Radar Fields: An Extension of Radiance Fields to SAR Supplementary material

Thibaud Ehret[1] Roger Marí[1] Dawa Derksen[2] Nicolas Gasnier[2] Gabriele Facciolo[1]
[1]Université Paris-Saclay, CNRS, ENS Paris-Saclay, Centre Borelli, France
[2]CNES, Toulouse, France

1. Results without noise

We show visual results without noise in Figures 1 and 2.

2. Additional unseen views for the examples shown in the main paper

We show additional unseen views in Figures 3 and 4.

3. Impact of the number of views

We show the impact of the number of views in Figure 5. It shows that more views is better but also that the performance is already very good with six views.

4. Performance when learning the specularity parameter θ

We show in Table 1 the impact of learning θ in terms of mean absolute error with respect to the ground truth. It shows that learning the specularity coefficient θ at the same time does not impact much the overall performance. Surprisingly, learning θ can sometimes (marginally) improve the results. Our intuition is that it is because learning θ at the same time allows to mitigate the impact of the noise. In that case it is probable that it is θ that "overfits" the noise and not the surface.

	θ known and fixed	θ learned (initialized to truth)	θ learned (initialized to any)
keops	0.22606274	0.14900368	0.1464539
stack	0.22562906	0.15893842	0.15840319
Fuji	0.52781403	2.0981526	2.1618936
Piton de la Fournaise	0.22874455	0.28307915	0.28888458

Table 1. Impact of learning θ in terms of mean absolute error (in m) with respect to the ground truth.

