A Symmetric Objective Function for ICP

SZYMON RUSINKIEWICZ, Princeton University, USA

SUPPLEMENTAL MATERIAL

This supplemental document contains one additional per-iteration convergence result, to complement Figure 4 of the original paper. In particular, it illustrates a "best-case scenario" of a smooth shape: an ellipsoid. In this case, we see the behavior of the different error metrics with no or minimal influence from outlier rejection or grossly incorrect correspondences.

The point-to-point metric shows linear convergence, as it does in other cases. Point-to-plane displays essentially quadratic convergence, though it actually diverges for relatively large initial misalignment. The quadratic method avoids this poor behavior, but for smaller misalignment does no better on average than point-to-plane. The two-plane method does better in this case: it avoids blowup at large misalignment, while transitioning from quadratic to cubic convergence as misalignment decreases.

The symmetric metric, and its rotated-normals variant, both show consistent cubic convergence, with a per-iteration constant that is a factor of 10-20 smaller than the two-plane variant. This is perhaps not surprising, given that the symmetric metric is exactly minimized for second-order surfaces, and so its performance is dictated by third-order deviations.

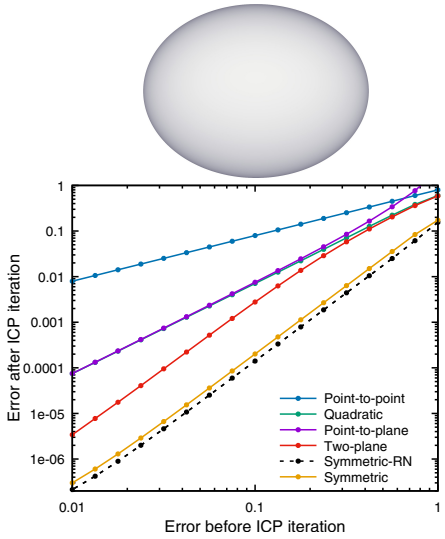

Fig. 4supp. Error decrease due to one ICP iteration on the `ellipsoid` model, aligned to itself. Ground-truth errors before and after the ICP iteration are shown on the x and y axes, respectively, of this log-log plot.