$\mathcal{W}$-PoseNet: Dense Correspondence Regularized Pixel Pair Pose Regression

Zelin Xu*, Ke Chen*, and Kui Jia

School of Electronic and Information Engineering
South China University of Technology
Guangzhou 510641, P.R. China
eexuzelin@mail.scut.edu.cn, {chenk,kuijia}@scut.edu.cn

Abstract

Solving 6D pose estimation is non-trivial to cope with intrinsic appearance and shape variation and severe inter-object occlusion, and is made more challenging in light of extrinsic large illumination changes and low quality of the acquired data under an uncontrolled environment. This paper introduces a novel pose estimation algorithm $\mathcal{W}$-PoseNet, which densely regresses from input data to 6D pose and also 3D coordinates in model space. In other words, local features learned for pose regression in our deep network are regularized by explicitly learning pixel-wise correspondence mapping onto 3D pose-sensitive coordinates as an auxiliary task. Moreover, a sparse pair combination of pixel-wise features and soft voting on pixel-pair pose predictions are designed to improve robustness to inconsistent and sparse local features. Experiment results on the popular YCB-Video and LineMOD benchmarks show that the proposed $\mathcal{W}$-PoseNet consistently achieves superior performance to the state-of-the-art algorithms.

1. Introduction

The problem of 6D pose estimation aims to predict a rotation and translation of an object instance in 3D space relative to a canonical CAD model, which plays a vital role in a number of applications such as augmented reality [20, 46], grasp and manipulation in robotics [37, 36, 47], and 3D semantic analysis [44, 34, 13]. Such a problem remains challenging in view of intrinsic inconsistent texture and shape of objects and inter-object occlusion in the cluttered scenes, in addition to extrinsic varying illumination and sensor noises. In light of this, obtaining a discriminative and robust feature representation for each object instance is essential for predicting its 6 degree-of-freedom pose.

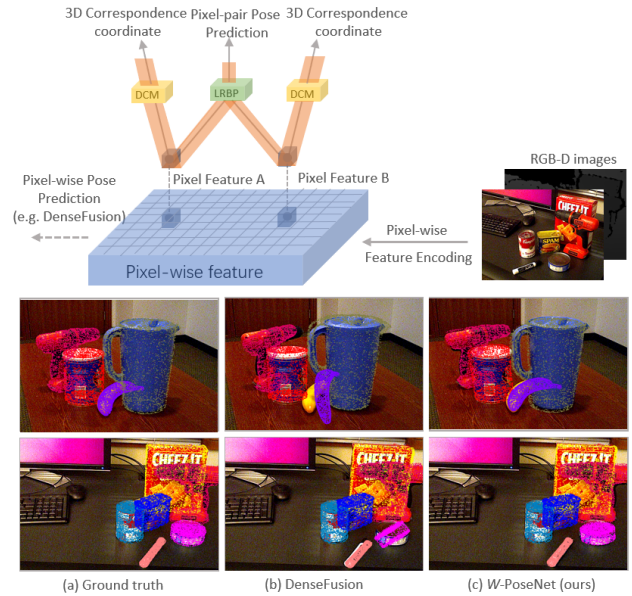*These authors contributed equally to this work.

Figure 1. Visualization of our concept of the proposed $\mathcal{W}$-PoseNet in comparison with its backbone DenseFusion [39]. The key differences lie in that our $\mathcal{W}$-PoseNet introduces pixel pair pose estimation (highlighted in green block) and dense correspondence mapping from each pixel to its 3D coordinate (highlighted in yellow blocks), which geometric structure is similar to the font $\mathcal{W}$. Illustrative examples at the bottom two rows are selected from the YCB-Video benchmark [43].

Owing to recent success of deep learning in visual recognition [31, 16], data-driven deep methods are introduced for pose estimation on RGB images via either direct regression on input images to 6D pose [38, 42, 33, 22] or exploiting 2D-3D correspondences discovered by detecting 2D projection in RGB images of sparse 3D keypoints [23, 24, 25, 30, 26]. Compared to those pose estimation methods relying on object texture extracted from RGB images, RGB-D data can provide extra depth information to mitigate the suffering of lack of textural appearance of

an object or fail to extract texture information given low-quality data.

Most of existing RGB-D methods [43, 33, 12] rely on iterative-closest point (ICP [1]) to refine pose predictions, which leads to less efficient inference and thus can be less favorable for real-time applications such as robotic grasp and planning. DenseFusion [39] has recently been developed to combine both textural and geometric features in a pixel-wise fusion manner, which can obtain accurate pose estimation in a real-time inference processing. Cheng *et al*. [4] further exploit both intra-modality and inter-modality correlation to learn more discriminative local feature based on DenseFusion [39]. However, these state-of-the-art dense regression-based pose estimation algorithms heavily depend on the quality and resolution of the acquired data. Different from well-established feature learning in the 2D domain [38, 42, 33, 22], learning discriminative feature representations from RGB-D data, containing both texture and geometries of object models, remains open to cope with heavy occlusion, varying texture and shape changes of objects and cross-modality misalignment.

We consider that *extracting discriminative dense features in local regions* and *generating a robust global representation of each object instance* are important to alleviate the aforementioned challenges in 6D pose estimation. To this end, based on pixel-wise regression methods such as DenseFusion [39] or Correlation Fusion [4], this paper is the first attempt to explore pixel-pair pose regression – $\mathcal{W}$-PoseNet by developing two key components of deep networks: a joint loss function and a pixel-pair feature encoding layer, which are shown in Figure 1.

On one hand, different from the loss function in other pixel-wise pose regressors [39, 4] on minimizing pose predictions and their corresponding ground truth, the proposed $\mathcal{W}$-PoseNet additionally encodes the textural and geometric information favoring for reconstructing pose-sensitive point sets, whose points' 3D coordinates are generated by transforming transforming the observed point cloud sampling from object models with the inverse of ground truth pose (see Sec. 3.3). In details, our $\mathcal{W}$-PoseNet incorporates an auxiliary task of dense correspondence mapping from each pixel of input data to its corresponding pose-sensitive coordinate in 3D space, in order to regularize local feature learning for pose regression by extra pixel-wise supervision self-generated from object models (*e.g.* CAD models or their mesh representation) and ground truth pose. Intuitively, compared to sparse keypoints in keypoint-based methods [23, 24, 25, 30, 26], dense correspondence mapping in our scheme treats each pixel as a keypoint to regress its corresponding 3D coordinates in object model space which makes each pixel-wise feature more discriminative, and thus is more robust to occlusion.

