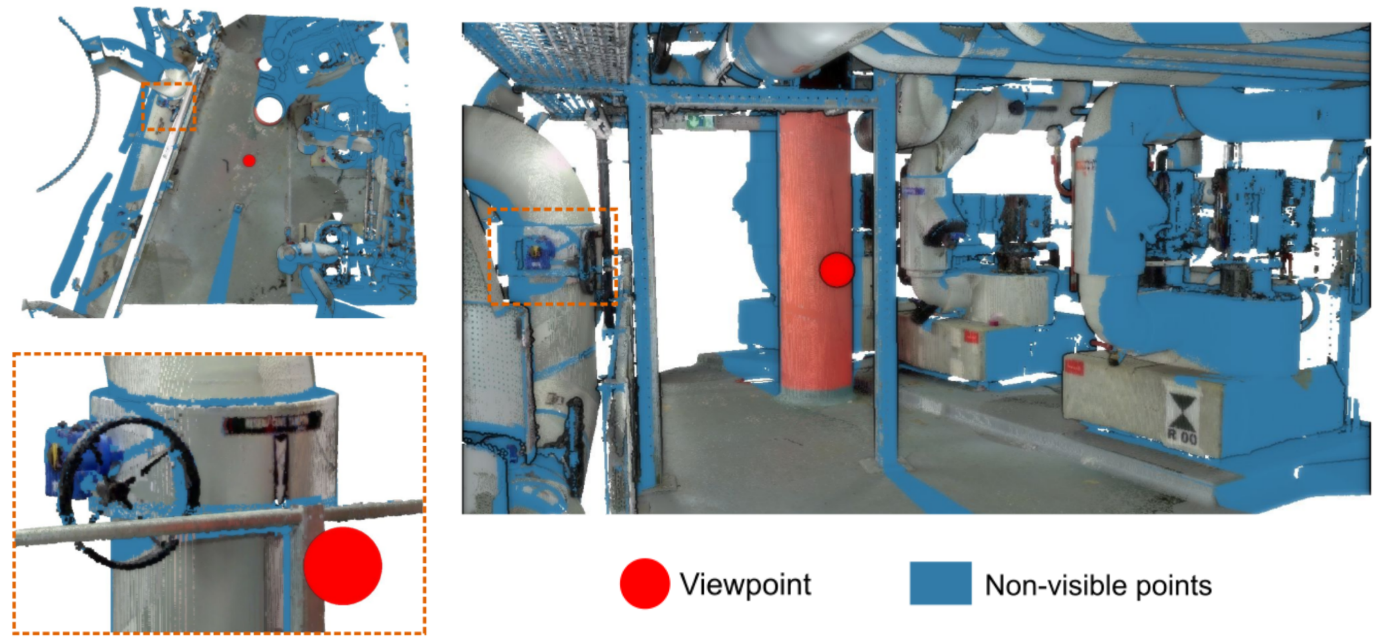

Point cloud and ground truth at a given viewpoint in a subscene of the dataset.

Dataset characteristics

- Unlimited amount of viewpoints thanks to an automatic ground truth construction process.
- Complex and detailed objects such as thin pipes, valves and manometers.
- LiDAR acquisitions with real artifacts, not sampled on a mesh.

Ground truth construction

- **1. Points-to-mesh pairing:** Remove the points too far from the mesh.
- **2. Mesh-to-points pairing:** Divide and remove the mesh triangles too far from the cloud. In (a) and (b), the corrected mesh (orange) is well paired with the points compared to the original mesh (violet).
- **3. Raycasting from the viewpoint to all the points:** The points that are too far from the first hit-point are set as hidden (red) (c).

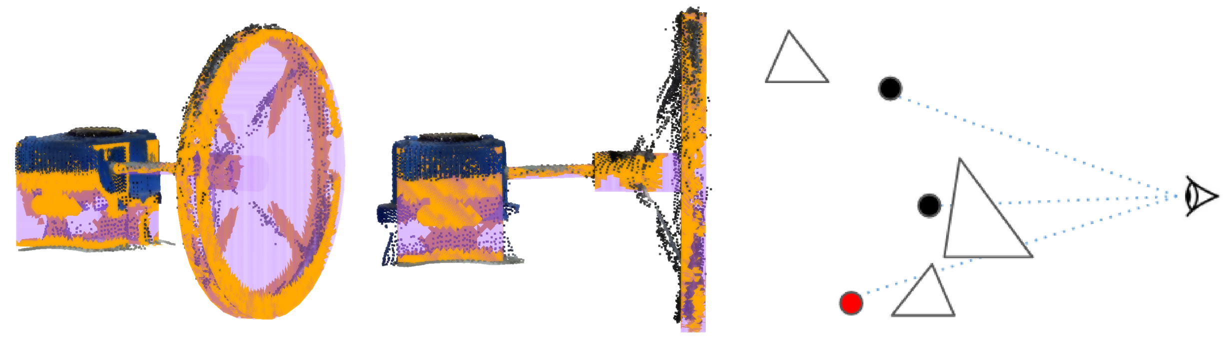

(a) (b) (c)

