



Robot-to-robot relative pose estimation using humans as markers

Md Jahidul Islam¹ · Jiawei Mo¹ · Junaed Sattar¹

Received: 28 October 2019 / Accepted: 9 May 2021 / Published online: 16 June 2021
© The Author(s), under exclusive licence to Springer Science+Business Media, LLC, part of Springer Nature 2021

Abstract

In this paper, we propose a method to determine the 3D relative pose of pairs of communicating robots by using human pose-based key-points as correspondences. We adopt a ‘leader-follower’ framework, where at first, the leader robot visually detects and triangulates the key-points using the state-of-the-art pose detector named OpenPose. Afterward, the follower robots match the corresponding 2D projections on their respective calibrated cameras and find their relative poses by solving the perspective-n-point (PnP) problem. In the proposed method, we design an efficient person re-identification technique for associating the mutually visible humans in the scene. Additionally, we present an iterative optimization algorithm to refine the associated key-points based on their local structural properties in the image space. We demonstrate that these refinement processes are essential to establish accurate key-point correspondences across viewpoints. Furthermore, we evaluate the performance of the proposed relative pose estimation system through several experiments conducted in terrestrial and underwater environments. Finally, we discuss the relevant operational challenges of this approach and analyze its feasibility for multi-robot cooperative systems in human-dominated social settings and feature-deprived environments such as underwater.

Keywords Underwater human–robot cooperation · Marine robotics · Underwater visual perception

1 Introduction

Accurate computation of relative pose is essential in multi-robot estimation problems such as cooperative tracking, localization (Kim & Eustice, 2013; Rekleitis et al., 2002), mapping (Johnson-Roberson et al., 2017; Se et al., 2005), path planning (Landa-Torres et al., 2017), and more. Unless global positioning information (*e.g.*, GPS, USBL) is available, the robots need to estimate their positions and orientations relative to each other based on their exteroceptive sensory measurements and noisy odometry (Zhou & Roumeliotis, 2008). This process is necessary for registering their measurements to a common frame of reference to maintain coordination during task execution.

In a cooperative setting, robots with visual sensing capabilities solve the relative pose estimation problem by triangulating mutually visible local features and landmarks. A lack of salient features significantly affects the accuracy of this estimation (Valgren & Lilienthal, 2010), which eventually hampers the overall success of the operation. Such difficulties often arise in poor visibility conditions underwater due to a lower number of point-based salient features and landmarks (Damron et al., 2018; Sattar et al., 2008). Nevertheless, being a low-power passive sensor, cameras have been the choice for exteroceptive perception in many important applications such as inspection of ship hulls and coral reefs (Kim & Eustice, 2013; Dunbabin et al., 2019), 3D reconstructions of archaeological sites (Johnson-Roberson et al., 2017), and human–robot collaborative missions in general (Islam et al., 2018). An important observation is that the proximity of human divers to robots is a fairly common occurrence in these applications and other monitoring and surveying tasks at shallow waterbodies (Kalaitzakis et al., 2020; Manderson et al., 2018). Besides, humans are frequently present and visible in many social scenarios (Islam et al., 2019; Kümmerle et al., 2013) where natural landmarks are not reliably identifiable due to repeated textures, noisy visual conditions, etc. Hence, the problem of having limited natural landmarks can

✉ Md Jahidul Islam
islam034@umn.edu

Jiawei Mo
moxxx066@umn.edu

Junaed Sattar
junaed@umn.edu

¹ Interactive Robotics and Vision Laboratory, Department of CSE, Minnesota Robotics Institute, University of Minnesota-Twin Cities, Minneapolis, MN 55455, USA




Fig. 1 A simplified illustration of 3D relative pose estimation between robot 1 and robot 2 (3). The robots know the transformations between their intrinsically-calibrated cameras and respective global frames, *i.e.*, {1}, {2}, and {3}. Robot 1 is considered as the leader (equipped with a stereo camera) and its pose in global coordinates (${}^1R_G, {}^1t_G$) is known. Robot 2 (3) finds its unknown global pose by cooperatively localizing itself relative to robot 1 using the human pose-based key-points as common landmarks

be alleviated by using mutually visible humans as *markers* (*i.e.*, features correspondences), particularly in human–robot collaborative applications. Despite the potential, the feasibility of using human presence or body-pose for robot-to-robot relative pose estimation has not been explored in the literature.

In this paper, we propose a method for computing six degrees-of-freedom (6-DOF) robot-to-robot transformation between pairs of communicating robots by using mutually detected humans’ pose-based key-points as correspondences. As illustrated in Fig. 1, we adopt a *leader-follower* framework where one of the robots (equipped with a stereo camera) is assigned as a leader. First, the leader robot detects and triangulates 3D positions of the key-points in its own frame of reference. Then the follower robot matches the corresponding 2D projections on its intrinsically calibrated camera and localizes itself by solving the perspective-n-point (PnP) problem (Zheng et al., 2013). It is to be noted that this entire process of *extrinsic calibration* is automatic and does not require prior knowledge about the robots’ initial positions. Additionally, it is straightforward to extend the leader-follower framework for multi-robot teams from the pairwise solutions. Furthermore, if the leader robot has global positioning information, *i.e.*, has a GPS or an USBL receiver, the follower robots can use that information to localize themselves in the global frame as well.

In addition to the conceptual design, we present an end-to-end system with efficient solutions to the practicalities involved in the proposed robot-to-robot pose estimation method (see Sect. 3). As mentioned, we use OpenPose (Cao et al., 2017) for detecting human body-poses in the image space. Although it provides reliable detection performance,

the extracted 2D key-points across different views do not necessarily associate as a correspondence. We propose a twofold solution to this:

- First, we design an efficient person re-identification module by evaluating the hierarchical similarities of the key-point regions in the image space (see Sect. 3.2). It takes advantage of the consistent human pose structures across viewpoints and evaluates their pair-wise similarities for fast body-pose association. We also demonstrate that the state-of-the-art (SOTA) appearance-based person re-identification models fail to provide acceptable performance under single-board real-time constraints.
- Subsequently, we formulate an iterative optimization algorithm to refine the noisy key-point correspondences by further exploiting their local structural properties in respective images (see Sect. 3.3). We demonstrate that the pair-wise key-point refinement is crucial to ensure their validity in a perspective geometric sense.

This two-stage process facilitates efficient and robust key-point associations across viewpoints for accurate robot-to-robot relative pose estimation (see Sect. 4). In this paper, we primarily focus on these two novel modules because the rest of the computational aspects are generic to all multi-robot cooperative pose estimation systems. Nevertheless, we present a fast implementation of the proposed system and evaluate its end-to-end performance over several terrestrial and underwater field experiments. Lastly, we analyze its practical feasibility and discuss various operational considerations (in Sect. 4.5).

2 Related work

2.1 Robot-to-robot relative Pose estimation

The problem of robot-to-robot relative pose estimation has been thoroughly studied for 2D planar robots, particularly for range and bearing sensors. Analytic solutions for determining 3-DOF robot-to-robot transformation using mutual distance and/or bearing measurements involve solving an over-determined system of nonlinear equations (Zhou & Roumeliotis, 2008; Trawny & Roumeliotis, 2010). Similar solutions for the 3D case, *i.e.*, for determining 6-DOF transformation using inter-robot distance and/or bearing measurements, has been proposed as well (Zhou & Roumeliotis, 2011; Trawny et al., 2010). In practice, these analytic solutions are used as an initial estimate for the relative pose, and then iteratively refined by optimization techniques (*e.g.* nonlinear weighted least-squares) to account for the noisy observation and uncertainty in robot motion.

