Point Pair Feature Matching: Evaluating Methods to Detect Simple Shapes

Markus Ziegler, Martin Rudorfer, Xaver Kroischke, Sebastian Krone, and Jörg Krüger

Technische Universität Berlin, PTZ 5, Pascalstr. 8-9, 10587 Berlin

contact: m.f.z@t-online.de, martin.rudorfer@iat.tu-berlin.de

ICVS THESSALONIKI, 2019

Introduction & Background

Point Pair Feature Matching (PPFM) has first been proposed by (Drost et al., 2010) for 6d object detection in point clouds. A recent benchmark, the SIXD Challenge 2017 (Hodan et al., 2018), shows that PPFM is still one of the most relevant methods. The method relies on Point Pair Features (PPF), which describe simple geometric relations between a reference point p_r and a target point p_t (see Figure). A description for the 3d object model is computed by using all permutations of PPF, which are then organized in a hash table. The hash key is the discretized feature itself and the value efficiently encodes the feature's pose with respect to the model. During matching, scene PPFs are computed and matched to the corresponding model PPFs. Using the model's hash table, each feature can vote for plausible poses in a Hough-like manner. There is one voting scheme for each scene reference point, and the resulting pose hypotheses are clustered to identify the hypothesis with highest weight.

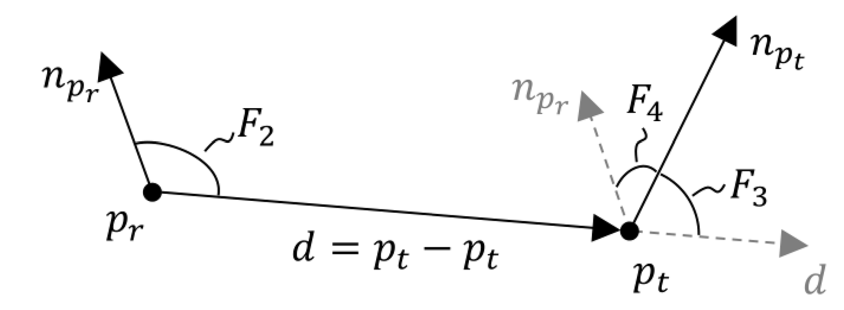

$$F(p_r, p_t) = (\|d\|, \angle(n_{p_r}, d), \angle(n_{p_t}, d), \angle(n_{p_r}, n_{p_t}))$$

Problems & Methods for Detecting Simple Shapes

In manufacturing, objects often have simple shapes, which exhibit many planar or constant-curvature surfaces and rotational symmetries. For such objects there will be many different point pairs which yield the same feature. In consequence, this feature suggests many different object poses, rendering the feature indistinctive. Various countermeasures have been proposed:

- use more discriminative features (Kim et al., 2011; Tuzel et al., 2014)
- incorporate edge or color information (Choi et al., 2012; Drost et al. 2012)
- identify and handle an axis of symmetry (de Figuereido et al., 2015)
- reduce the number of votes a feature can cast (Hinterstoisser et al., 2016).

Further topics of improvement are the sensitivity to sensor noise, particularly for the estimation of the surface normals, the computational complexity of the method, and the verification of resulting pose hypotheses.

Methodology

We implement several extensions to improve the detection of simple shapes. The two most important approaches are:

- Detecting an arbitrary number of discrete rotational symmetries by matching the object to itself
- Using a more discriminative feature which adds the four values $v_{out,1\ldots4}$ incorporating information about the features location within the model (by casting four rays originating from the target point and recording the distance required to either leave the surface or hit the nearest intersection).