On the other hand, inspired by the point-pair feature [6, 10] on point clouds in 3D domain, this paper proposes a novel pixel pair feature encoding layer based on low-rank bilinear pooling [15, 40] to generate pixel-pair features from two sampled pixels' vectors in feature map output of the encoder, which are then mapped onto 6-DoF pose. Similar to the concept in [6, 10], our pixel-pair features describe the relative geometric structure of two sampled 3D points on the object models, but also contain texture information of their 2D projection on input RGB images. Global representation for each instance consists of all pixel-pairs' features, each of which outputs a pose prediction and its corresponding confidence in a pairwise manner. Soft voting on top confidential pose candidates is employed to generate final pose predictions, which can achieve stable and robust performance even given sparsely sampled pixel-pair features.

The main contributions of this paper are three-fold.

- This paper proposes a novel pixel-pair pose regression for 6d pose estimation, based on low-rank bilinear feature pooling. Global representation can achieve robust performance with sparse pixel-pair features.
- This paper also designs an auxiliary task of dense correspondence mapping from input data to 3D coordinates sensitive to 6D pose, which can regularize feature encoding in deep pose regression to improve pixel-wise local feature discrimination.
- Extensive experiments are conducted on the popular YCB-Video and LineMOD benchmarks, whose results can demonstrate that the proposed $\mathcal{W}$-PoseNet consistently outperforms the state-of-the-art algorithms.

Source codes and pre-trained models will be released at `https://github.com/xzlscut/W-PoseNet`.

2. Related Works

Keypoint-based 6D Pose Estimation – The algorithms belonging to keypoint-based pose estimation share a two-stage pipeline: first localizing 2D projection of predefined 3D keypoints and then generate pose predictions via 2D-to-3D keypoint correspondence mapping with a PnP [7]. Existing methods can be categorized into two groups: object detection based [30, 35] and dense heatmap based [24, 25]. The former focuses on solving the problem of sparse keypoint localization via object detection on the whole object to reduce negative effects of background [30, 35], but are sensitive to occlusion [24]. The latter group of methods [24, 25] pay more attention to discovering latent correlation across all the keypoints and thus are more robust to inter-object occlusion. Recently, PVNet [26] is proposed to detect keypoints via voting on pixel-wise predictions of the directional vector that points to keypoints and is robust to truncation and occlusion. In contrast, dense correspondence mapping in our method share similar concepts as
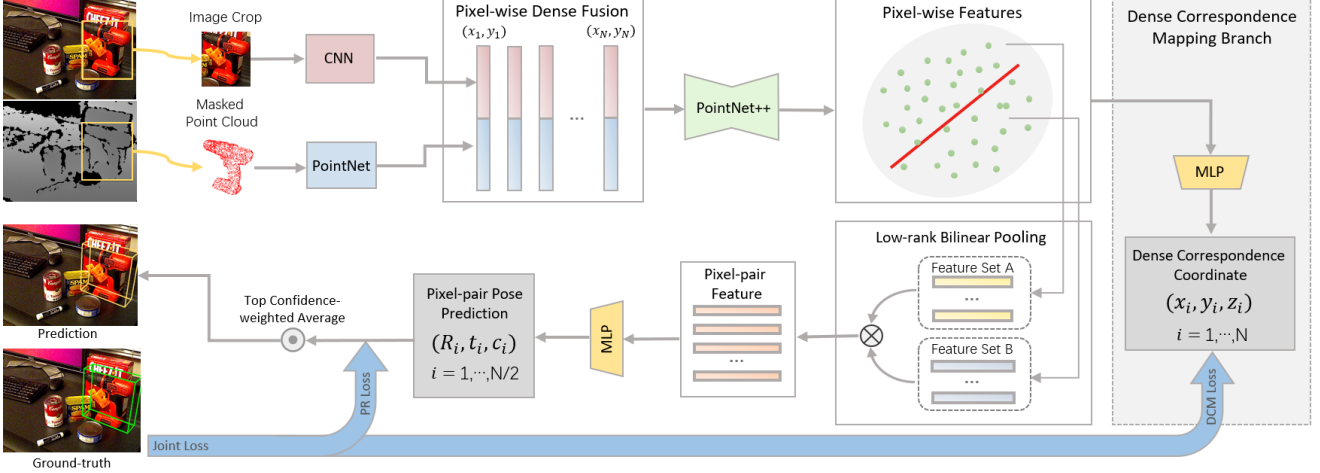
Figure 2. Pipeline of the proposed $\mathcal{W}$-PoseNet. The method first detects and segments the foreground containing object instances on RGB images. RGB and depth images are respectively fed into feature encoders and then fused with the PointNet++. Pixel-wise features are sparsely sampled and combined to generate pixel pair features, which produces 6D pose $[\boldsymbol{R}|\boldsymbol{t}]$ and confidence c for each pixel pair. The branch about Dense Correspondence Mapping is to regress pose-specific 3D coordinates from per-pixel features, providing addition geometric constraints for feature learning. A joint loss on dense correspondence mapping and pixel-pair pose regression are used to supervise network training.

keypoint-based methods but makes pixel-wise predictions on 3D keypoints directly instead of their projection in 2D images, which performs robustly to handle with occlusion.

Dense 6D Pose Estimation – An alternative group of algorithms is to produce dense pose predictions for each pixel or local patches with hand-crafted features [18, 32], CNN patch-based feature encoding [5, 13] and CNN pixel-based feature encoding [39, 4], whose final pose output is selected via a voting scheme. In [2, 21], random forests are adopted to regress 3d coordinates for each pixel and discover 2D-3D correspondence to produce pose predictions. The proposed $\mathcal{W}$-PoseNet is designed to combine both dense pose predictions and 3d coordinates in a unique framework. In general, our method follows the same pipeline – dense pose predictions and voting as [39, 4], but the key difference lies in incorporating an extra branch to constrain feature encoding with dense correspondence mapping from input RGB-D images to its corresponding 3D coordinates in the model coordinate system.