Robots that rely on visual perception (*i.e.*, use cameras as exteroceptive sensors) solve the relative pose estimation problem by triangulating mutually visible features and landmarks (Wang & Wilso, 1992). Therefore, it reduces to solving the PnP problem by using sets of 2D-3D correspondences between geometric features and their projections on respective image planes (Zheng et al., 2013). Although high-level geometric features (*e.g.*, lines, conics) have been proposed, point-based features are typically used in practice for relative pose estimation (Janabi-Sharifi & Marey, 2010). Moreover, the PnP problem is solved either using iterative approaches by formulating the over-constrained system ($n > 3$) as a nonlinear least-squares problem, or by using sets of three non-collinear points ($n = 3$) in combination with Random Sample Consensus (RANSAC) to remove outliers (Fischler & Bolles, 1981). Besides, vision-based approaches often use temporal-filtering methods, the extended Kalman-filter (EKF) in particular, to reduce the effect of noisy measurements in order to provide near-optimal pose estimates (Wang & Wilso, 1992; Janabi-Sharifi & Marey, 2010). On the other hand, it is also common to simplify the relative pose estimation by attaching specially designed calibration-patterns on each robot (Rekleitis et al., 2006). However, this requires that the robots operate at a sufficiently close range, and remain mutually visible.

2.2 Human body-Pose detection

Visual detection of 2D human pose has made significant progress over the last decade. The SOTA methodologies can be categorized into the top-down and bottom-up approaches. The top-down approaches (Gkioxari et al., 2014; Pishchulin et al., 2012) detect the humans in the image space first, and then perform localization and association of their body-parts. One major limitation of these approaches is that their run-times are proportional to the number of persons in the image. Additionally, the robustness of the pose estimation largely depends on the accuracy of their person detectors. In contrast, the bottom-up approaches (Cao et al., 2017; Pishchulin et al., 2016) do not suffer from these two issues. However, they require solving a more computationally challenging inference problem of learning global contextual cues for simultaneous body-part detection and association.

The classical approaches typically use pictorial structures (Ferrari et al., 2008; Andriluka et al., 2009) to model the appearance of human body-parts. A set of densely sampled shape descriptors are used for localizing the body-parts and then classifiers such as AdaBoost, SVMs, etc., are used for detection. Associating the detected body-parts is rather challenging; a mixture of tree-based models are typically used to learn separate pairwise relationships for different body-part configurations (Johnson & Everingham, 2011). Graph-based connectivity models are then used to formulate the inference

(association) as a graph-cut problem. These pairwise connectivity models can be further generalized (Pishchulin et al., 2013) to capture the anatomical relationships among multiple body-parts. Recently proposed approaches use Deep Neural Networks (DNNs) to learn the human pose detection from large training datasets to perform fast and accurate global inference. DeepPose (Toshev & Szegedy, 2014), for instance, formulates the problem as a regression problem and uses a cascade of DNNs to learn the inference in a holistic fashion. On the other hand, OpenPose (Cao et al., 2017) jointly learns to detect and associate using *pose machines* (Ramakrishna et al., 2014). In contrast to DNNs, each module of a pose machine is trained locally; the sequential predictions of these modules are then refined to perform a hierarchical joint inference. Such hierarchical structures facilitate fast inference for multi-person pose estimation in addition to achieving SOTA performance. Due to these compelling reasons, we use OpenPose in this work.

2.3 Human-aware robot control

Human-awareness is important for autonomous mobile robots operating in social settings and human-robot collaborative applications. A large body of literature and systems exist (Islam et al., 2018; Mead & Matarić, 2017) which focus on the areas of understanding human motion, instructions, behaviors, etc. Additionally, tracking human pose relative to a robot is particularly common in applications such as person tracking or following (Islam et al., 2019; Montemerlo et al., 2002), collaborative manipulation (Mainprice & Berenson, 2013), behavior imitation (Lei et al., 2015), etc. However, the feasibility of using humans' presence or their body-poses as markers for robot-to-robot relative pose estimation has not been explored in the literature.

3 System design and methodology

Our proposed robot-to-robot relative pose estimation system incorporates several computational components: detection of human body-poses in images captured from different views (by leader and follower robots), pair-wise association of the detected humans across viewpoints, geometric refinement of the key-point correspondences, and 3D pose estimation of the follower robot relative to the leader. We present a snapshot of the end-to-end computational pipeline in Fig. 2. As in standard multi-robot cooperation, the proposed system requires synchronized communication between the leader and follower robots. From a follower robot's perspective, the primary challenge is to identify the mutually visible humans and then correctly associate their body-poses. Subsequently, geometric refinements of those associated pose-based keypoints are essential for accurate relative pose estimation in

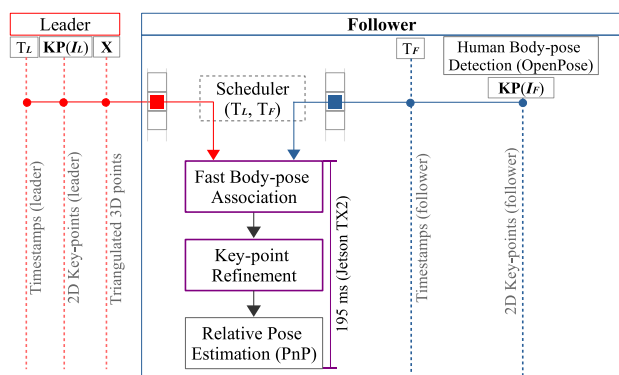


Fig. 2 The end-to-end computational pipeline is outlined from the perspective of a follower robot which shares a clock with the communicating leader robot by using a timestamp-based buffer scheduler for synchronized data registration. The mutually visible human body-pose based key-points are then associated and refined for relative pose estimation. We design these two novel components (marked in purple boxes) to establish robust and accurate key-point correspondences at a fast rate (195 milliseconds per estimation on NVIDIA™ Jetson TX2)

the wild. We design robust and efficient modules to meet these operational requirements in a reasonable computational overhead. In the following sections, we present methodological details of these modules and discuss the relevant design choices.

3.1 Human body-pose detection

OpenPose (Cao et al., 2017) is an open-source library for real-time multi-human 2D pose detection in images, originally developed using Caffe and OpenCV libraries¹. We use a Tensorflow implementation² based on the *MobileNet model* that provides faster inference compared to the original model (also known as the *CMU model*). Specifically, it processes a 368×368 image in 180 milliseconds on the embedded computing board named Jetson TX2 (NVIDIA™, 2014), whereas the original model takes multiple seconds.