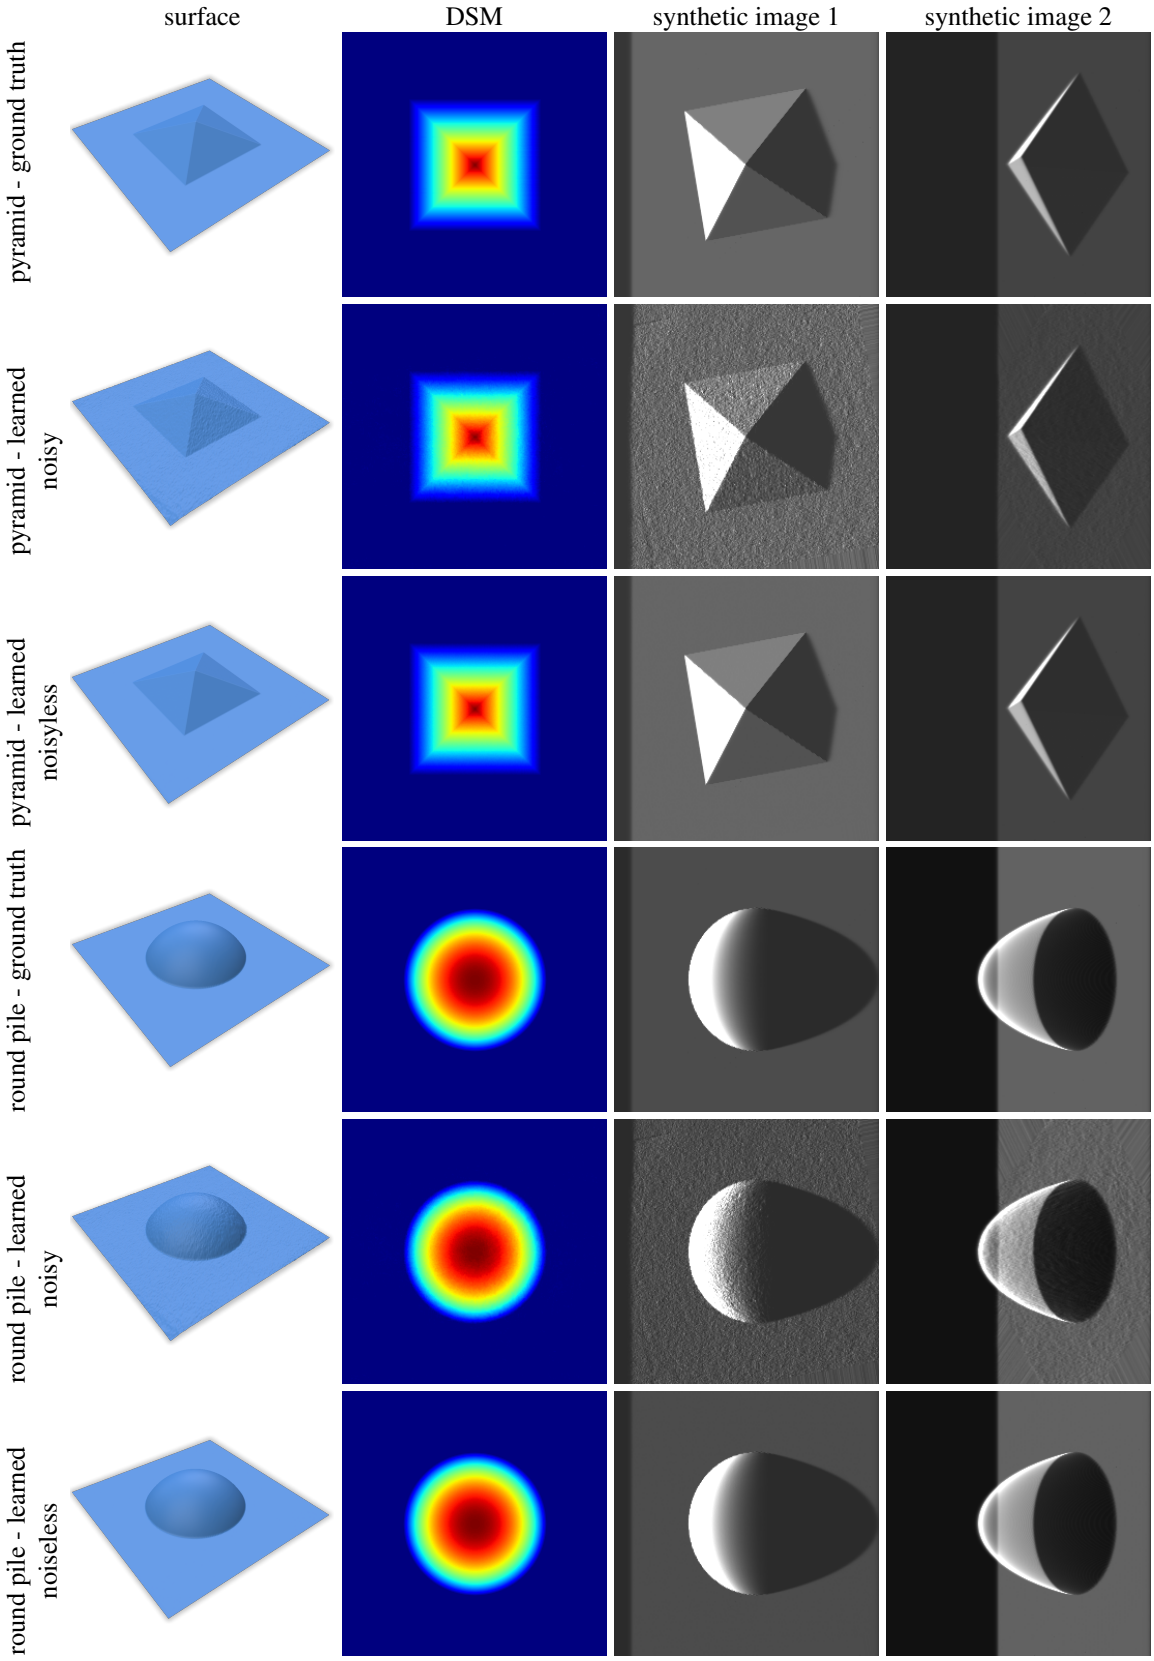

Figure 1. Results of surface learning using five SAR images on two toy examples. The first represents a pyramid (first row is the ground truth, second row the learned model with the noisy views and third row with the noiseless views) and the second a round pile (ground truth in the fourth row, learned model with noisy views in the fifth row and learned model with noiseless views in the last row). From left to right: visualization of the surface model, DSM corresponding to the surface model, and two images corresponding to two views used during training.