Bilinear Pooling in CNNs – The bilinear pooling [19] is an effective tool to calculate second-order statistics of local features in visual recognition but suffers from expensive computational cost due to its high-dimensionality. A number of its variants have been proposed to cope with the challenge via approximation with Random Maclaurin [11] or Tensor Sketch [27]. A low-rank bilinear model is proposed by Kong *et al.* [15] to avoid computing the bilinear feature matrix, which is approximated via two decomposed low-rank matrices. Similarly, [41] introduces Grassmann Pooling based on SVD decomposition to approximate feature maps. A quadratic transformation with the low-rank

constraint is exploited on pairwise feature interaction either from spatial locations [17] or from feature maps across network layers [45]. Our pixel-pair pose regression shares similar concepts as bilinear models to construct pair-wise features on each pixel, but a very sparsely sampled set of pixel-pair features in $\mathcal{W}$-PoseNet can achieve robust estimation, owing to pair-wise combination on local features.

3. Methodology

The problem of 6D pose estimation given RGB-D images is to detect object instances in the scenes and estimate their rotation $\boldsymbol{R} \in SO(3)$ and translation $\boldsymbol{t} \in \mathbb{R}^3$. In detail, 6D pose can be defined as a rigid transformation $\boldsymbol{p} = [\boldsymbol{R}|\boldsymbol{t}]$ from the object coordinate system with respect to the camera coordinate system. This paper aims to improve discrimination and robustness of global representation of object instances having large variations of texture and shape, in the perspective of pairwise feature interaction on dense local features.

Figure 2 illustrates the whole pipeline of our $\mathcal{W}$-PoseNet, which consists of several main stages in addition to iterative pose refinement as DenseFusion [39] (see Sec. 3.4). Encouraging by the recent success of dense pose regression on RGB-D data, we employ the same feature encoding and fusion part of DenseFusion [39] to extract and fuse pixel-wise features from heterogeneous data (see Sec. 3.1). Pixel-wise features are sampled to generate a sparse set of pixel-pair features via bilinear pooling on local features, which produces pose candidates and their confidences for voting (see Sec. 3.2). For robust pixel-wise features, an auxiliary task of learning a dense correspondence mapping

3

from each pixel to 3D coordinate together with pixel-pair pose regression is designed in Sec. 3.3.

3.1. Semantic Segmentation and Feature Encoding

We adopt the identical modules of DenseFusion [39] and briefly introduce the following steps: 1) semantic segmentation on RGB images and point clouds converted by depth images; 2) dense feature extraction and 3) pixel-wise feature fusion, which is shown in the top row of Figure 2.

Following the segmentation algorithm adopted in [39, 43], an RGB image is first fed into the auto-encoder-based segmentation network to produce $N + 1$ binary masks belonging to N object classes and the background class respectively. The bounding box of each mask is used to crop the corresponding depth image, and the depth value in the object mask is converted to a point cloud.

Cropped image patches and point clouds are fed into 2D CNN based and PointNet [28] based feature encoder to extract texture and geometric features from heterogeneous data sources. Specifically, given $H \times W$ RGB-D images, the texture branch aims to mapping $H \times W \times 3$ images to $H \times W \times d_{\text{rgb}}$-dimensional feature maps, which correspond to a feature vector $\boldsymbol{f} \in \mathbb{R}^{d_{\text{rgb}}}$ at each spatially localized pixel. The other branch extracts geometric features from orderless points via two shared-MLP and produces $P \times d_{\text{depth}}$-dimensional feature maps where P denotes the number of sample points in the object point cloud obtained from the depth image. The color feature $\boldsymbol{f}$ corresponding to these P points in RGB image patches will then be extracted and used in pixel-wise dense fusion.

To capture correlation across RGB and depth modalities, geometric and appearance features belonging to the same pixels are concatenated and fed into a MLP-based feature fusion network (*i.e.* PointNet++ [29]) to generate P pixel-wise fusion feature $\boldsymbol{x} \in \mathbb{R}^{d_{\text{fusion}}}$. Instead of concatenating the pixel-wise fusion feature with a global feature obtained by average pooling as in DenseFusion[39], we adopted the hierarchical pooling of PointNet++ to aggregate local neighboring features. Our motivation without using global features here is to get rid of global information about object instance to verify the effectiveness of a combination of local features as a global description.

3.2. Pixel Pair Pose Regression

This section presents the key components of our pixel-pair pose regression: 1) pixel-pair feature generation; 2) pixel-pair feature encoding; and 3) pixel-pair pose regression and confidence-weighted voting, which are shown in the bottom row of Figure 2.

Generation on Pixel Pairs – In 3D domain, local descriptors for point clouds can encode the neighboring geometric structure of each point but suffer from sensor noises and sparse point distribution. In view of this, multiple point-

pair features [6] consisting of relative position and orientation of two oriented points, are collected as a global representation of an object model. Such a feature representation can still perform robustly to occlusion and sparse point clouds, which is suitable for practical applications. Inspired by robustness and low computational cost of point-pair features [6], we introduce a novel feature encoding layer inserted between pose regression module and feature fusion module investigated in the previous section. Specifically, we divided the pixels corresponding to P pixel-wise fusion features into two sets, and let each pixel of one set form a pair with a pixel of another set, and each pixel can only be selected once. After this processing, the P pixels corresponding to the P pixel-wise fusion feature will eventually generate $P/2$ pixel-pairs. Other more dense strategies such as generation by any two pixels can also be employed, but we choose the aforementioned generation owing to its simplicity.

Pixel-Pair Feature (ppf) Encoding – We have the obtained pixel pairs and now need to generate pixel-pair features. A naive approach to combine local features (*i.e.* $\boldsymbol{x}_k$ where $k = 1, 2, \ldots, P$) into a compact one is to direct concatenate them, which reveal the 1-order statistical information of local features. Encouraged by bilinear models in visual recognition, the key idea of pixel-pair feature encoding is designed based on low-rank bilinear pooling [14]. Specifically, given two pixel-wise feature vectors (*e.g.* $\boldsymbol{a}$ and $\boldsymbol{b}$ from the set $\{\boldsymbol{x}\}^P$), the second-order pooling method [3] can be written as

$$\mathcal{G}_{\text{ppf}}(\boldsymbol{a}, \boldsymbol{b}) = \text{vec}(\boldsymbol{a}\boldsymbol{b}^T) \in \mathbb{R}^{d_{\text{fusion}}^2}. \tag{1}$$

In view of high-dimensionality of $\boldsymbol{a}\boldsymbol{b}^T$, the low-rank bilinear pooling [14] can be adopted to obtain two low-rank matrices $\boldsymbol{U}$ and $\boldsymbol{V}$ to avoid computing $\boldsymbol{a}\boldsymbol{b}^T$ as

$$\mathcal{G}_{\text{ppf}}(\mathbf{a}, \mathbf{b}) = \boldsymbol{P}^T \sigma(\boldsymbol{U}^T \boldsymbol{a} \circ \boldsymbol{V}^T \boldsymbol{b}) \tag{2}$$

where $\circ$ denotes the Hadamard product and we follow [14] use relu nonlinearity after Hadamard product which shows a slightly better result.