OpenPose generates 18 key-points pertaining to the nose, neck, shoulders, elbows, wrists, hips, knees, ankles, eyes, and ears of a human body. As shown in Fig. 3, a subset of these 2D key-points and their pair-wise anatomical relationships are generated for each human. We represent the key-points $\mathbf{KP}(I)$ by a $N_I \times 18$ array where N_I is the number of detected humans in an image I . If a particular key-point is occluded or not detected, then the values are left as $(-1, -1)$. We configure $\mathbf{KP}(I)$ in a way that the first row belongs to the left-most person, the second row belongs to the next left-most person, and gradually the last row belongs to the right-most person in the image. This way of sorting the key-points helps to speed




Fig. 3 Multi-human 2D body-pose detection using OpenPose in various human–robot collaborative settings

up the process of associating the rows of $\mathbf{KP}(I_{leader})$ and $\mathbf{KP}(I_{follower})$. That is, the follower robot needs to make sure that it is pairing the key-points of the *same* individuals. This is important because in practice they might be looking at different individuals, or the same individuals in a different spatial order. Associating multiple persons across different images is a well-studied problem known as *person re-identification (ReId)*.

3.2 Person Re-identification using hierarchical similarities

Although several existing deep visual models provide very good solutions for person ReId (Ahmed et al., 2015; Li et al., 2014), we design a simple and efficient model to meet the real-time single-board computational constraints. The idea is to avoid using a computationally demanding feature extractor by making use of the hierarchical anatomical structures that are already embedded in the key-points. First, we bundle the subsets of key-points in several spatial bounding boxes (BBox) as follows:

- Face BBox: nose, eyes, and ears;
- Upper-body BBox: neck, shoulders, and hips;
- Lower-body BBox: hips, knees, and ankles;
- Left-arm BBox: left shoulder, elbow, and wrist;
- Right-arm BBox: right shoulder, elbow, and wrist;
- Full-body BBox: encloses all the key-points.

Figure 4 illustrates the spatial hierarchy of these BBoxes and their corresponding key-points. They are extracted by spanning the corresponding key-points' coordinate values in both the x and y dimensions. We use an offset (of additional 10% length) in each dimension to capture more spatial information around the key-points. A BBox is discarded if its area

¹ <https://www.github.com/CMU-Perceptual-Computing-Lab/openpose>.

² <https://www.github.com/ildoonet/tf-pose-estimation>.




Fig. 4 An illustration of how the hierarchical body-parts are extracted for person ReID based on their structural similarities; once the persons are associated, the pair-wise key-points are refined and used as correspondences

falls below an empirically chosen threshold of 600 square pixels. We found that BBox areas below this resolution are not always informative and are prone to erroneous results. This happens when the corresponding body-part is either not detected or very far from the camera.

Once the BBox areas are selected, we exploit their pair-wise structural properties as features for person ReID; specifically, we compare the structural similarities between image patches pertaining to the face, upper-body, lower-body, left-arm, right-arm, and the full body of a person. Based on their aggregated similarities, we evaluate the pair-wise association between each person as seen by the leader (in I_{leader}) and by the follower (in $I_{follower}$). The structural similarity (Wang et al., 2004) for a particular pair of single-channel rectangular image-patches (\mathbf{x}, \mathbf{y}) is evaluated based on three properties: luminance $l(\mathbf{x}, \mathbf{y}) = 2\mu_x\mu_y / (\mu_x^2 + \mu_y^2)$, contrast $c(\mathbf{x}, \mathbf{y}) = 2\sigma_x\sigma_y / (\sigma_x^2 + \sigma_y^2)$, and structure $s(\mathbf{x}, \mathbf{y}) = \sigma_{xy} / \sigma_x\sigma_y$; here, μ_x (μ_y) denotes the mean of image patch \mathbf{x} (\mathbf{y}), σ_x^2 (σ_y^2) denotes the variance of \mathbf{x} (\mathbf{y}), and σ_{xy} denotes the cross-correlation between \mathbf{x} and \mathbf{y} . The structural similarity metric (SSIM) is then defined as:

$$SSIM(\mathbf{x}, \mathbf{y}) = l(\mathbf{x}, \mathbf{y})c(\mathbf{x}, \mathbf{y})s(\mathbf{x}, \mathbf{y}) = \frac{2\mu_x\mu_y}{\mu_x^2 + \mu_y^2} \times \frac{2\sigma_{xy}}{\sigma_x^2 + \sigma_y^2}.$$

In order to ensure numeric stability, two standard constants $c_1 = (255k_1)^2$ and $c_2 = (255k_2)^2$ are added as:

$$SSIM(\mathbf{x}, \mathbf{y}) = \frac{2\mu_x\mu_y + c_1}{\mu_x^2 + \mu_y^2 + c_1} \times \frac{2\sigma_{xy} + c_2}{\sigma_x^2 + \sigma_y^2 + c_2}. \tag{1}$$

We use $k_1 = 0.01$, $k_2 = 0.03$, and an 8×8 sliding window in our implementation. Additionally, we resize the patches extracted from I_{leader} so that their corresponding pairs in $I_{follower}$ have the same dimensions. Then, we apply Eq. 1 on every channel (R, G, B) and use their average value as the similarity metric on a scale of $[0, 1]$. Specifically, we use this metric for person ReID as follows:

- We only consider the mutually visible body-parts for evaluating the pair-wise SSIM values. This choice is important to enforce meaningful comparisons; otherwise, it is equivalent to using only the full-body BBox, which we found to be highly inaccurate.
- Each person in $I_{follower}$ is associated with the most similar person corresponding to the maximum SSIM value in I_{leader} . However, the association is discarded if that value is less than a threshold $\delta_{min} = 0.4$ which is chosen by an AUC (area under the curve)-based analysis (see Sect. 4.2). This reduces the risk of inaccurate associations, particularly when there are mutually exclusive people in the scene.

3.3 Key-point refinement


Once the specific persons are identified, *i.e.*, the rows of $KP(I_{leader})$ and $KP(I_{follower})$ are associated, the mutually visible key-points are paired together to form correspondences. Although the key-points are ordered and OpenPose localizes them reasonably well, they cannot be readily used as geometric correspondences due to perspective distortions and noise. We attempt to solve this problem by designing an iterative optimization algorithm that refines the noisy correspondences based on their structural properties in a 32×32 neighborhood. By denoting $\phi_I(\mathbf{p})$ as the 32×32 image-patch centered at $\mathbf{p} = [p_x, p_y]^T$ in image I , we define a loss for each correspondence ($\mathbf{p}_l \in I_{leader}$, $\mathbf{p}_f \in I_{follower}$) as:

$$L(\mathbf{p}_l, \mathbf{p}_f) = 1 - SSIM(\phi_{I_{leader}}(\mathbf{p}_l), \phi_{I_{follower}}(\mathbf{p}_f)). \tag{2}$$


Then, we refine each initial key-point correspondence ($\mathbf{p}_l^0, \mathbf{p}_f^0$) by minimizing the following function:

$$\mathbf{p}_f^* = \underset{\mathbf{p}}{\operatorname{argmin}} L(\mathbf{p}_l^0, \mathbf{p}) \text{ s. t. } \|\mathbf{p} - \mathbf{p}_f^0\|_\infty < 32. \tag{3}$$

