FlowNet3D++: Geometric Losses For Deep Scene Flow Estimation (Supplementary Material)

Zirui Wang
ryan@robots.ox.ac.uk

Shuda Li
shuda@robots.ox.ac.uk

Henry Howard-Jenkins
henryhj@robots.ox.ac.uk

Victor Adrian Prisacariu
victor@robots.ox.ac.uk

Min Chen
min.chen@oerc.ox.ac.uk

1. Further FlowNet3D++ Results With Our Dynamic Reconstruction Benchmark

We present more qualitative results to demonstrate that integrating FlowNet3D++ is useful for high-level dynamic reconstruction applications. We utilise our dynamic reconstruction benchmark pipeline to make qualitative comparisons between dynamic reconstruction results with FlowNet3D++ integrated and without FlowNet3D++. For dynamic reconstruction baseline, we implement SobolevFusion[4], which is an improved version of KillingFusion[3].

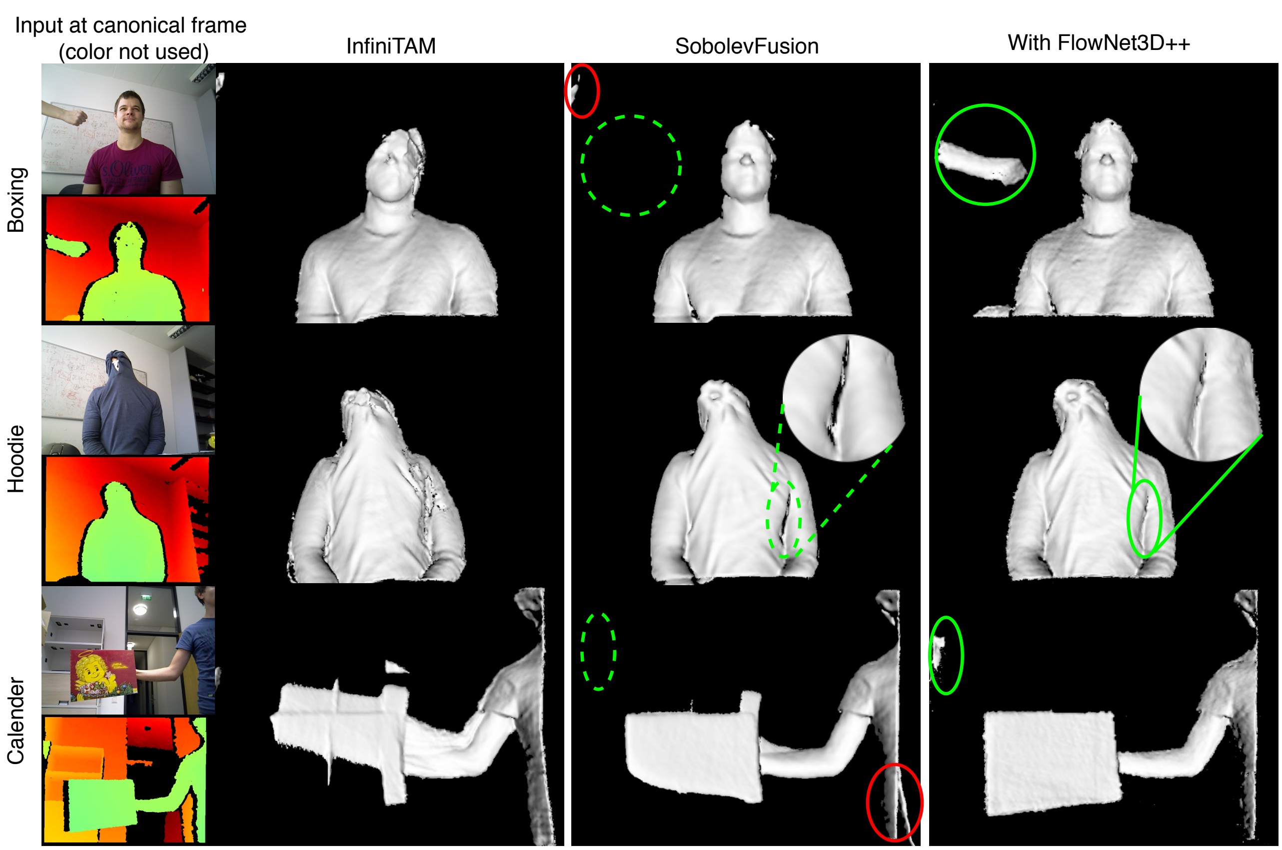

Figure 1. Results of running our benchmark on VolumeDeform [1] Dataset. Aside from obvious differences, we use red circles to highlight places where integrating FlowNet3D++ produces less distortion and green circles to highlight places integrating FlowNet3D++ preserves more detail. Note there is no foreground mask, the background is missing because of volume size. For scenes with more rigid components (the hoodie sequence), we have similar results with SobolevFusion because the rigid tracker lifts the most weight.