Pixel-Pair Pose Regression and Confidence-Weighted Voting – All $P/2$ pixel-pair features $\mathcal{G}_{\text{ppf}}$ are fed into a MLP and each pixel-pair feature generates an object's 6D pose. As a result, a set of $P/2$ pose predictions can be collected. Similar to DenseFusion, we used a self-supervision way to calculate confidence for each pose prediction. This confidence will then be used to generate the final output pose of each object from $P/2$ pose predictions. Specifically, self-supervised confidence c are incorporated as an extra term on pose regression loss which can be depicted as

$$L_{\text{PR}} = \frac{1}{N} \sum_i (L_i^{pair} c_i - w \log(c_i)), \tag{3}$$
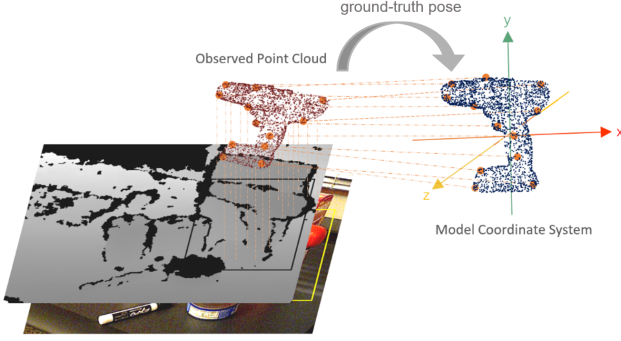
Figure 3. Dense Correspondence Mapping. This module aims to regularize feature learning in pose regression with 3D coordinates under a canonical pose. Specifically, the point cloud in blue is generated by transforming the point clouds sampled from depth image into the model coordinate system, which are used as supervision signals in dense correspondence mapping.

where w is a balancing hyper-parameter and L^{pair} denotes the error of pose predictions of each pixel pair feature as follows [39]:

$$L_i^{pair} = \frac{1}{M} \sum_j \min_{0<k<M} ||(\boldsymbol{R}x_j + \boldsymbol{t}) - (\hat{\boldsymbol{R}}_i x_k + \hat{\boldsymbol{t}}_i)|| \quad (4)$$

where x_j denotes the j-th point in a point set randomly sampled from the object model, M is the size of the point set, $[\boldsymbol{R}|\boldsymbol{t}]$ is the ground-truth pose and $[\hat{\boldsymbol{R}}|\hat{\boldsymbol{t}}]$ is the predicted pose output by each pixel-pair feature. Inspired by the hard voting scheme in 3D point pairs in [6] for feature matching, we used the confidence-weighted average of k predicted pose with the highest confidence to get the final output pose.

3.3. Auxiliary Dense Correspondence Mapping

For obtaining discriminative and robust features per pixels, our method designs an auxiliary dense correspondence mapping task to regularize feature learning with direct mapping from each pixel to a 3D coordinate. As shown in Figure 3, dense correspondence mapping is achieved via regressing per-pixel features $\boldsymbol{x}$ to its corresponding 3D coordinate in the model coordinate system. To this goal, the object point cloud obtained by depth input image is transformed with the inverse of ground truth pose $[\boldsymbol{R}|\boldsymbol{t}]$ to obtain its corresponding point cloud in the model coordinate system. Each 3D coordinate in the transformed point cloud is used as a label of its corresponding output in the dense correspondence mapping branch. Object function, termed as Dense Correspondence Mapping (DCM) Loss, can be thus written as

$$L_{\text{DCM}} = \frac{1}{M} \sum_i ||p_i - \hat{p}_i|| \quad (5)$$

where M represents the points sampled from the observed point cloud, $\hat{p}$ is the prediction of the DCM branch of our

network, and p represents the ground truth coordinate corresponding to $\hat{p}$.

A joint lose to combine both terms for pose regression and dense correspondence mapping can be written as:

$$L_{\text{joint}} = L_{\text{PR}} + \lambda L_{\text{DCM}} \quad (6)$$

where λ is the trade-off parameter between two terms.

3.4. Iterative Pose Refinement

Iterative refinement in DenseFusion [39] has gained significant performance improvement for 6D pose estimation, which is designed based on the concept to correct its pose estimation error in an iterative manner. Similar to DenseFusion [39] and CorrelationFusion [4], we also adopt a refinement network to improve final pose predictions, which encode the resulting point clouds after transformation with previous pose predictions to gradually refining pose in a residual learning manner.

4. Experiments

4.1. Settings

Datasets – To evaluate $\mathcal{W}$-PoseNet comprehensively, we conduct experiments on two popular benchmarks – the YCB-Video dataset [43] and the LineMOD dataset [9]. The YCB-Video dataset consists of 21 objects with different textures and sizes and has 92 videos in total. Following recent work [43, 39, 4], We adopt 80 videos for training, 2949 keyframes from the other 12 videos for testing, and also used an additional 80,000 synthetic images provided by [43]. For a fair comparison, we used the semantic segmentation mask provided by PoseCNN in evaluation by following [43, 39]. The LineMOD dataset contains 15,783 images belonging to 13 low-texture objects placed in a different cluttered environment suffering from the challenges of occlusion, clutter and illumination changes. We follow prior work [39, 4] use 15% of images for training and use the other 85% of images for testing without using additional synthetic data.