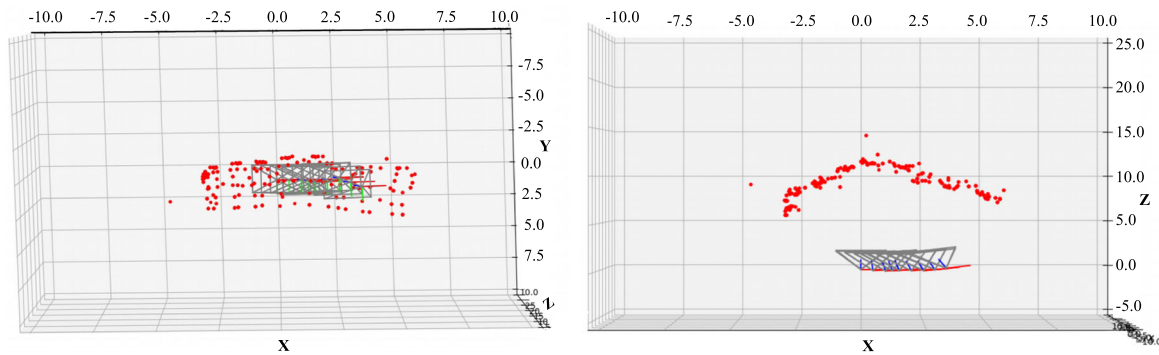
As Eq. 3 suggests, we fix $\mathbf{p}_l = \mathbf{p}_l^0$ and refine $\mathbf{p}_f = \mathbf{p}_f^0$ to maximize $SSIM(\phi_{I_{leader}}(\mathbf{p}_l), \phi_{I_{follower}}(\mathbf{p}_f))$. In our implementation, we adopt a gradient-based refinement algorithm



(a) A group of people seen from multiple views and their 2D body-poses (detected by OpenPose).



(b) Person association and pose-based key-point correspondences for a particular image pair; a unique identifier is assigned to each association, matched key-points are shown in green lines for the right-most person.



(c) The reconstructed 3D key-points of the humans' structure and the estimated camera poses (up-to scale).

Fig. 5 Results of estimating *structure from motion* using only human pose-based key-points as features

that performs the following iterative update:

$$\mathbf{p}_f^{t+1} = \mathbf{p}_f^t - \eta \cdot \nabla L(\mathbf{p}_f^t, \mathbf{p}_f^t). \quad (4)$$

We follow the procedures suggested in (Avanaki, 2009; Otero and Vrscay, 2014) for computing the gradient of SSIM. For fast processing, we vertically stack all the key-points and their gradients to perform the optimization simultaneously with a fixed learning rate of $\eta = 0.003$ for a maximum iteration of 100. We present empirical validations for the choices of the refinement resolution and other hyper-parameters in Sect. 4.2.

3.4 Robot-to-robot pose estimation

Once the mutually visible key-points are associated and refined, the follower robot uses the corresponding 3D positions (provided by the leader) to estimate its relative pose by solving a PnP problem. Thus, we require that the leader

robot is equipped with a stereo camera (or an RGBD camera) so that it can triangulate the refined key-points using epipolar constraints (or use the depth sensor) to represent the key-points in 3D. Let \mathbf{x}_l denote the 3D locations of the key-points in the leader's coordinate frame, and \mathbf{p}_f denote their corresponding 2D projections on the follower's camera. Then, assuming the cameras are synchronized, the PnP problem is formulated as follows:

$$\mathbf{T}_f^l = \operatorname{argmin}_{\mathbf{T}_f^l} \|\mathbf{p}_f - \mathbf{K}_f \mathbf{T}_f^l \mathbf{x}_l\|^2. \quad (5)$$

Here, \mathbf{K}_f is the intrinsic matrix of the follower's camera and \mathbf{T}_f^l is its 6-DOF transformation relative to the leader. In our implementation, we follow the standard iterative solution for PnP using RANSAC (Zheng et al., 2013).