Complex visibility metrics

We designed new metrics to focus on the areas with a high variability in depth of the visible points in the ground truth, often located at the visibility boundaries.

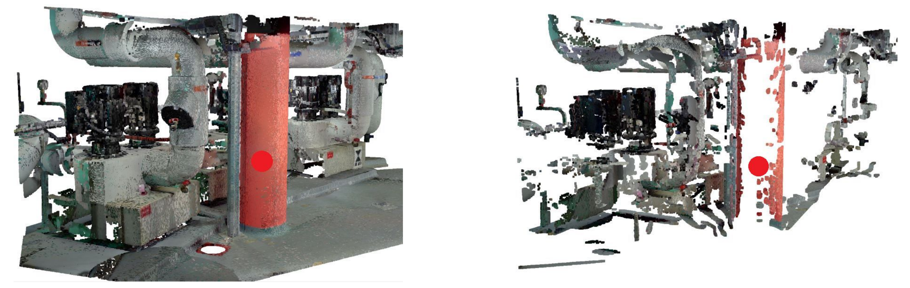

Input cloud Complex visibility areas

Visual results - Sparse cloud

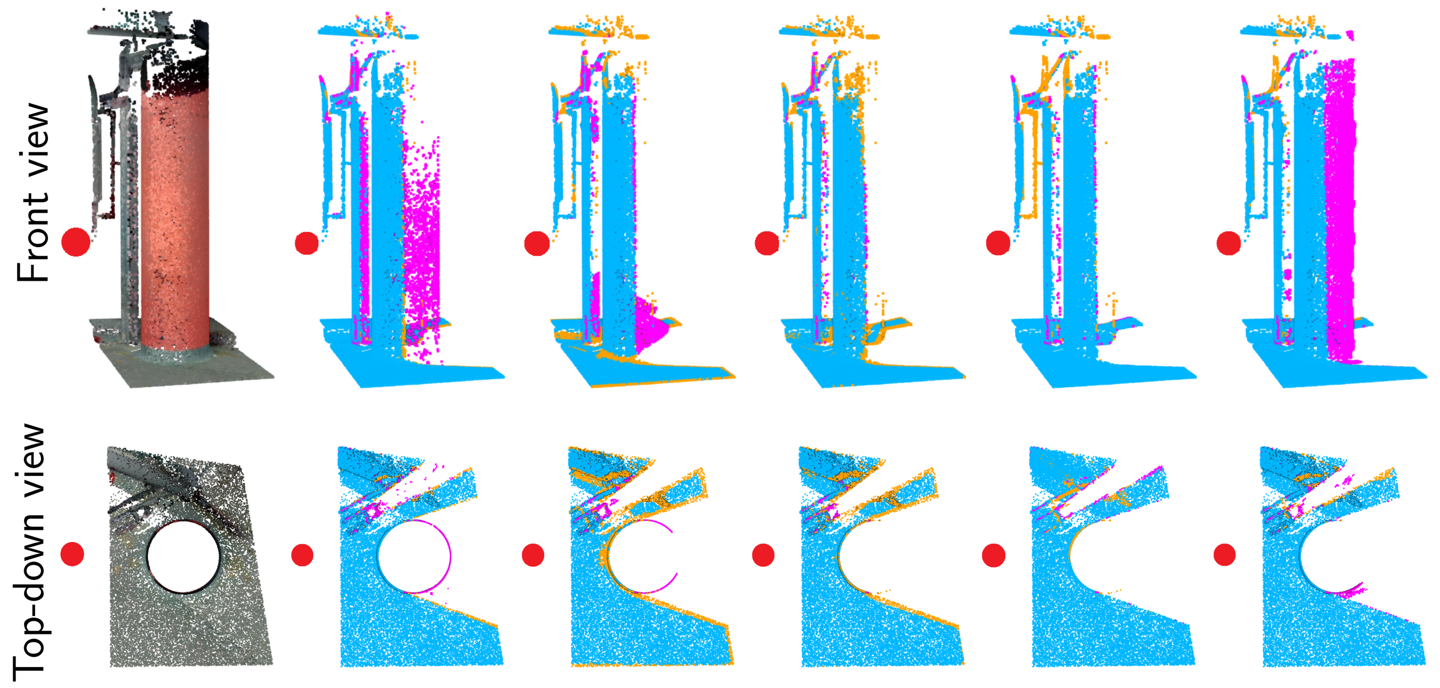

(a) Simple scene with a pole.