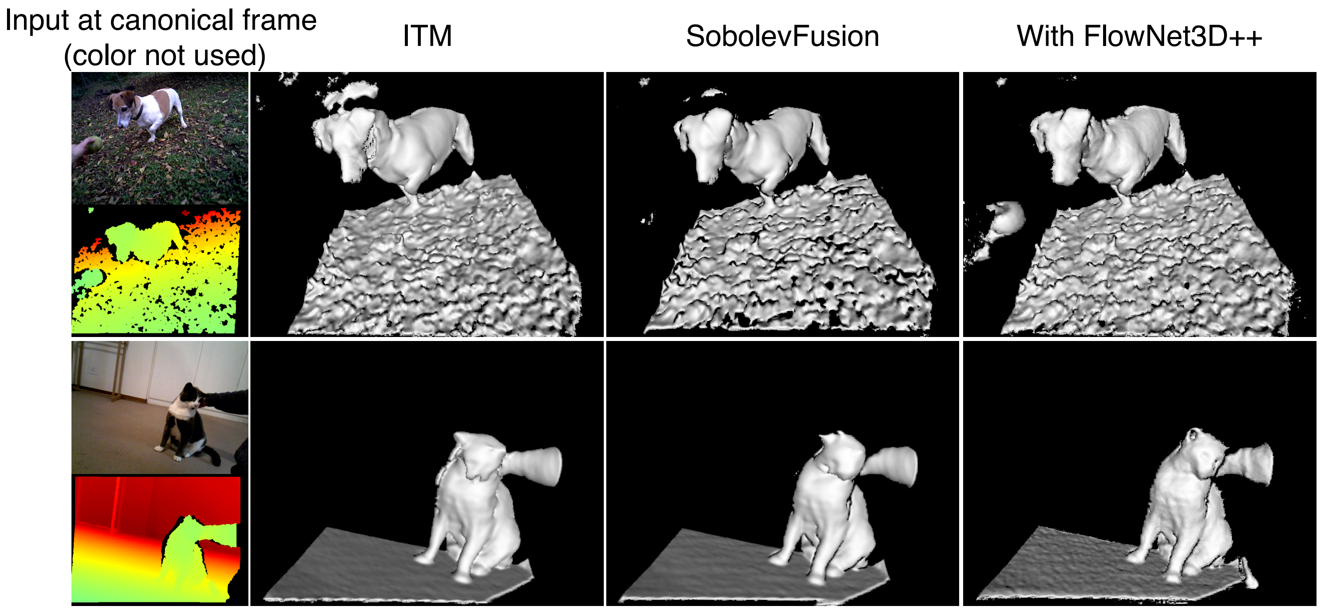

Figure 2. Results of running our benchmark on a dataset used in [2]. Our results have fewer artefacts and more details than SobolevFusion after few hundreds of frames.