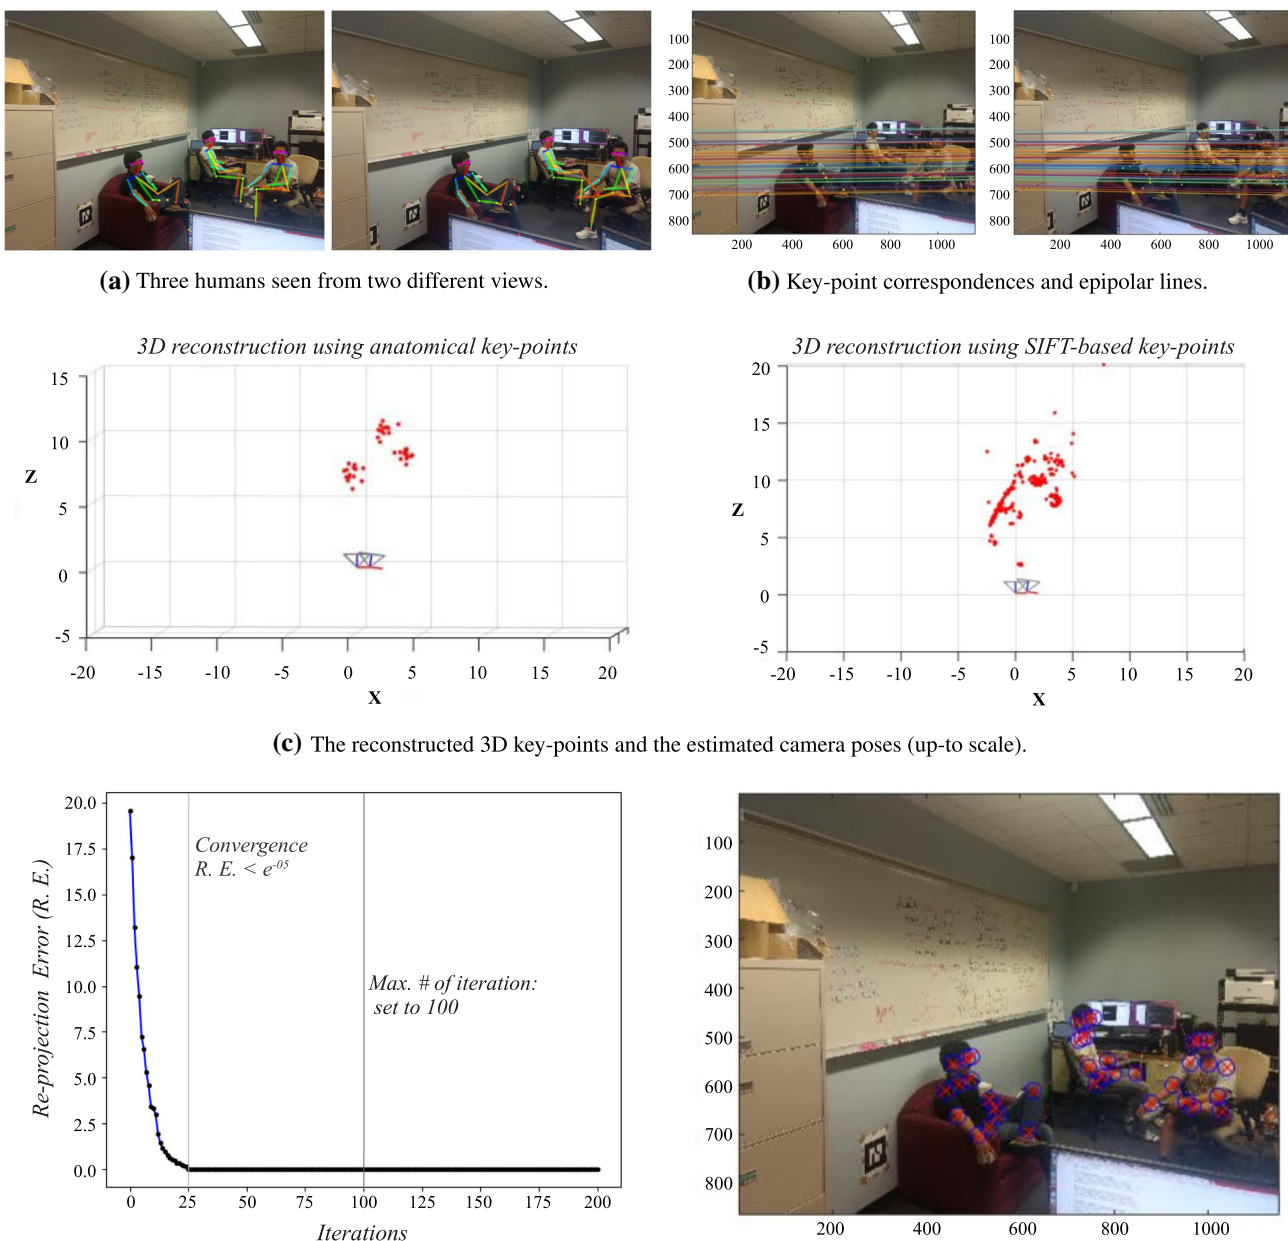


Fig. 6 Structure from motion for a two-view case using only human pose-based key-points as features (Color figure online)

4 Experimental analysis

We conduct several experiments with 2-DOF and 3-DOF robots to evaluate the performance and feasibility of the proposed relative pose estimation method. We present these experimental details, analyze the results, and discuss various operational considerations in the following sections.

4.1 Proof of concept: structure from motion

At first, we perform experiments to validate that the human pose-based key-points can be used as geometric correspondences for relative pose estimation. As illustrated in Fig. 5a, we emulate an experimental set-up for structure from motion with humans; we use an intrinsically calibrated monocular camera to capture a group of nine (static) people from multiple views. Here, the goal is to estimate the camera poses and


Table 1 A quantitative performance comparison for various person ReId models on standard datasets; a set 150 test images are used for comparison from each dataset

Person ReId models	Market-1501 Dataset		CUHK-03 Dataset		FPS (Jetson TX2)
	Rank-1 Acc. (%)	MAP (%)	Rank-1 Acc. (%)	MAP (%)	
Aligned ReId	92.90	90.12	67.15	68.03	0.67
Deep person ReId	85.86	68.24	62.26	65.15	0.33
Tripled-loss ReId	85.25	74.88	72.75	60.27	0.62
Proposed person ReId	75.67	72.26	57.82	54.91	7.45


Table 2 Effectiveness of the proposed person ReId method on real-world data; each set contains 100 images of multiple humans in ground and underwater scenarios

Person ReId models	Set A (1–2 humans per image)		Set B (3–5 humans per image)	
	Rank-1 Acc. (%)	FPS (Jetson TX2)	Rank-1 Acc. (%)	FPS (Jetson TX2)
Aligned ReId	62.75	0.62	56.65	0.48
Deep person ReId	55.32	0.29	42.36	0.12
Tripled-loss ReId	55.15	0.58	44.85	0.44
Proposed person ReId	76.55	6.81	71.56	5.45

Bold indicates best score



(a) ROC curve for the person ReID performance (AUC = 0.938): a total of 20 thresholds ($\delta_{min} \in [0, 1]$) are considered in the evaluation; the point marked in red corresponds to $\delta_{min} = 0.4$.



(b) Fraction of key-point correspondences that fall within various resolutions of respective refinement regions; the vertical red line corresponds to a refinement region of 32×32 .


Fig. 7 Empirical selection of hyper-parameters: **a** SSIM threshold for pose association in the proposed person ReId module and **b** resolution of the key-point refinement region. The evaluation is performed on the

combined set of 250 images containing a total of 687 person associations with 8256 key-point correspondences

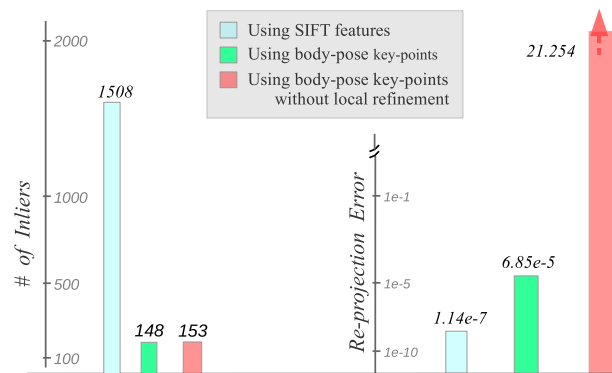
reconstruct the 3D structures of the humans using only their body-poses as features.

In the evaluation, we first use OpenPose to detect the human pose-based 2D key-points in the images (Fig. 5a). Then, we utilize the proposed person ReId and key-point refinement modules to obtain the feature correspondences across multiple views (Fig. 5b). Subsequently, we follow the standard procedures for structure from motion (Hartley & Zisserman, 2003): fundamental matrix computation using 8-point algorithm with RANSAC, essential matrix computation, camera pose estimation by enforcing the Cheirality

constraint, and linear triangulation. Finally, the triangulated 3D points and camera poses are refined using bundle adjustment. As demonstrated in Fig. 5c, the spatial structure of the reconstructed points on the human bodies and the camera poses are consistent with our setup. Results of another experiment for a *two-view* case are shown in Fig. 6, which further validate that the estimated camera poses are comparable to the ground truth, *i.e.*, analogous SIFT feature-based estimation. Next, we demonstrate the effectiveness of our proposed refinement modules in ensuring this robust pose estimation performance.



(a) Inaccurate 3D reconstruction due to using raw key-point correspondences (i.e., without the local refinement processes).



(b) Quantitative performance compared to using SIFT features (in terms of number of inliers and re-projection error).