Performance Metrics – We adopt AUC of ADD(-S) and ADD(-S)<2cm as performance metrics in the YCB-Video dataset following recent work [43, 39] for comparative evaluation. For asymmetric objects, we use the average distance (ADD) between 3D model points transformed by ground-truth poses and predicted poses to measure the pose estimation error. For symmetric objects, the average closest point distance (ADD-S) is employed to measure the mean error. 6d pose predictions are considered to be correct if the error is smaller than a predefined threshold. We vary the distance threshold from 0 to 10cm to plot an accuracy-threshold curve and compute the area under the curve (AUC) for pose evaluation. Here, ADD(-S) denotes both ADD and ADD-S

Table 1. Comparative evaluation of 6D pose estimation on the YCB-Video Dataset in terms of the ADD(-S)<2cm and the AUC of ADD(-S) metrics. We compare DenseFusion [39], PoseCNN+ICP [43] with the proposed $\mathcal{W}$-PoseNet with and without iterative pose refinement in Sec. 3.4. Objects in bold are symmetric. Results are report in units of %.

Method	without Refinement				with Refinement					
	DenseFusion		$\mathcal{W}$-PoseNet		PoseCNN+ICP		DenseFusion		$\mathcal{W}$-PoseNet	
	AUC	<2cm	AUC	<2cm	AUC	<2cm	AUC	<2cm	AUC	<2cm
002_master_chef_can	**70.7**	**70.7**	69.2	65.9	68.1	51.1	**73.3**	**72.3**	72.0	68.6
003_cracker_box	86.8	88.6	**87.8**	**90.0**	83.4	73.3	**94.2**	**98.2**	91.3	93.7
004_sugar_box	90.8	96.8	**91.5**	**98.5**	**97.5**	99.5	96.5	**100.0**	95.1	99.8
005_tomato_soup_can	84.7	82.8	**87.4**	**84.2**	81.8	76.6	85.5	83.0	**88.9**	**84.6**
006_mustard_bottle	90.9	94.1	**93.4**	**100.0**	**98.0**	98.6	94.7	96.1	96.5	**100.0**
007_tuna_fish_can	**79.5**	**58.5**	77.0	55.9	**83.9**	**72.1**	81.9	62.2	78.8	61.4
008_pudding_box	89.4	94.4	**91.8**	**98.1**	96.6	**100.0**	93.2	98.6	94.5	**100.0**
009_gelatin_box	**95.7**	**100.0**	94.6	**100.0**	98.1	**100.0**	96.7	**100.0**	96.0	**100.0**
010_potted_meat_can	**79.6**	76.9	79.0	**77.8**	83.5	77.9	**83.6**	79.9	82.6	**80.4**
011_banana	76.8	60.2	**87.9**	**87.3**	91.9	88.1	83.7	88.4	**92.8**	**98.9**
019_pitcher_base	87.1	87.2	**92.0**	**100.0**	**96.9**	97.7	**96.9**	**100.0**	95.0	**100.0**
021_bleach_cleanser	**87.5**	**85.4**	85.2	77.4	**92.5**	**92.7**	89.7	90.8	89.5	89.0
024_bowl	86.1	**61.3**	**86.2**	49.8	81.0	54.9	**89.5**	**95.1**	87.7	93.6
025_mug	83.9	**80.5**	**84.9**	79.7	81.1	55.2	**88.9**	88.8	88.2	**90.3**
035_power_drill	83.7	83.1	**91.1**	**98.8**	97.7	99.2	92.7	96.5	93.6	**99.5**
036_wood_block	**89.4**	**98.8**	86.3	96.3	87.6	80.2	**92.8**	**100.0**	87.0	97.5
037_scissors	77.1	50.8	**91.5**	**99.5**	78.4	49.2	77.5	48.6	**90.7**	**97.8**
040_large_marker	89.1	90.6	**90.9**	**96.3**	85.3	87.2	**93.0**	**100.0**	92.5	99.4
051_large_clamp	**71.5**	**78.0**	71.4	74.0	**75.2**	74.9	72.5	78.7	70.8	**79.2**
052_extra_large_clamp	**70.1**	**72.0**	68.0	60.4	64.4	48.8	**69.9**	**74.9**	69.6	72.7
061_foam_brick	92.2	**100.0**	92.5	**100.0**	**97.2**	**100.0**	91.9	**100.0**	92.9	**100.0**
MEAN	83.9	81.5	**85.7**	**85.2**	86.6	79.9	87.6	88.2	**87.9**	**90.8**

metrics for asymmetric and symmetric objects respectively. The ADD(-S)<2cm metric considers the pose as correct if the average distance is smaller than 2cm, which is the minimum tolerance for robotic manipulation. Similarly, for the LineMOD dataset, we use the ADD-S distance for symmetric objects (*i.e.* eggbox and glue) and the ADD distance for the remaining objects having an asymmetric geometry as [39, 30].

Implementation Details – As DenseFusion [39], in our experiments, CNN used for textural feature encoding is composed of Resnet-18 [8] followed by 4 up-sampling as the decoder. RGB and depth images are respectively encoded into 128-dimensional vectors, which are fused into a pixel-wise dense feature x_{fusion}. PointNet++ [29] aggregating pixel-wise features consists of two set abstraction and two feature propagation layers. For training our model, learning rate is set 10^{-4} with Adam optimizer. 23 and 27 epochs are set for training the $\mathcal{W}$-PoseNet and the refinement network respectively, and we refine pose predictions with 2 iterations.

4.2. Comparison with State-of-the-art Methods

We compare the proposed $\mathcal{W}$-PoseNet and the state-of-the-art methods on the YCB-Video and LineMOD datasets, and the results are visualized in Tables 1 and 2. In general,

our method can significantly outperform its competitors in both with and without iterative pose refinement on both performance metrics. Specifically, on the YCB-Video, our network consistently performs better than its backbone Dense-Fusion especially for asymmetric objects with and without post-processing refinement, while similar results are observed on the LineMOD dataset. As the dense feature extraction and fusion part of our $\mathcal{W}$-PoseNet and DenseFusion are identical, the performance gain can only be explained by our design on a pixel-pair combination of local features and auxiliary dense correspondence mapping.

4.3. Ablation Studies

In order to verify the effectiveness of each module of $\mathcal{W}$-PoseNet, we conducted a number of ablation experiments on the LineMOD dataset without iterative pose refinement.