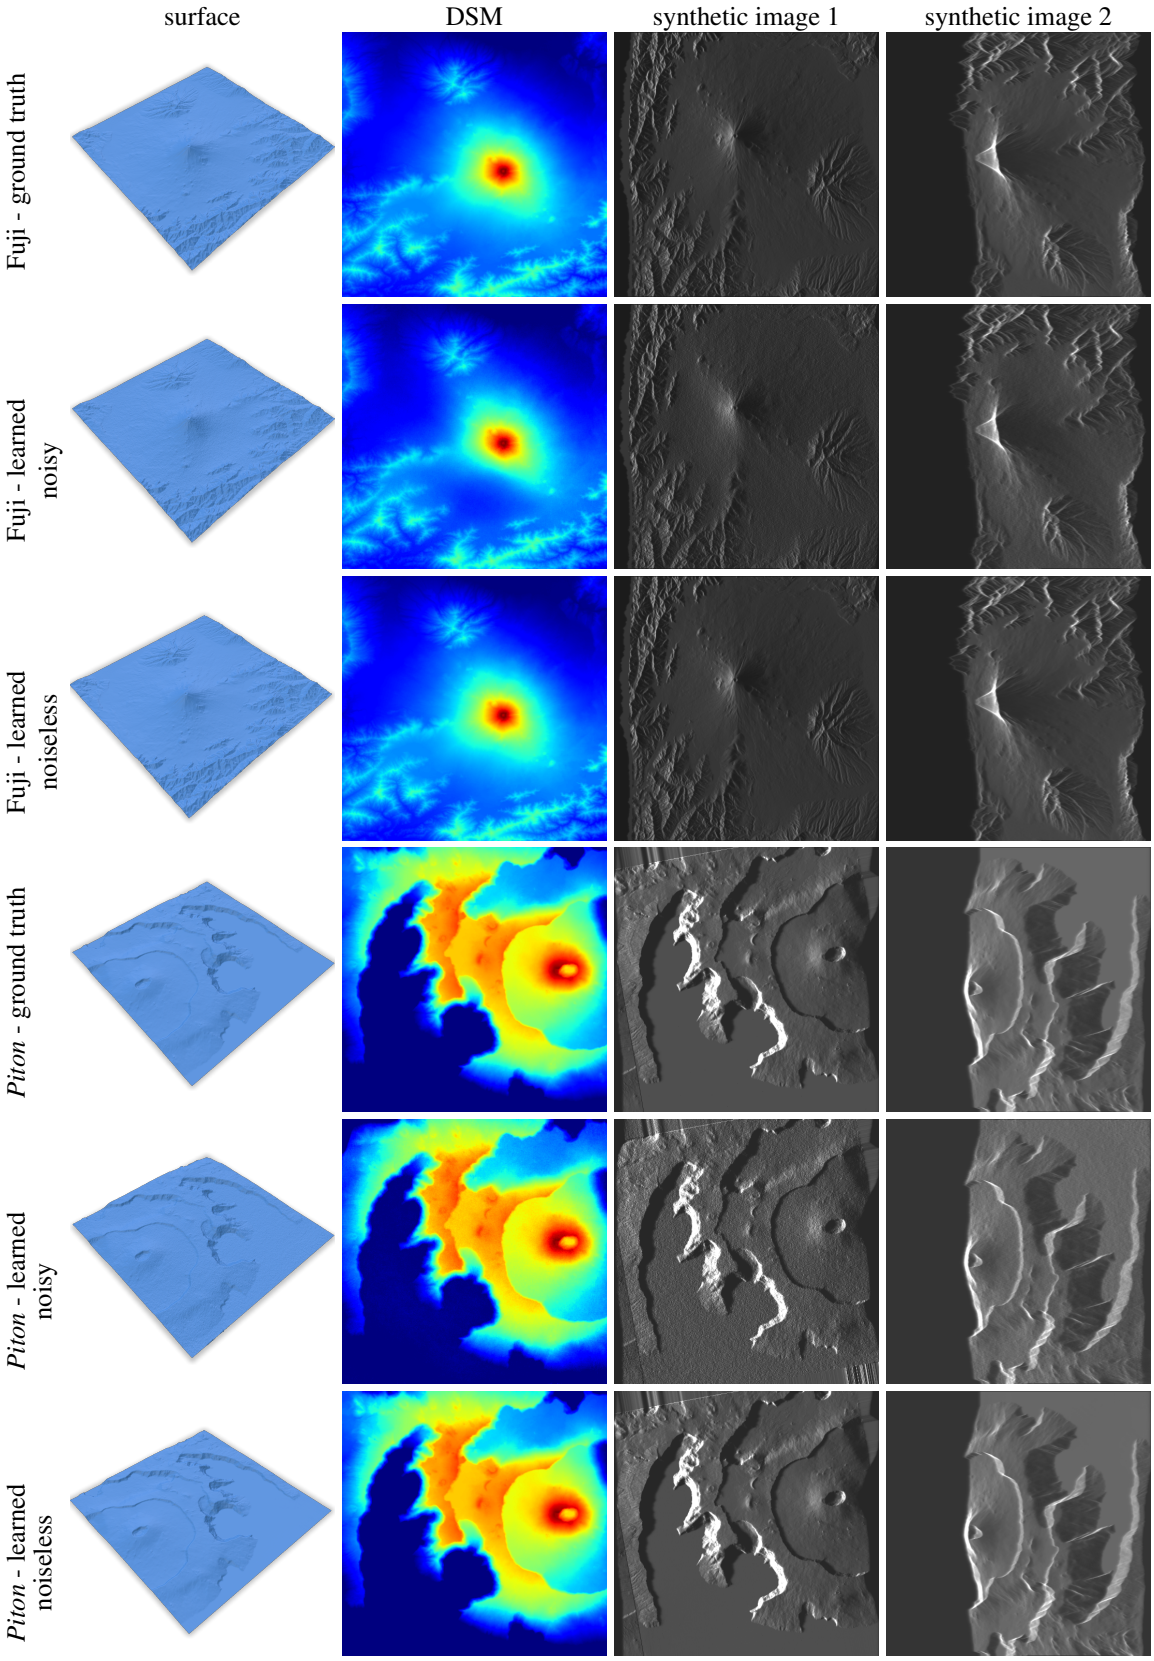

Figure 2. Results of surface learning using five SAR images on simulated data from two real DSMs. The first represents mount Fuji (first row is the ground truth, second row the learned model with the noisy views and third row with the noiseless views) and the second the *Piton de la Fournaise* (ground truth in the fourth row, learned model with noisy views in the fifth row and learned model with noiseless views in the last row). From left to right: visualization of the surface model, DSM corresponding to the surface model, and two images corresponding to two views used during training.