Fig. 8 Necessity of the proposed key-point refinement process; results correspond to the experiment illustrated in Fig. 5

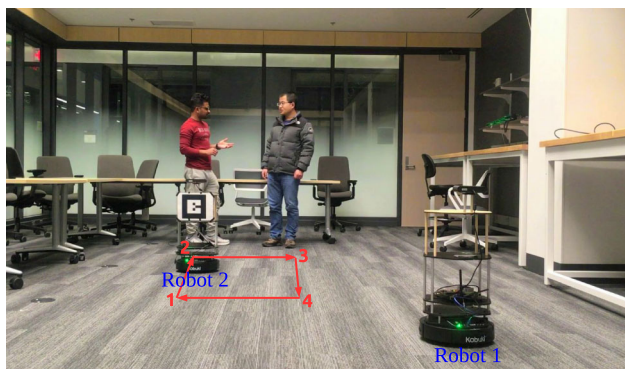



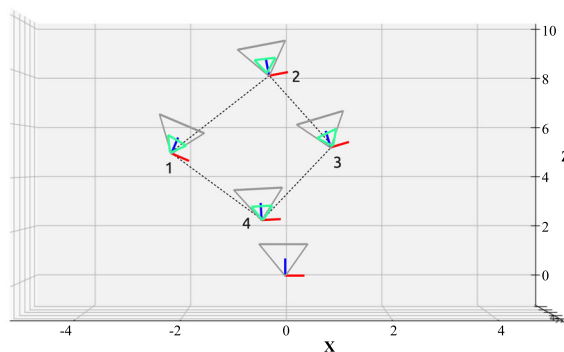
Fig. 9 3-DOF ground experiment: one leader and one follower robot; the follower robot’s trajectory is shown by red arrows (Color figure online)

4.2 Effectiveness of the body-pose association and key-point refinement modules

It is easy to notice that person ReID is essential for associating mutually visible persons across different views. As mentioned in Sect. 3.2, we focus on achieving fast association by making use of the local structural properties around the anatomical key-points in the image space. In contrast, the SOTA person ReID approaches adopt deep visual feature extractors that are computationally demanding. In Table 1, we quantitatively evaluate the SOTA models named Aligned ReID (Zhao et al., 2017), Deep person ReID (Li et al., 2014), and Triple-loss ReID (Zheng et al., 2011) based on accuracy and mean averaged precision (mAP) on two standard datasets. Specifically, a test-set containing 150 instances from the Market-1501 and CUHK-03 datasets are used for the evaluation; also, their run-times on a NVIDIA™ Jetson TX2 are shown for comparison. The results indicate that although these models (once trained on similar data) per-



(a) The leader robot detects the pose-based key-points and shares the 3D locations.



(b) Estimated poses of the follower relative to the leader (the green cones represent the respective ground truth).

Fig. 10 An experiment to evaluate the accuracy of 2D relative pose estimation with two planar robots and two mutually visible humans (Color figure online)

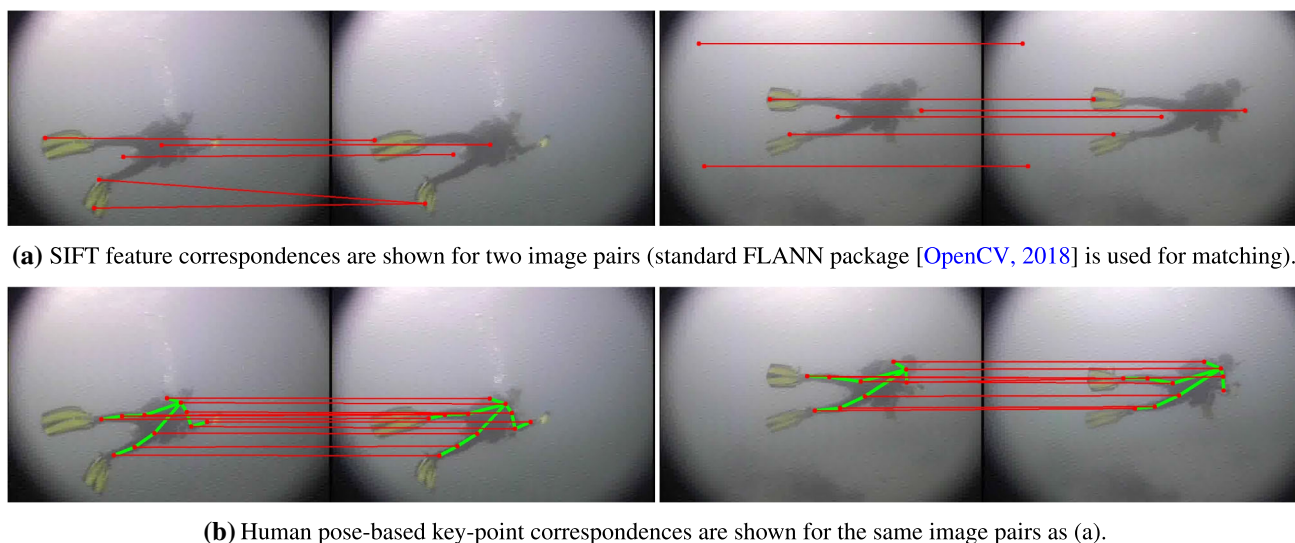


Fig. 11 Illustrations where a lack of natural landmarks limits the utility of standard feature detectors. As seen, presence of a single human in the scene facilitates considerably more anatomical key-point correspondences than the point-based features

form well on standard datasets, they are computationally too expensive for single-board embedded platforms.

Moreover, as demonstrated in Table 2, these off-the-shelf models do not perform that well on high-resolution real-world images. Although their performance can be improved by training on more comprehensive real-world data, the computational complexity remains a barrier. To this end, the proposed person ReID module provides significantly faster run-time and better portability as it does not require rigorous large-scale training. Its only hyper-parameter is the SSIM threshold δ_{min} (see Sect. 3.2), which we select by standard AUC-based analysis of ROC (receiver operating characteristic) curve. As shown in Fig. 7a, we choose $\delta_{min} = 0.4$, which corresponds to 83.5% true-positive and 6.5% false-positive rates for person ReID on the combined test set of 250 images containing 687 person associations. Additionally, we select the key-point refinement resolution through an ablation experiment with 8256 key-point correspondences. We observe that the optimal key-point location is found within 25×25 pixels of the initial estimate by OpenPose for over 96% of the cases. As shown in Fig. 7b, we make a more conservative choice of 32×32 refinement region in our implementation.

Finally, we evaluate the utility and effectiveness of the proposed key-point refinement algorithm based on re-projection errors and compare the results with traditional SIFT feature-based reconstruction. As Fig. 8a demonstrates, the 3D reconstruction and camera pose estimation with raw key-points are inaccurate as the unrefined correspondences are invalid in a perspective geometric sense. As Fig. 8b shows, the average re-projection error for the refined key-points reduces to $6.85e^{-5}$ pixels, which is acceptable considering




Fig. 12 6-DOF underwater experiment: one leader and two follower robots (aerial view)

the fact that there are ten times less anatomical key-points than SIFT feature-based key-points. This evaluation corresponds to the experiment presented in Fig. 5, which shows that the refined key-points constitute accurate scene reconstruction and camera pose estimation. Another qualitative validation of the iterative key-point refinement algorithm and its convergence behavior can be found in Fig. 6.

4.3 Robot-to-robot 3-DOF pose estimation

We also perform experiments for 3-DOF robot-to-robot relative pose estimation with 2D robots. In the particular scenario shown in Fig. 9, we use two planar robots (one leader and one follower) and two mutually visible humans in the scene. The robot with an AR-tag on its back is used as the follower robot while the other robot is used as the leader. The AR-tag is used to obtain the follower's ground truth relative pose for comparison. On the other hand, the leader robot is

equipped with an RGBD camera; it communicates with the follower and shares the 3D locations of the mutually visible key-points. Specifically, it detects the human pose-based 2D key-points and associates the corresponding depth information to represent them in 3D. Subsequently, the follower robot uses this information to localize itself relative to the leader by following the proposed estimation method.

