Supplementary Materials:
3DRegNet: A Deep Neural Network for 3D Point Registration

G. Dias Pais[1], Srikumar Ramalingam[2], Venu Madhav Govindu[3],
Jacinto C. Nascimento[1], Rama Chellappa[4], and Pedro Miraldo[1]

[1]Instituto Superior Técnico, Lisboa [2]Google Research, NY
[3]Indian Institute of Science, Bengaluru [4]University of Maryland, College Park

In these supplementary materials, we start by showing additional figures illustrating the 3DRegNet vs. FGR, with and without ICP for refinement, (see Sec. A). In Sec. C, we discriminate the results obtained in Tab. 5 of the paper.

A. Additional Results

We show some new figures to better illustrate the advantages of the 3DRegNet against previous methods (i.e., Tab. 5 of the main document).

We start by showing additional experimental results on the 3D scan alignment to complement the results shown in Fig. 5 of the paper. Two sequences were used, MIT and BROWN, from the SUN3D dataset. **Please note that the 3DRegNet was not trained using these sequences; these are used for testing only**. These experiments are similar to the ones in Fig. 5 of the paper. However, instead of only showing a pair of 3D scans (required by each of the methods), we show the registration of 10 3D scans. We compute the 3D alignment in a pairwise manner, i.e., we compute the transformation from Scan 1 to Scan2, from Scan 2 to Scan 3, ..., and Scan 9 to Scan 10. Then, we apply transformations to move all the 3D Scans 2, 3, ..., 10 into the first one, which we selected for the reference frame. We consider the cumulative transformation from the first to i^{th} 3D scan, i.e., we pre-multiplied all the transformations from 1 to i to move all the point clouds into the first (common) reference frame. We used the methods: (i) 3DRegNet, (ii) 3DRegNet + ICP, (iii) FGR, and (iv) FGR + ICP. These results are shown in Fig. A.1. We show an additional column with the ground-truth transformation for comparison. We use the network trained for the results in Tab. 5(b) of the paper.

As we can see from Fig. A.1, for both the Brown and the MIT sequences, the registration results for the 10 scans given by the 3DRegNet method are much closer to the ground-truth than the FGR. When running the ICP after the 3DRegNet, while for the Brown, we see some improvements (compare the door in 3DRegNet vs. 3DRegNet + ICP), for the MIT we see some degradation on the results. When comparing FGR with 3DRegNet, for the Brown sequence, we see that the 3DRegNet is performing better than the FGR, even for the case in which we use ICP for the FGR refinement. For the MIT sequence, we see that, while the 3DRegNet is performing better than the FGR, the ICP for refinement after both is leading to the same final 3D registration. However, we can also observe that the 3DRegNet is giving better results than 3DRegNet + ICP and FGR + ICP (see the cabinets in the environment).

We further evaluate the use of 3dRegNet against the current state-of-the-art FGR method by showing the trajectories obtained from each of the methods. The results for 20 frames in two sequences are shown in Fig. A.2. The point clouds shown in this figure are registered using the ground-truth transformations, and the paths shown are computed directly from 3DRegNet + ICP and FGR + ICP. From the top of the Fig. A.2 (Harvard sequence), it can be seen that we are performing better than the FGR + ICP, i.e., 3DRegNet + ICP provides a trajectory estimate that is closer to the ground-truth. For the Brown dataset (bottom of Fig. A.2), we see that both trajectories perform similarly. However, we stress that the 3DRegNet is faster than the competing methods, as shown in the Tab. 5(b) of the paper.