Effects of Pixel-Pair Pose Prediction – Compared to pixel-wise pose estimation in DenseFusion [39], pixel-pair pose regression is verified its effectiveness in Tabel 2. In detail, the only difference between $\mathcal{W}$-PoseNet without DCM Loss and DenseFusion lies in the usage of pixel-pair pose regression, and thus 2% improvement gain can be credited to the effects of Pixel-Pair Pose Prediction.

Table 2. Comparative evaluation of 6D pose estimation in terms of ADD(-S) on the LineMOD dataset. Objects with bold name are symmetric. All methods use depth image. Results are report in units of %.

Method	without Refinement				with Refinement			
	DenseFusion [39]	$\mathcal{W}$-PoseNet w/o DCM Loss	$\mathcal{W}$-PoseNet w/o Pixel Pair Prediction	$\mathcal{W}$-PoseNet	Implicit +ICP[33]	SSD-6D +ICP[12]	DenseFusion [39]	$\mathcal{W}$-PoseNet
ape	79.5	80.1	82.9	**86.3**	20.6	65.0	92.3	**92.8**
bench vi.	84.2	91.7	95.8	**97.4**	64.3	80.0	93.2	**99.6**
camera	76.5	82.7	94.9	**98.3**	63.2	78.0	94.4	**99.0**
can	86.6	88.7	94.8	**97.5**	76.1	86.0	93.1	**99.3**
cat	88.8	87.1	94.2	**97.3**	72.0	70.0	96.5	**99.0**
driller	77.7	82.4	89.5	**95.9**	41.6	73.0	87.0	**97.8**
duck	76.3	81.3	83.2	**93.3**	32.4	66.0	92.3	**96.2**
eggbox	**99.9**	**99.9**	99.6	**99.9**	98.6	**100.0**	99.8	99.9
glue	99.4	99.7	99.5	**99.8**	96.4	**100.0**	**100.0**	99.9
hole p.	79.0	80.9	90.1	**95.7**	49.9	49.0	92.1	**97.3**
iron	92.1	88.1	**98.0**	96.6	63.1	78.0	97.0	**98.6**
lamp	92.3	93.5	96.3	**99.2**	91.7	73.0	95.3	**99.8**
phone	88.0	89.9	93.4	**96.4**	71.0	79.0	92.8	**98.3**
MEAN	86.2	88.2	93.2	**96.4**	64.7	77.0	94.3	**98.2**

Evaluation on the DCM Loss – We compare DenseFusion and $\mathcal{W}$-PoseNet without pixel-pair regression in Table 2, which only differs on the usage of additional dense correspondence mapping branch. It is evident that, owing to introducing auxiliary DCM branch, the $\mathcal{W}$-PoseNet can significantly outperform the state-of-the-art DenseFusion by 7%, which indicates that the DCM loss can effectively improve the quality of per-pixel features. Moreover, both modules in $\mathcal{W}$-PoseNet can further boost the performance in a joint learning manner, which can demonstrate our motivation.

Table 3. Ablation studies in sparsity of pixel-pair features. Models are trained on the LineMOD dataset in terms of ADD(-S) metric.

Number of Pixel-pair	Accuracy
100	95.6%
250	96.4%
500	95.9%
750	95.4%
1000	94.9%

Robust against Sparseness of Pixel-Pair Features – In the LineMOD dataset, we conduct an experiment on the effects of pixel-pair feature sparsity on performance, whose results are shown in Table 3. In detail, we respectively generate $\{100, 250, 500, 750, 1000\}$ pixel-pair features, which are feed into pixel-pair pose regression as input. We find out that the size of pixel-pair features had limited effects on estimation performance, which demonstrates the effectiveness of our light combination of pixel-pair features.

Robustness against Inter-Object Occlusion – To verify the robustness of $\mathcal{W}$-PoseNet against inter-object occlusion, we first calculate the invisible surface percentage for each object instance as DenseFusion [39] did, then evaluate the proposed $\mathcal{W}$-PoseNet with different degrees of occlusion. Specifically, we sampled a certain number of points on object models, projected these points onto their image plane by using the ground-truth poses and the camera intrinsic parameters. Given the observation (*i.e.* depth $d(p)$ on each pixel) in depth images and 2D projection $d(\hat{p})$ of 3D coordinates, if any pixel satisfying $|d(p) - d(\hat{p})| > h$ (h is set to 20mm in our experiment), the pixel is considered under occlusion and can be marked as an invisible pixel. The invisible surface percentage for each instance can be generated via the ratio between the size of invisible points and the total number of sampled points. Figure 5 shows comparative evaluation of several methods in terms of the ADD(-S)<2cm metric. Our $\mathcal{W}$-PoseNet can consistently outperform PoseCNN+ICP [43] and DenseFusion [39] with increasing the invisible surface percentage. On one hand, both $\mathcal{W}$-PoseNet and DenseFusion are more robust than PoseCNN+ICP against inter-object occlusion in view of dense pose predictions on local features. On the other hand, the performance gain of our $\mathcal{W}$-PoseNet sharing the identical feature encoding module of DenseFusion can demonstrate its superior robustness of our dense correspondence regularized pixel pair pose regression to pixel-wise regression in DenseFusion.