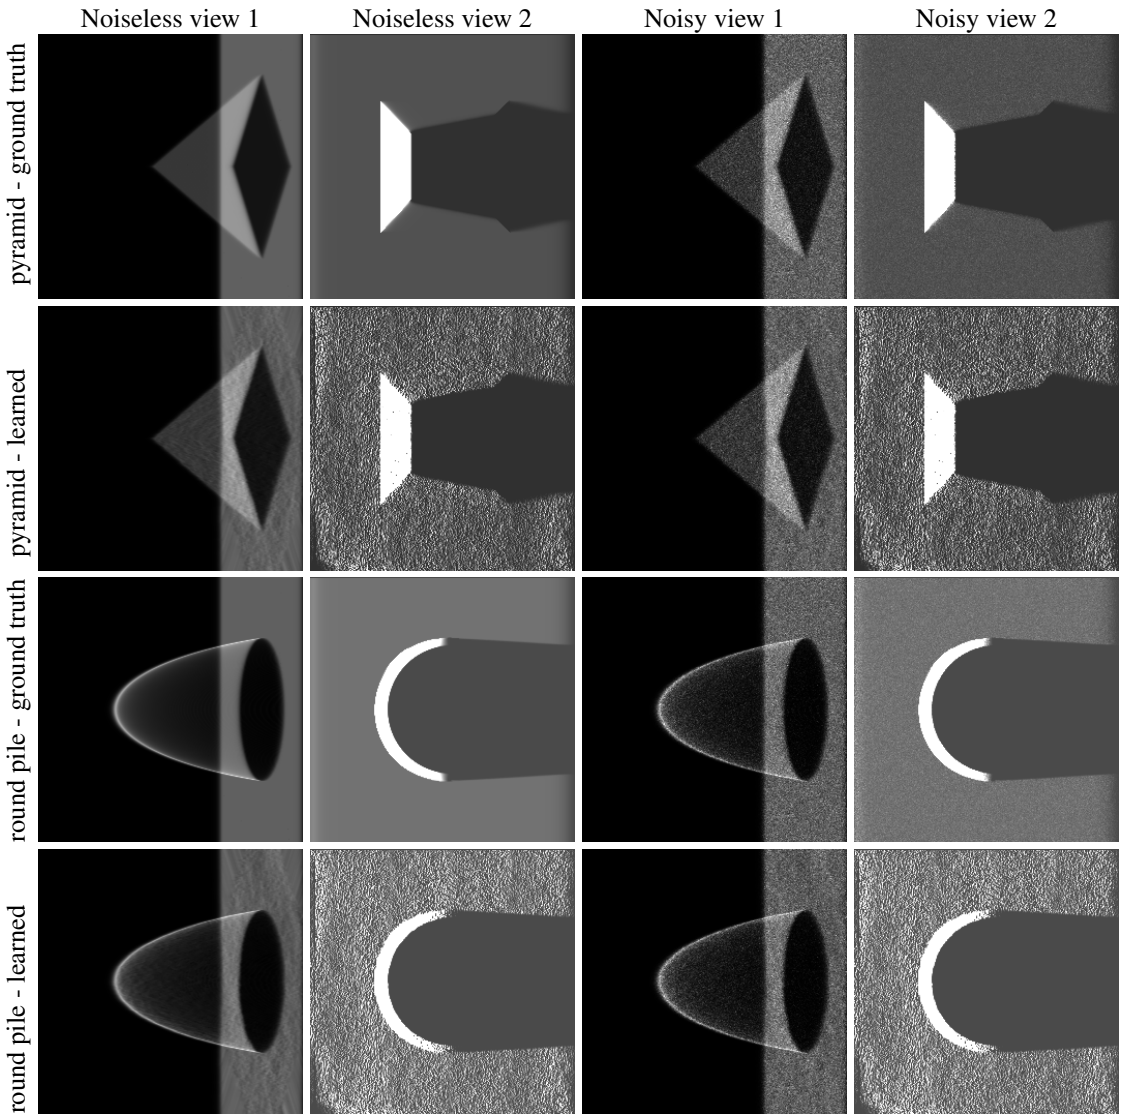

Figure 3. Additional unseen views corresponding to Figure 3 of the main paper. The first two columns corresponds to noiseless views while the last two are with noise. First row is the pyramid ground truth, second row the learned pyramid model, third row the round pile ground truth and last row the learned round pile model.