B. Cumulative Distribution Function for SUN3D

To better illustrate the performance of 3DRegNet against FGR, the cumulative distribution function of the rotation errors was computed for the SUN3D sequences as shown in Fig. B.3. It can be seen that FGR performs better than 3DRegNet until 2.5 degrees error. Also, 3DRegNet is remarkably better when compared to the FGR + ICP, exhibiting superior performance around 4 degrees error. This implies that FGR does a better job for easier problems. However, for a larger number of cases, it has high error (also higher than that of 3DRegNet). In other words, FGR has a lower median error and higher mean error compared to

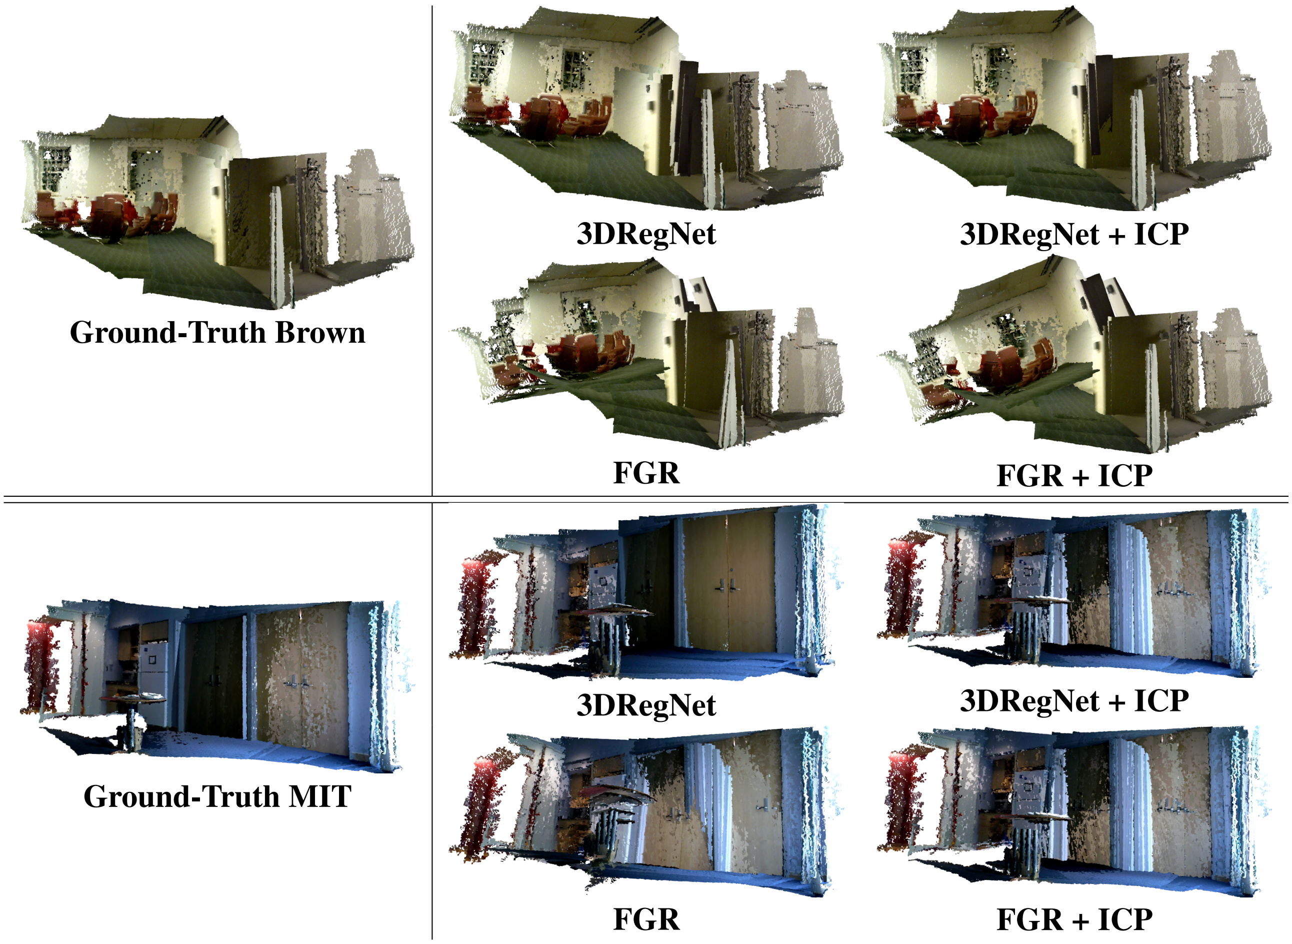

Figure A.1: *Results for the alignment of 20 3D scans using the 3DRegNet, 3DRegNet + ICP, FGR, and FGR + ICP. We consider just the transformations computed using the respective methods, i.e., we are not removing the drift from the estimation. No transformation averaging for final refinement was used.*