True Positives (blue), False Positives (purple), False Negatives (orange).
The positive predictions are the visible points.

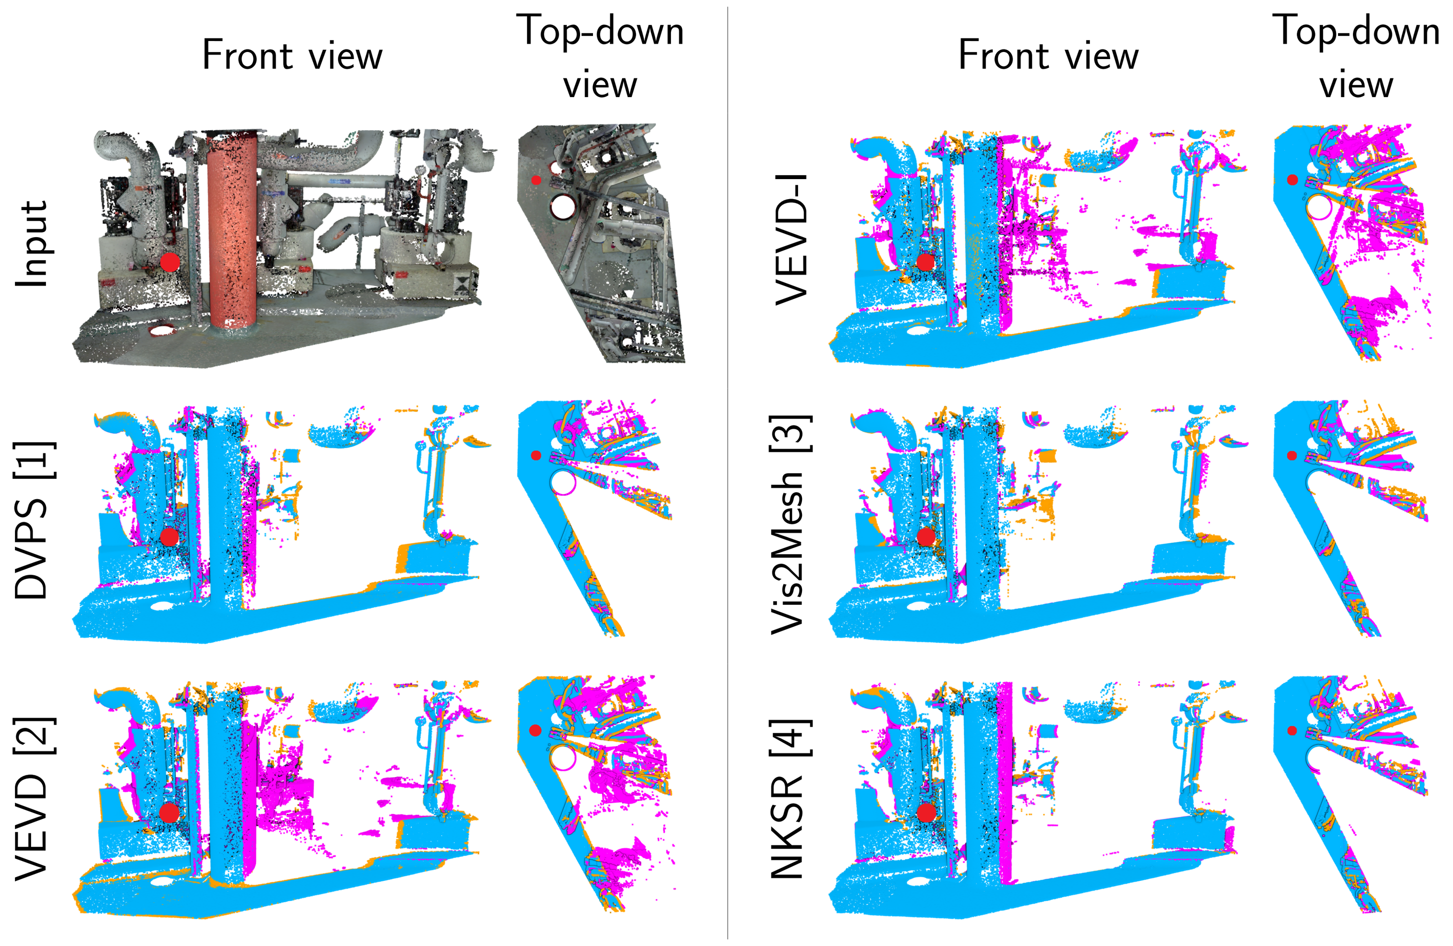

(b) Complex scene with objects behind the pole.

- (a): All methods remove most of the points behind the pole. DVPS and NKSR show more FPs while VEVD, VEVD-I and Vis2Mesh provide more FNs.
- (b): Similar to (a) except for VEVD and VEVD-I which present many FPs behind the pole.
- VEVD-I provides less FPs than VEVD but still more than the other methods.
- The quantitative results on the complex visibility metrics are lower than on the standard metrics and show a considerable scope for improvement.

Quantitative results - Sparse cloud

Method	DVPS [1]	VEVD [2]	VEVD-I	Vis2Mesh [3]	NKSR [4]
Surface reconstruction	X	X	X	✓	✓
t(s)	10^1	10^2	10^2	10^1	10^1
TP	27.27	**31.59**	22.51	27.74	28.75
FP	**1.87**	18.86	9.38	7.02	3.53
FN	7.72	**3.40**	12.48	7.25	6.24
TN	**63.14**	46.16	55.63	57.99	61.49
Precision	**93.58**	62.62	70.59	79.81	89.07
Recall	77.95	**90.28**	64.33	79.29	82.16
Accuracy	**90.41**	77.74	78.14	85.74	90.23