As demonstrated in Fig. 9, we move the follower robot in a rectangular pattern and evaluate the 3-DOF pose estimates relative to the static leader robot. We present the qualitative results in Fig. 10; it shows that the follower robot's pose estimates are very close to their respective ground truth. Overall, we observe an average error of 0.0475% in translation (cm) and a 0.8625° average error in rotation, which is reasonably accurate. We obtain similar qualitative and quantitative performance with a dynamic leader as well. Next, we present field experimental validations of the relative pose estimation performance in feature-deprived underwater scenarios.

4.4 Robot-to-robot 6-DOF pose estimation in adverse underwater visual conditions

As seen in Fig. 11, standard point-based feature detectors fail to generate a large pool of reliable correspondences when there are very few salient features and landmarks in the scene. Consequently, the sampling-based parameter estimation techniques (*e.g.*, RANSAC) often generate inaccurate results in feature-deprived underwater scenarios. However, we demonstrate that human pose-based key-points can still be refined to establish reliable geometric correspondences for robot-to-robot relative pose estimation. Moreover, we get a reasonably large pool of correspondences with only one or two humans in the scene, which is fairly common in cooperative underwater missions.

We perform several field experiments in human–robot collaborative setups; Fig. 12 shows the setup of a particular underwater experiment where we capture human body-poses from different perspectives to estimate the 6-DOF transformations of two follower robots relative to a leader robot. The leader robot is equipped with a stereo camera; hence, the 3D information of the human pose-based key-points is obtained by using stereo triangulation technique. Subsequently, we find the corresponding 2D projections on the follower robots' cameras using the proposed person ReId and key-point refinement processes. Finally, we estimate the follower-to-leader relative poses from their respective PnP solutions.

We present a particular snapshot in Fig. 13a; it illustrates the leader and follower robots' perspectives and the associated human pose-based key-points. Subsequently, Fig. 13b, c demonstrate the geometric validity of those key-point correspondences and the reconstructed 3D points are shown in Fig. 13d. As seen, the estimated 3D structure is consistent

with the mutually visible humans' body-poses. Finally, the estimated 6-DOF poses of the follower robots relative to the leader robot are shown in Fig. 13e.


Such leader-to-follower pose estimates are useful in cooperative diver following (Islam et al., 2019), convoying (Shkurti et al., 2017), and other interactive tasks while operating in close proximity. The robust performance and efficient implementation of the proposed modules make it suitable for use by visually-guided underwater robots in human–robot collaborative applications. However, there are a few practicalities involved which can affect the performance; next, we discuss these aspects and their possible solutions based on our experimental findings.

4.5 Discussion: operational challenges and practicalities


Synchronized cooperation: A major operational requirement of multi-robot cooperative systems is the ability to register synchronized measurements in a common frame of reference, which can be quite challenging in practice. For problems such as ours, an effective solution is to maintain a buffer of time-stamped measurements and register those as a batch using a temporal *sliding window*. We effectively used such timestamp-based buffer schedulers (ROS.org, 2018) in our implementation with reasonable robustness. However, the challenge remains in finding instantaneous relative poses, especially when both robots are in motion. Nevertheless, these aspects are independent of the choice of features/key-points for relative pose estimation and more generic requirements to multi-robot cooperative systems.

Trade-off between robustness and efficiency: It is quite challenging to ensure a fast yet robust performance for visual feature-based body-pose estimation and person ReId on limited computational resources of embedded platforms. This trade-off between robustness and efficiency led us to design fast body-pose association and refinement modules. These efficient modules enable us to achieve an average end-to-end run-time of 375–420 milliseconds for relative pose estimation on Jetson TX2. Note that the proposed person ReId and key-point refinement account for only 195–240 milliseconds (up to nine humans in the scene). Hence, faster human body-pose detectors (than OpenPose) can significantly boost the end-to-end run-time of the system.


In Sect. 4.2, we demonstrated that the proposed person ReId model performs reasonably well in practice despite its simplistic design. One operational benefit in our application is that the humans are seen at once from every perspective; hence, both their appearances and body-poses remain consistent. We provide a demonstration of this benefit in Fig. 14; it shows three *queries* for ReId where humans with similar suit/wearable are seen at various distances and orientations from the camera. Although the *gallery* images contain sev-




(a) A group of people seen from multiple perspectives; the detected key-points and their associations are annotated in respective images.



(b) Feature correspondences for the leader's stereo image pair is shown: (left) SIFT features; (right) human pose-based key-points. As seen, 20 anatomical key-point matches are found compared to only two SIFT feature matches.



(c) The anatomical key-points are pair-wise associated, refined, and then used to project epipolar lines for the stereo image pair.



(d) Stereo triangulation of the human pose-based key-points (seen by the leader robot). (e) Estimated relative poses of the followers robots.

Fig. 13 An underwater experiment for 3D relative pose estimation using one leader and two follower robots

eral humans with similar appearances, we observe that the top three results correspond to best matches both in terms of human appearance and body-pose. We postulate that computing aggregated similarity scores on local pose-based BBoxes contribute to these results. Since the general-purpose person ReId problem is significantly harder and requires more sophisticated computational pipelines (Zhao et al., 2017; Li et al., 2014), our proposed module seems to take advantage of the body-pose consistency across viewpoints for a faster run-time.

Number of humans and relative viewing angle: We observed a couple of other practical issues during the field experiments. First, the presence of multiple humans in the scene helps to ensure reliable pose estimation performance.

We found that two or more mutually visible humans are ideal for establishing a large pool of reliable correspondences. Additionally, the pose estimation performance is affected by the relative viewing angle; specifically, it often fails to find correct associations when the $\angle_{leader-human-follower}$ is larger than (approximately) 135° . This results in a situation where the robots are exclusively looking at opposite sides of the person without enough common key-points. Moreover, other than temporal lags, we did not observe significant deviations in pose estimation performance with an increasing number of robots within this viewing angle; note that we used up to three follower robots in our experiments.

Practical use cases: Other than the operational challenges discussed above, the proposed system does not incur