Effects of Iterative Pose Refinement – We employ the identical refine network with the same settings as DenseFusion. From Tables 1 and 2, iterative pose refinement can further boost estimation performance of $\mathcal{W}$-PoseNet. More importantly, even with iterative refinement, the proposed network can consistently perform better than DenseFusion
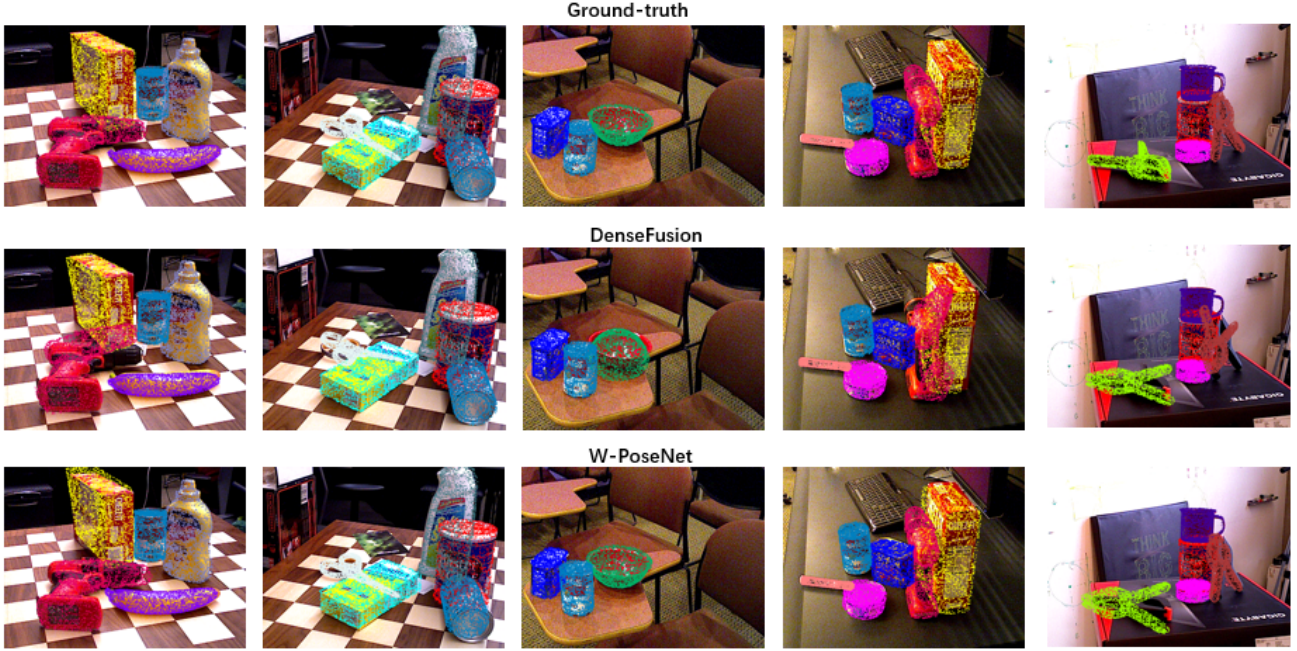
Figure 4. Qualitative results on the YCB-Video dataset from DenseFusion and $\mathcal{W}$-PoseNet. In the first row, 2D images are obtained by projecting the CAD model transformed by ground-truth pose onto the image plane. Images in the second and third rows are obtained by projecting the CAD model transformed by DenseFusion's output pose and $\mathcal{W}$-PoseNet's output pose respectively.
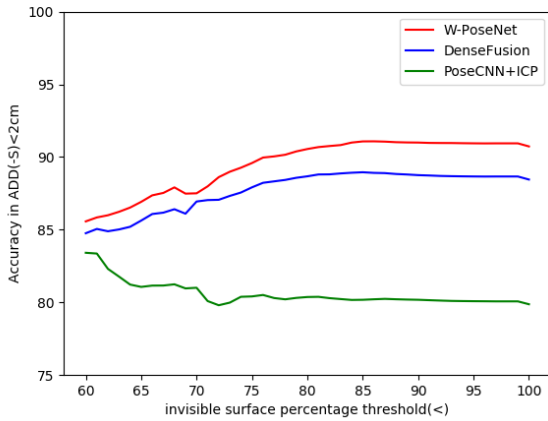

Figure 5. Performance comparison of our $\mathcal{W}$-PoseNet and two state-of-the-art methods under different degree of occlusion.

as well as other competitors.

4.4. Visualization

Figure 4 shows more qualitative results of DenseFusion and $\mathcal{W}$-PoseNet on the YCB-Video dataset. As illustrated in the figure, our method is more robust in cluttered scenes, texture-less objects and occlusion. The fifth column shows an example of a failure due to the fact that 051_large_clamp and 052_extra_large_clamp are completely identical except for their size, making it difficult for the network to distinguish between these two objects.

5. Conclusion

This paper introduces a novel 6D pose estimation network based on two key observations: *extracting discriminative dense features in local regions* and *generating a robust global representation of each object instance*. To this end, we propose two key modules based on the state-of-the-art DenseFusion, pixel-pair pose prediction and dense corresponding mapping utilizing object models, which are verified their effectiveness in the experiments respectively.

Acknowledgements

References

[1] PJ Besl and Neil D McKay. A method for registration of 3-d shapes. *TPAMI*, 1992.

[2] Eric Brachmann, Alexander Krull, Frank Michel, Stefan Gumhold, Jamie Shotton, and Carsten Rother. Learning 6d object pose estimation using 3d object coordinates. In *ECCV*, 2014.

[3] Joao Carreira, Rui Caseiro, Jorge Batista, and Cristian Sminchisescu. Semantic segmentation with second-order pooling. In *ECCV*, 2012.

[4] Yi Cheng, Hongyuan Zhu, Cihan Acar, Wei Jing, Yan Wu, Liyuan Li, Cheston Tan, and Joo-Hwee Lim. 6d pose estimation with correlation fusion. *arXiv preprint arXiv:1909.12936*, 2019.

[5] Andreas Doumanoglou, Rigas Kouskouridas, Sotiris Malassiotis, and Tae-Kyun Kim. Recovering 6d object pose and predicting next-best-view in the crowd. In *CVPR*, 2016.

[6] Bertram Drost, Markus Ulrich, Nassir Navab, and Slobodan Ilic. Model globally, match locally: Efficient and robust 3d object recognition. In *CVPR*, 2010.

[7] Martin A Fischler and Robert C Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 1981.

[8] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016.

[9] Stefan Hinterstoisser, Stefan Holzer, Cedric Cagniart, Slobodan Ilic, Kurt Konolige, Nassir Navab, and Vincent Lepetit. Multimodal templates for real-time detection of texture-less objects in heavily cluttered scenes. In *ICCV*, 2011.

[10] Stefan Hinterstoisser, Vincent Lepetit, Naresh Rajkumar, and Kurt Konolige. Going further with point pair features. In *ECCV*, 2016.

[11] Purushottam Kar and Harish Karnick. Random feature maps for dot product kernels. In *AISTATS*, 2012.

[12] Wadim Kehl, Fabian Manhardt, Federico Tombari, Slobodan Ilic, and Nassir Navab. Ssd-6d: Making rgb-based 3d detection and 6d pose estimation great again. In *ICCV*, 2017.

[13] Wadim Kehl, Fausto Milletari, Federico Tombari, Slobodan Ilic, and Nassir Navab. Deep learning of local rgb-d patches for 3d object detection and 6d pose estimation. In *ECCV*, 2016.