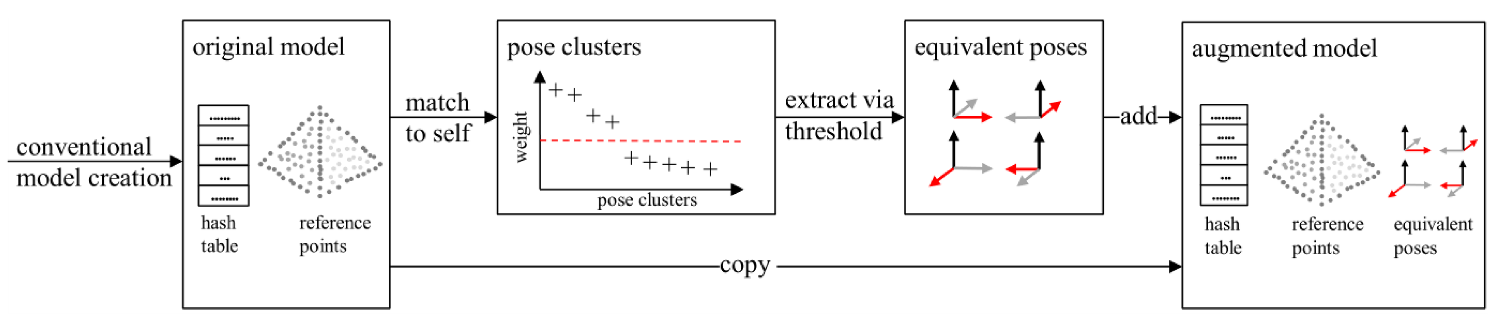

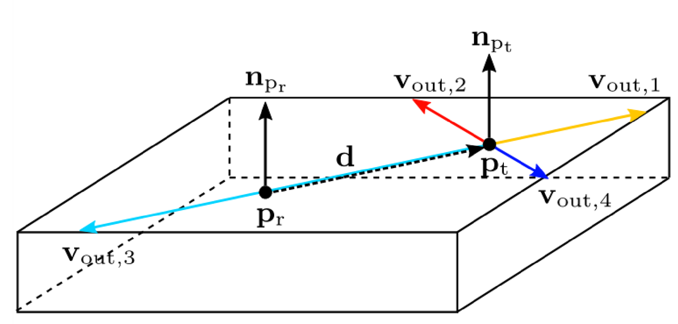

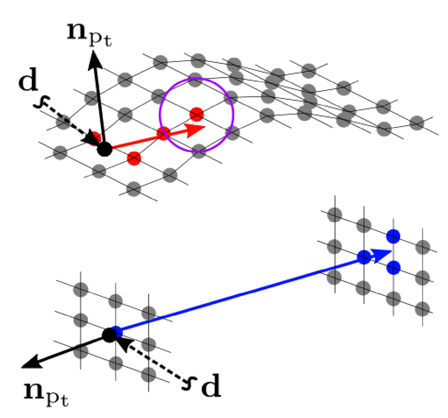

Results

While both of the main approaches improve recognition rates, the rotational symmetry handling also reduces computation time. In contrast, the augmented feature increases computational complexity. Tests have been conducted on our geometric primitive dataset. Results for our other extensions and some typical failure cases are shown in the following figures.

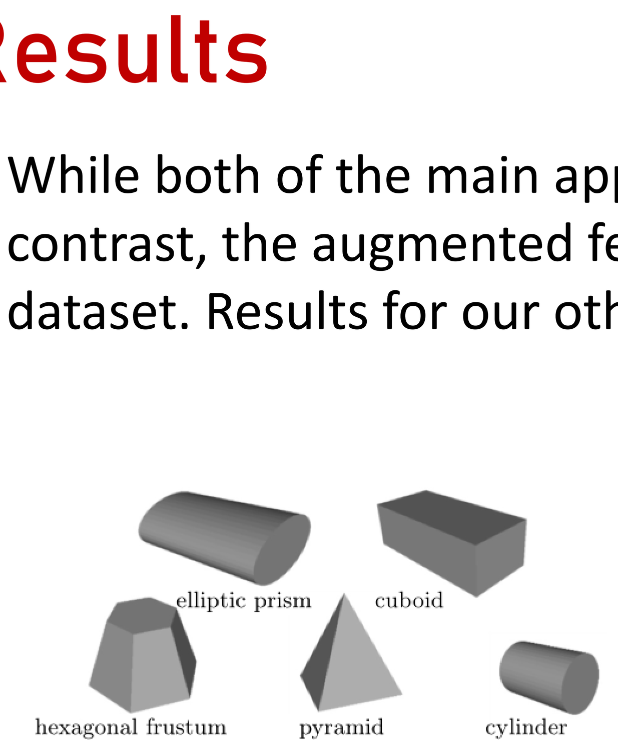

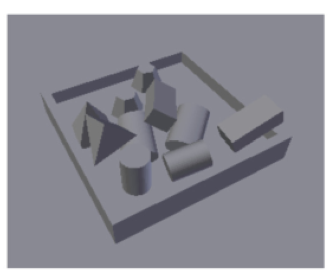

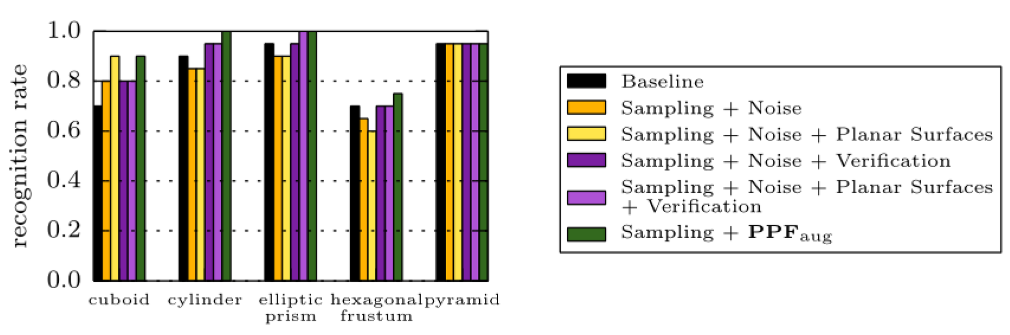

Baseline: Drost-PPFM + angle refinement + RS cluster&hash.
Sampling: Uses voting spheres.
Noise: Storing the adjacent PPFs in the model and voting for ajdacent pose rotation angels.
Planar Surfaces: Flag array and bottom-up clustering.
Verification: Postprocessing pose clusters with verifier.
$\mathbf{PPF}_{aug}$: Augmented feature with model-specific configuration.