3DRegNet, as evident from Tab. 1. As the complexity of the problem increases, 3DRegNet + ICP becomes the best algorithm, which is confirmed by the line of FGR + ICP. At smaller degrees of errors, both lines are very similar (as confirmed in the previous image for the MIT sequence), which indicates that they converge to the same place. However, when the rotation error increases, this difference is more significant, and our method provides a much better solution to the registration problem.

C. Discriminate Results for SUN3D

Although the main paper presents the overall mean and median for all the pairs in the three sequences of the SUN3D data set, the individual errors for each of the sequence vary significantly. This is because each sequence has its own characteristics. Here we show the discriminate results for each sequence of the SUN3D sequences (see Tab. 1).

From the results, we see that while ICP is performimng better than 3DRegNet for the MIT sequence, 3DRegNet is superior in Harvard (both with and without ICP or Umeyama). In the Brown sequence, we see that while we are beating the current state-of-the-art in the mean, without refinement, we are loosing for RANSAC and FGR in the median (though the differences are minor). When considering refinement (i.e. with Umeyama or ICP), in general, our proposal is the best method. Exception is the slightly better performance in the FGR + ICP where the estimated median and the translation are superior by a small margin. Overall, when we see these results, we can draw the same conclusions as the ones addressed in the paper. While both ICP and FGR perform well for less challenging scenarios (small transformations), our method is superior for larger transformations. In addition to these conclusions, we can easily see that the 3DRegNet is significantly faster than any other method, with and without refinement[1].