F1-score	85.05	73.95	67.32	79.55	**85.48**

Vis2Mesh and NKSR were run on GPU, the others on CPU.
The positive predictions are the visible points.

Sparse vs dense clouds

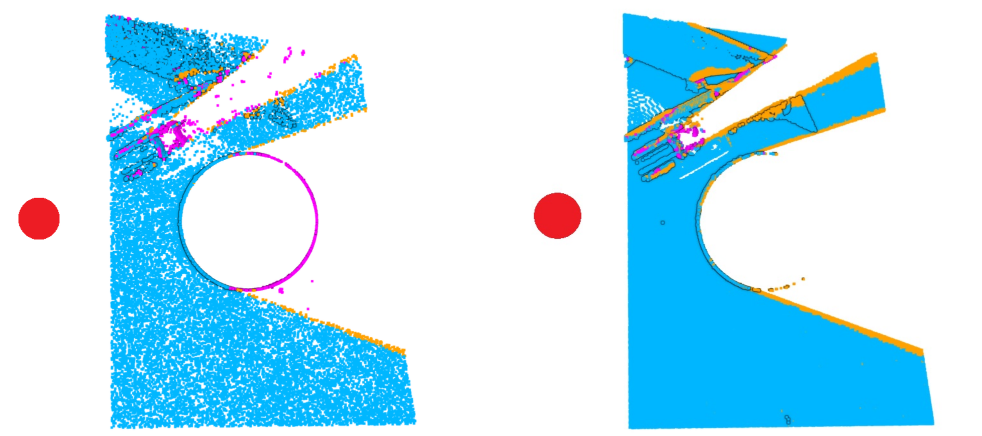

DVPS [1] predictions on a sparse and dense cloud (10 times denser). DVPS is the only method that provide significantly different visualizations accross densities. It shows less FPs on the dense cloud but more FNs. The other methods seem to be robust to the density.

Conclusion

- The current methods struggle to remove the points on challenging areas and densities.
- Some methods require a GPU to run in a reasonable time on large and dense point clouds.

References

[1] Katz, Sagi *et al.* "Direct visibility of point sets". *SIGGRAPH*, 2007.
[2] Biasutti, Pierre *et al.* "Visibility estimation in point clouds with variable density". *VISAPP*, 2019.
[3] Song, Shuang *et al.* "Vis2mesh: Efficient mesh reconstruction from unstructured point clouds of large scenes with learned virtual view visibility". *ICCV*, 2021.
[4] Huang, Jiahui *et al.* "Neural kernel surface reconstruction". *CVPR*, 2023.