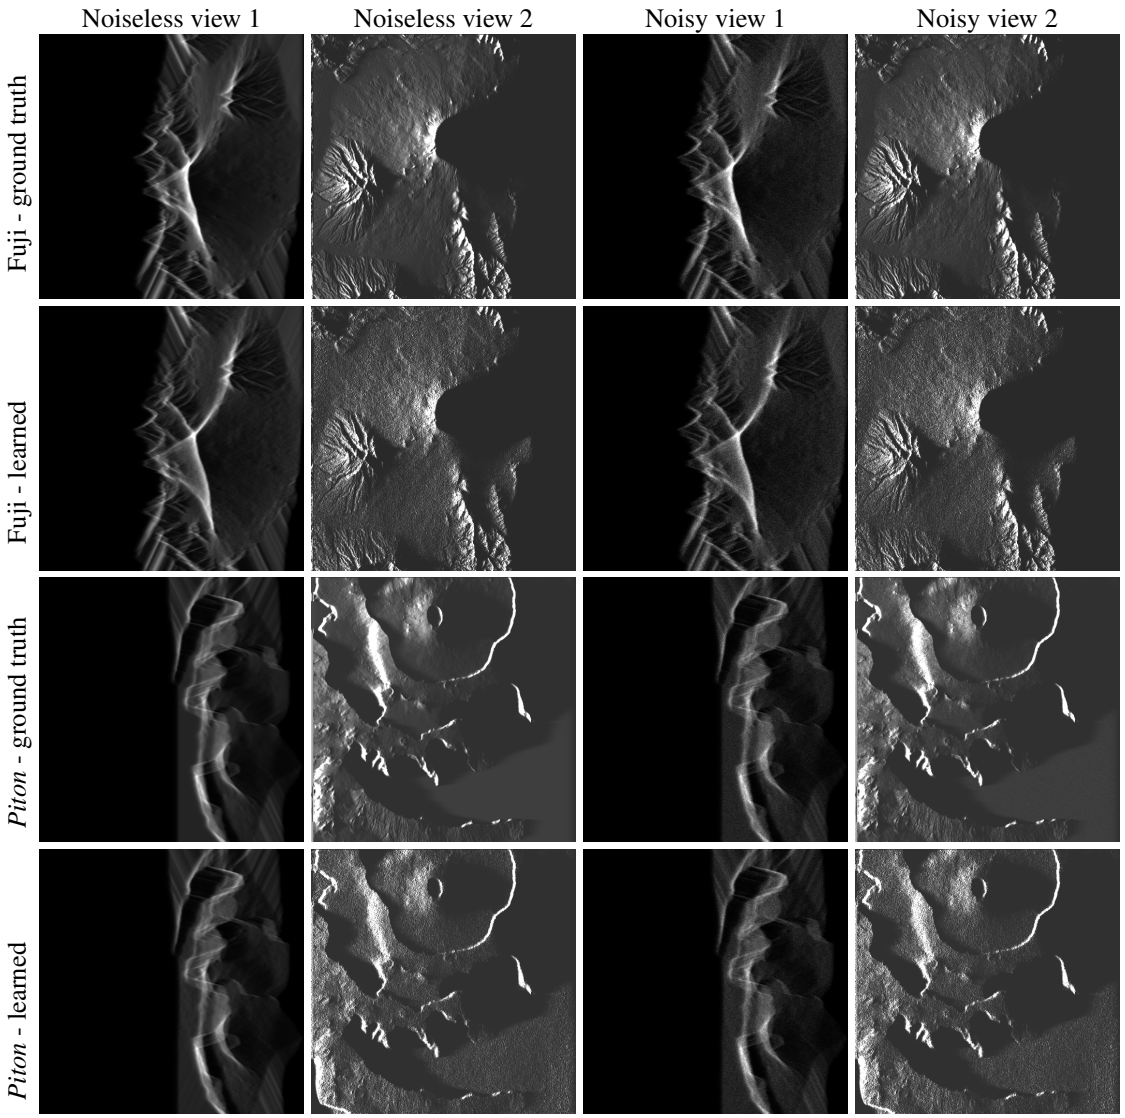

Figure 4. Additional unseen views corresponding to Figure 4 of the main paper. The first two columns corresponds to noiseless views while the last two are with noise. First row is Fuji ground truth, second row the learned Fuji model, third row