[1] We stress that all the methods are being run the same conditions, only

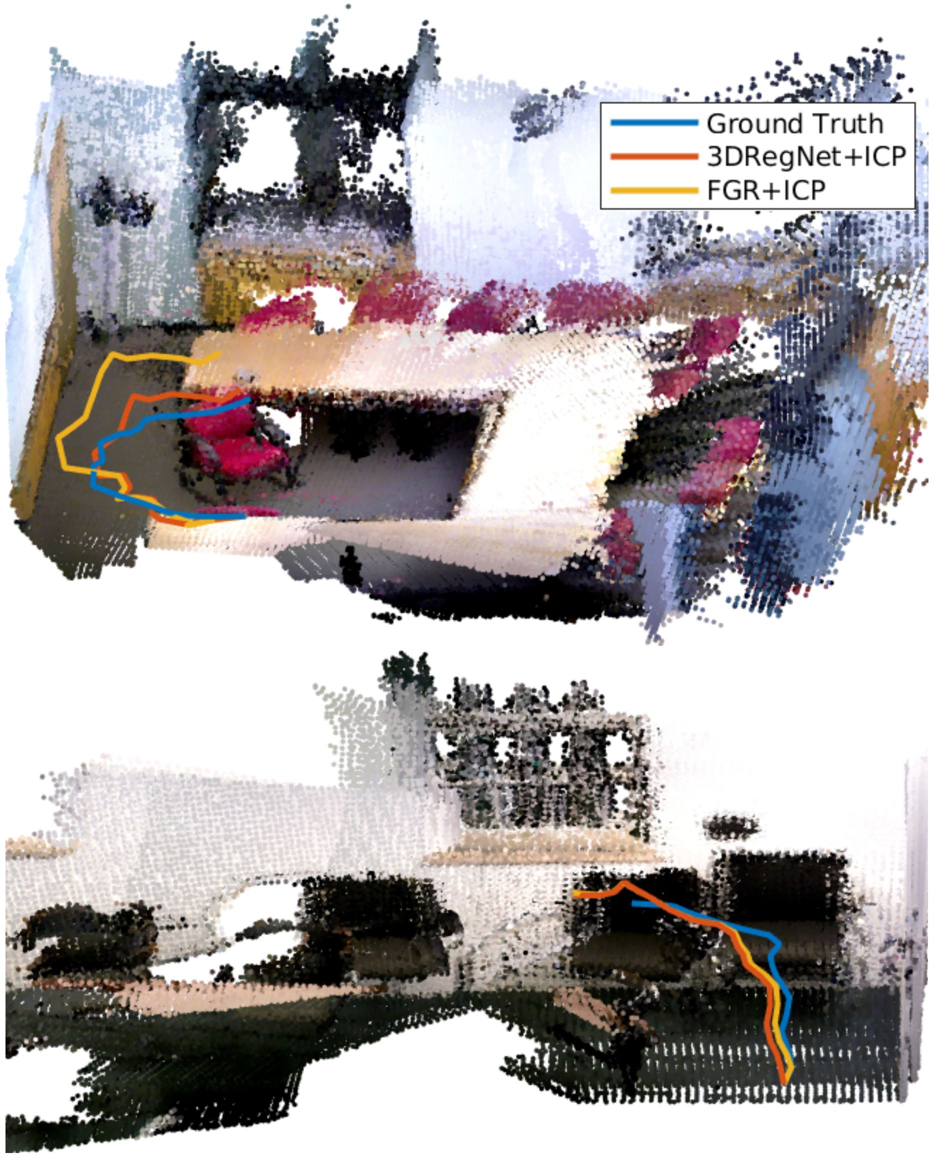

Figure A.2: *Two examples of trajectories obtained using the 3DRegNet + ICP vs. FGR + ICP against the Ground-Truth.*

References

[1] Martin A. Fischler and Robert C. Bolles. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. Commun. ACM, 24(6):381–395, 1981.

[2] Peter H. Schonemann. A generalized solution of the orthogonal procrustes problem. Psychometrika, 31(1):1–10, 1966.

[3] Qian-Yi Zhou, Jaesik Park, and Vladlen Koltun. Fast global registration. In European Conf. Computer Vision (ECCV), pages 766–782, 2016.

---

using CPU.