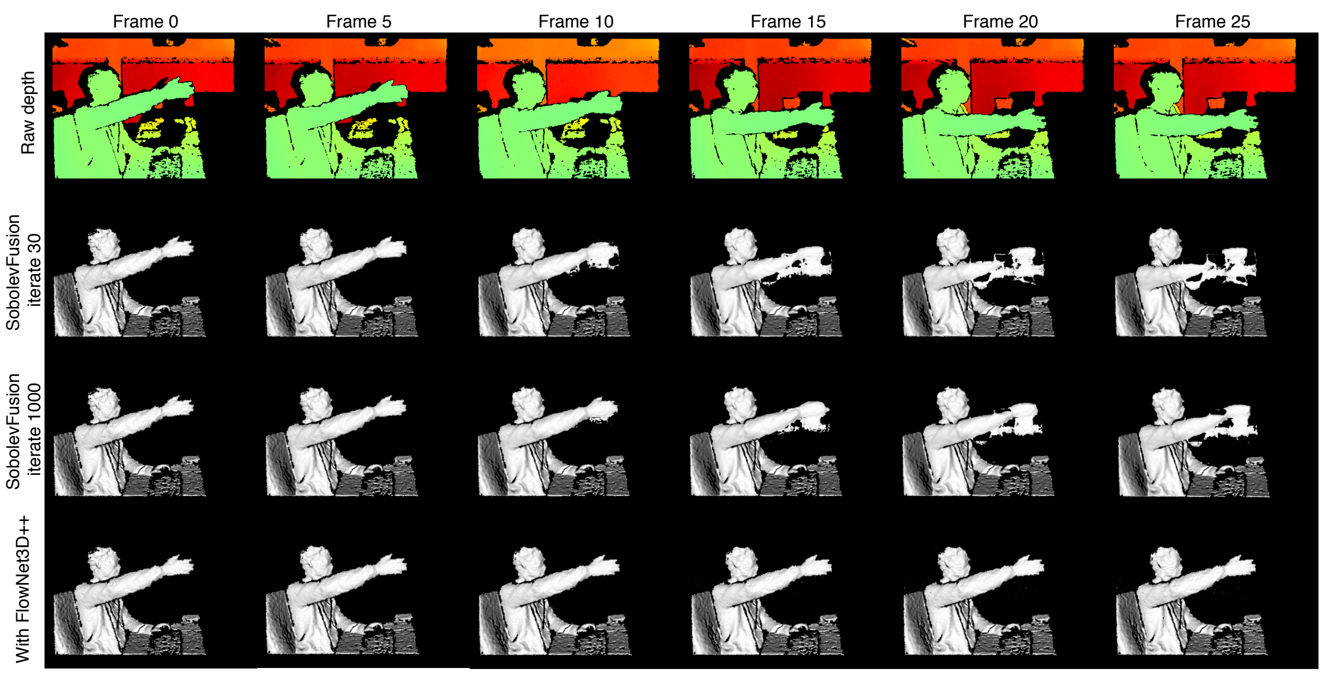

Figure 3. A comparison of the different number of iterations on a video we recorded. We show the canonical model at different time stamps. From the figure we can see that when the iterations increase, the SobolevFusion results improve. However, it is still difficult to register the live volume to the canonical model. In contrast, integrating FlowNet3D++ enables the handling of large motion.

2. FlowNet3D++ Non-Rigid Registration

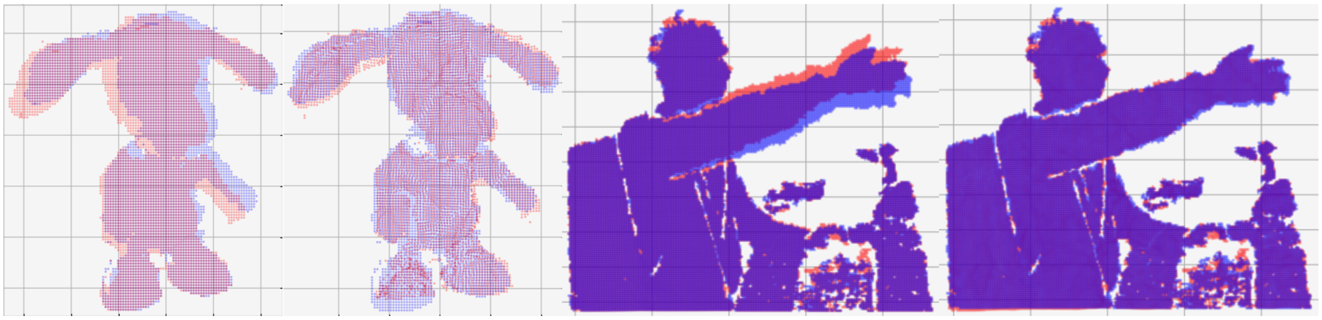

Figure 4. Point clouds warped using the deep scene flow. The left and right pairs are the results using the *Snoopy* sequence and *Waving* sequence, respectively. It can be seen that FlowNet3D++ has aligned the source point cloud (blue) with the target point cloud (red) effectively. Note that although the colour scheme for each of the pairs appears different, this is simply caused by the differing densities of their respective point clouds.

References

[1] M. Innmann, M. Zollhöfer, M. Nießner, C. Theobalt, and M. Stamminger. Volumedeform: Real-time volumetric non-rigid reconstruction. In *ECCV*, 2016. 1

[2] C. Malleson, J.-Y. Guillemaut, and A. Hilton. Hybrid modelling of non-rigid scenes from rgbd cameras. *IEEE Transactions on Circuits and Systems for Video Technology*, 2018. 2

[3] M. Slavcheva, M. Baust, D. Cremers, and S. Ilic. Killingfusion: Non-rigid 3d reconstruction without correspondences. In *CVPR*, 2017. 1

[4] M. Slavcheva, M. Baust, and S. Ilic. Sobolevfusion: 3d reconstruction of scenes undergoing free non-rigid motion. In *CVPR*, 2018. 1