[14] Jin-Hwa Kim, Kyoung-Woon On, Woosang Lim, Jeonghee Kim, Jung-Woo Ha, and Byoung-Tak Zhang. Hadamard product for low-rank bilinear pooling. 2016.

[15] Shu Kong and Charless Fowlkes. Low-rank bilinear pooling for fine-grained classification. In *CVPR*, 2017.

[16] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, 2012.

[17] Yanghao Li, Naiyan Wang, Jiaying Liu, and Xiaodi Hou. Factorized bilinear models for image recognition. In *ICCV*, 2017.

[18] Joerg Liebelt, Cordelia Schmid, and Klaus Schertler. independent object class detection using 3d feature maps. In *CVPR*, 2008.

[19] Tsung-Yu Lin, Aruni RoyChowdhury, and Subhransu Maji. Bilinear cnn models for fine-grained visual recognition. In *ICCV*, 2015.

[20] Eric Marchand, Hideaki Uchiyama, and Fabien Spindler. Pose estimation for augmented reality: a hands-on survey. *TVCG*, 2015.

[21] Frank Michel, Alexander Kirillov, Eric Brachmann, Alexander Krull, Stefan Gumhold, Bogdan Savchynskyy, and Carsten Rother. Global hypothesis generation for 6d object pose estimation. In *CVPR*, 2017.

[22] Arsalan Mousavian, Dragomir Anguelov, John Flynn, and Jana Kosecka. 3d bounding box estimation using deep learning and geometry. In *CVPR*, 2017.

[23] Alejandro Newell, Kaiyu Yang, and Jia Deng. Stacked hourglass networks for human pose estimation. In *ECCV*, 2016.

[24] Markus Oberweger, Mahdi Rad, and Vincent Lepetit. Making deep heatmaps robust to partial occlusions for 3d object pose estimation. In *ECCV*, 2018.

[25] Georgios Pavlakos, Xiaowei Zhou, Aaron Chan, Konstantinos G Derpanis, and Kostas Daniilidis. 6-dof object pose from semantic keypoints. In *ICRA*, 2017.

[26] Sida Peng, Yuan Liu, Qixing Huang, Xiaowei Zhou, and Hujun Bao. Pvnet: Pixel-wise voting network for 6dof pose estimation. In *CVPR*, 2019.

[27] Ninh Pham and Rasmus Pagh. Fast and scalable polynomial kernels via explicit feature maps. In *SIGKDD*, 2013.

[28] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *CVPR*, 2017.

[29] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In *NIPS*, 2017.

[30] Mahdi Rad and Vincent Lepetit. Bb8: A scalable, accurate, robust to partial occlusion method for predicting the 3d poses of challenging objects without using depth. In *ICCV*, 2017.

[31] Mohammad Rastegari, Vicente Ordonez, Joseph Redmon, and Ali Farhadi. Xnor-net: Imagenet classification using binary convolutional neural networks. In *ECCV*, 2016.

[32] Min Sun, Gary Bradski, Bing-Xin Xu, and Silvio Savarese. Depth-encoded hough voting for joint object detection and shape recovery. In *ECCV*, 2010.

[33] Martin Sundermeyer, Zoltan-Csaba Marton, Maximilian Durner, Manuel Brucker, and Rudolph Triebel. Implicit 3d orientation learning for 6d object detection from rgb images. In *ECCV*, 2018.

[34] Alykhan Tejani, Danhang Tang, Rigas Kouskouridas, and Tae-Kyun Kim. Latent-class hough forests for 3d object detection and pose estimation. In *ECCV*, 2014.

[35] Bugra Tekin, Sudipta N Sinha, and Pascal Fua. Real-time seamless single shot 6d object pose prediction. In *CVPR*, 2018.

[36] Andreas ten Pas, Marcus Gualtieri, Kate Saenko, and Robert Platt. Grasp pose detection in point clouds. *IJRR*, 2017.

[37] Jonathan Tremblay, Thang To, Balakumar Sundaralingam, Yu Xiang, Dieter Fox, and Stan Birchfield. Deep object pose estimation for semantic robotic grasping of household objects. *arXiv preprint arXiv:1809.10790*, 2018.

[38] Shubham Tulsiani and Jitendra Malik. Viewpoints and keypoints. In *CVPR*, 2015.

[39] Chen Wang, Danfei Xu, Yuke Zhu, Roberto Martín-Martín, Cewu Lu, Li Fei-Fei, and Silvio Savarese. Densefusion: 6d object pose estimation by iterative dense fusion. In *CVPR*, 2019.

[40] Xing Wei, Yue Zhang, Yihong Gong, Jiawei Zhang, and Nanning Zheng. Grassmann pooling as compact homogeneous bilinear pooling for fine-grained visual classification. In *ECCV*, 2018.

[41] Xing Wei, Yue Zhang, Yihong Gong, Jiawei Zhang, and Nanning Zheng. Grassmann pooling as compact homogeneous bilinear pooling for fine-grained visual classification. In *ECCV*, 2018.

[42] Yu Xiang, Wongun Choi, Yuanqing Lin, and Silvio Savarese. Data-driven 3d voxel patterns for object category recognition. In *CVPR*, 2015.

[43] Yu Xiang, Tanner Schmidt, Venkatraman Narayanan, and Dieter Fox. Posecnn: A convolutional neural network for 6d object pose estimation in cluttered scenes. *arXiv preprint arXiv:1711.00199*, 2017.

[44] Danfei Xu, Dragomir Anguelov, and Ashesh Jain. Pointfusion: Deep sensor fusion for 3d bounding box estimation. In *CVPR*, 2018.

[45] Chaojian Yu, Xinyi Zhao, Qi Zheng, Peng Zhang, and Xinge You. Hierarchical bilinear pooling for fine-grained visual recognition. In *ECCV*, 2018.

[46] Ying Kin Yu, Kin Hong Wong, and Michael Ming-Yuen Chang. Pose estimation for augmented reality applications using genetic algorithm. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, 2005.

[47] Menglong Zhu, Konstantinos G Derpanis, Yinfei Yang, Samarth Brahmbhatt, Mabel Zhang, Cody Phillips, Matthieu Lecce, and Kostas Daniilidis. Single image 3d object detection and pose estimation for grasping. In *ICRA*, 2014.