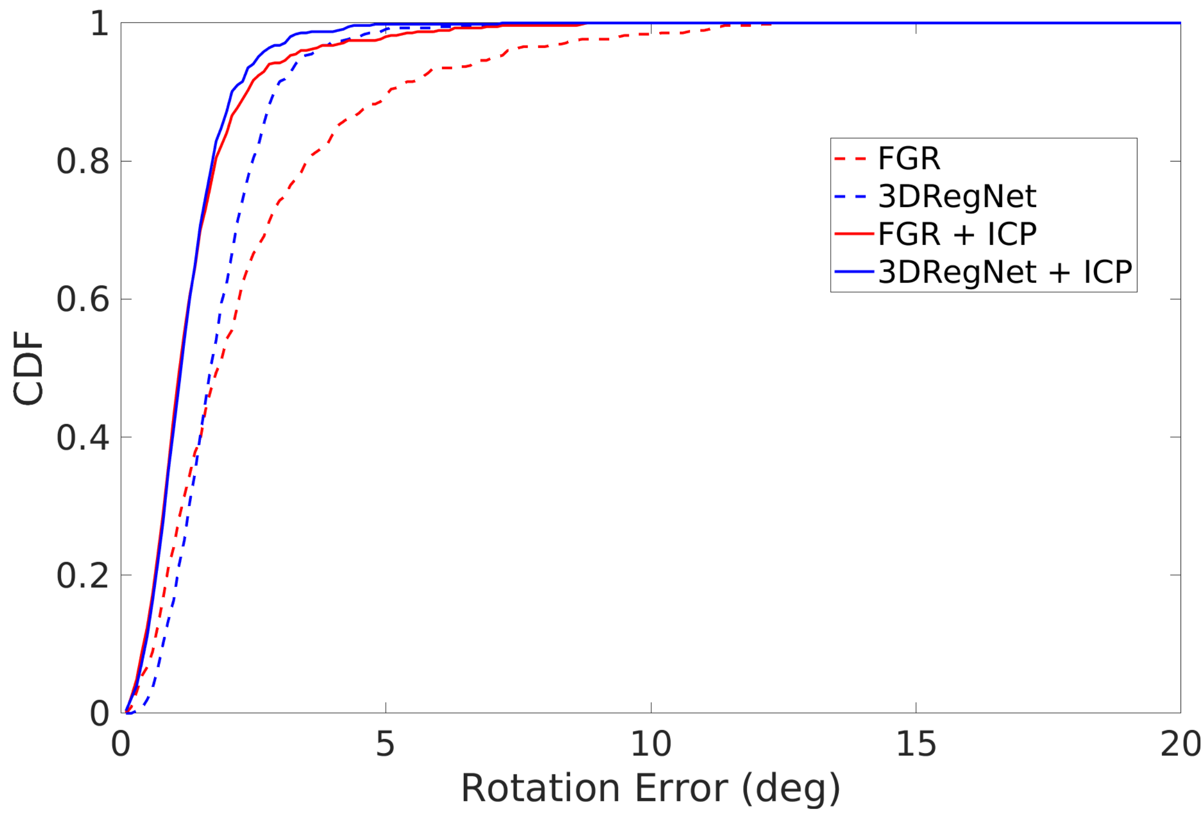

Figure B.3: *Cumulative Distribution Function (CDF) of the rotation errors on the SUN3D dataset.*

	Rotation [deg]		**Translation [m]**		
Method	**Mean**	**Median**	**Mean**	**Median**	**Time [s]**
FGR	1.96	1.58	0.083	0.055	0.16
ICP	**1.53**	**1.14**	**0.071**	**0.045**	0.086
RANSAC	1.90	1.64	0.080	0.065	2.28
3DRegNet	1.77	1.62	0.080	0.070	**0.023**
FGR + ICP	**1.01**	**0.38**	**0.038**	**0.021**	0.19
RANSAC + U	1.58	1.35	0.065	0.053	2.28
3DRegNet + ICP	1.10	1.04	0.047	0.039	0.062
3DRegNet + U	1.15	1.10	0.048	0.047	**0.023**

(a) **MIT**

	Rotation [deg]		**Translation [m]**		
Method	**Mean**	**Median**	**Mean**	**Median**	**Time [s]**
FGR	3.25	2.63	0.169	0.117	0.14
ICP	4.94	3.11	0.275	0.221	0.082
RANSAC	2.87	2.28	0.166	0.113	3.49
3DRegNet	**1.75**	**1.60**	**0.095**	**0.078**	**0.023**
FGR + ICP	1.59	1.30	0.112	0.067	0.18
RANSAC + U	2.54	1.82	0.149	0.092	3.49
3DRegNet + ICP	1.38	1.28	0.098	0.075	0.085
3DRegNet + U	**1.20**	**1.13**	**0.069**	**0.059**	**0.023**

(b) **Harvard**

	Rotation [deg]		**Translation [m]**		
Method	**Mean**	**Median**	**Mean**	**Median**	**Time [s]**
FGR	2.72	1.77	0.12	**0.060**	0.15
ICP	3.74	1.69	0.16	0.11	0.080
RANSAC	3.99	**1.66**	0.20	0.071	2.55
3DRegNet	**1.92**	1.78	**0.089**	0.082	**0.020**
FGR + ICP	1.64	1.14	0.079	**0.046**	0.19
RANSAC + U	3.77	1.48	0.182	0.059	2.55
3DRegNet + ICP	1.33	1.18	0.067	0.047	0.085
3DRegNet + U	**1.13**	**1.06**	**0.051**	0.048	**0.020**

(c) **Brown**

Table 1: *Comparison with the baselines: FGR [3]; and RANSAC-based approaches [1, 2].*