the *Piton de la Fournaise* ground truth and last row the learned *Piton de la Fournaise* model.

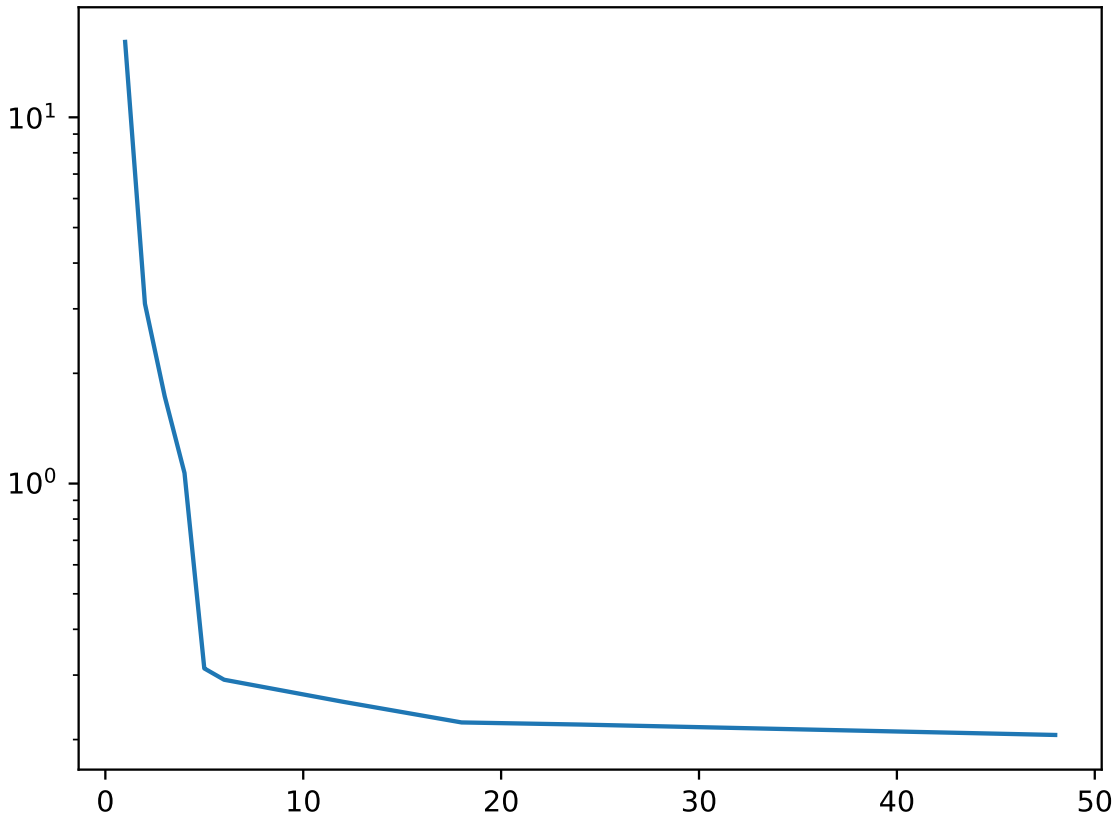

Figure 5. Mean absolute error (in m) between the ground gruth and the learned surface (in log, y axis) in function of the number of views (x axis).