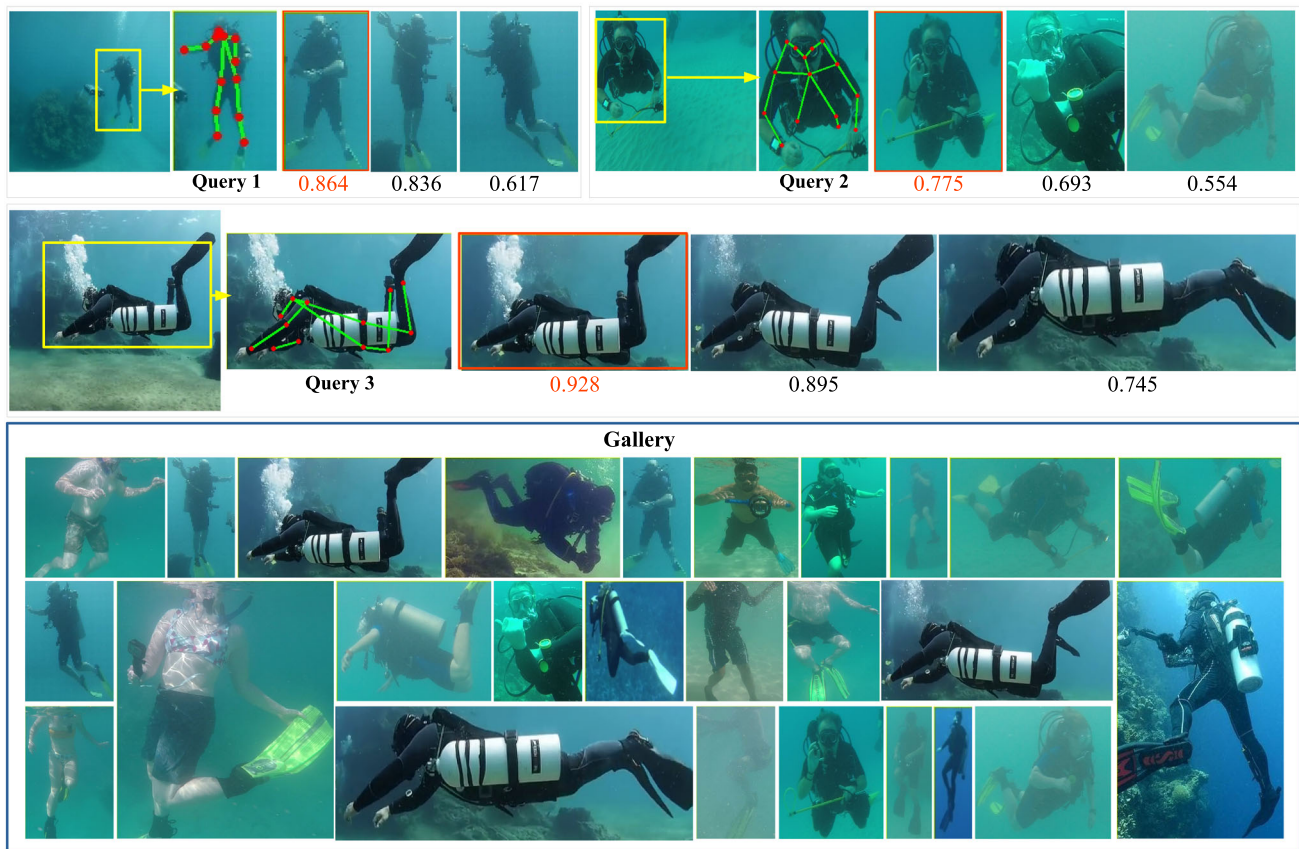


Fig. 14 Three test cases for the proposed person ReId module are shown: each query is matched with a gallery of candidate images (inside the blue box); the top three matches and respective scores are shown

alongside the query image. The scores represent averaged SSIM scores for the mutually visible body-part BBoxes (see Sect. 3.2)

additional constraints or require any prior knowledge such as global positioning information, robots' initial positions, 3D sensing capabilities of the follower robots, etc. As mentioned, it is useful for robot-to-robot pose estimation in human–robot collaborative applications where multiple robots need to maintain spatial coordination during task execution. The prominent use cases are multi-robot conveying, diver following, cooperative mapping, inspection, and surveying. Moreover, it is best suited in applications where the human body-poses are already evaluated for other purposes (*e.g.*, detection/tracking, human–robot communication). It is important to note that the proposed system is not intended to be used as a replacement of full-form visual SLAM or cooperative global localization. Nevertheless, if the leader robot has a GPS or an USBL receiver, it facilitates the follower robots to essentially localize themselves globally by using visual sensing alone.

5 Conclusions and future work

In this paper, we explore the feasibility of using human body-poses as markers to establish reliable multi-view geometric

correspondences and to eventually solve the robot-to-robot relative pose estimation problem. First, we use OpenPose for extracting the pose-based 2D key-points pertaining to the humans in the scene. Then we associate the humans seen from multiple views using an efficient person re-identification model. Subsequently, we refine the key-point correspondences using an iterative optimization algorithm based on their local structural similarities in the image space. Finally, we use the 3D locations of the key-points (triangulated by the leader robot) and their corresponding 2D projections (seen by the follower robot) to formulate a PnP problem and solve for the unknown pose of the follower robot relative to the leader. We perform extensive experiments in terrestrial and underwater environments to investigate the applicability of the proposed relative pose estimation method; the experimental results validate its effectiveness both for 2D and 3D robots. We also discuss the relevant operational challenges and propose efficient solutions to deal with them. In the future, we seek to improve the end-to-end run-time of the proposed system and plan to use it in practical applications such as multi-robot conveying, cooperative source-to-destination planning, etc. Additionally, we aim to investigate the appli-

cability of DensePose (Alp Güler et al., 2018) in our work, which can potentially provide significantly more key-point correspondences per-person compared to OpenPose.

Acknowledgements We would like to thank Hyun Soo Park (Assistant Professor, University of Minnesota) for his valuable insights which immensely enriched this paper. We gratefully acknowledge the support of the MnDrive initiative and thank NVIDIA Corporation for donating two Titan-class GPUs for this research. In addition, we are grateful to the Bellairs Research Institute of Barbados for providing us with the facilities for field experiments; we also acknowledge our colleagues at the IRVLab and the participants of the 2019 Marine Robotics Sea Trials for their assistance in collecting data and conducting the experiments.

References

- Ahmed, E., Jones, M., & Marks, T. K. (2015). An improved deep learning architecture for Person Re-identification. In *Conference on computer vision and pattern recognition (CVPR)*, (pp. 3908–3916). IEEE.
- Alp Güler, R., Neverova, N., & Kokkinos, I. (2018). DensePose: Dense human Pose estimation in the wild. In *Conference on computer vision and pattern recognition (CVPR)* (pp. 7297–7306). IEEE.
- Andriluka, M., Roth, S., & Schiele, B. (2009). Pictorial structures revisited: People detection and articulated Pose estimation. In *Conference on computer vision and pattern recognition (CVPR)* (pp. 1014–1021). IEEE.
- Avanaki, A. N. (2009). Exact global histogram specification optimized for structural similarity. *Optical Review*, 16(6), 613–621.
- Cao, Z., Simon, T., Wei, S.-E., & Sheikh, Y. (2017). Realtime multi-person 2d Pose estimation using part affinity fields. In *Conference on computer vision and pattern recognition (CVPR)* (pp. 7291–7299). IEEE.
- Damron, H., Li, A. Q., & Rekleitis, I. (2018). Underwater surveying via bearing only cooperative localization. In *International conference on intelligent robots and systems (IROS)* (pp. 3957–3963). IEEE/RSJ.
- Dunbabin, M., Dayoub, F., Lamont, R., & Martin, S. (2019). Real-time vision-only perception for robotic coral reef monitoring and management. In *ICRA workshop on underwater robotics perception*. IEEE.
- Ferrari, V., Marin-Jimenez, M., & Zisserman, A. (2008). Progressive search space reduction for human Pose estimation. In *Conference on computer vision and pattern recognition (CVPR)* (pp. 1–8). IEEE.
- Fischler, M. A., & Bolles, R. C. (1981). Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24(6), 381–395.
- Gkioxari, G., Hariharan, B., Girshick, R., & Malik, J. (2014). Using K-poselets for detecting people and localizing their keypoints. In *Conference on computer vision and pattern recognition (CVPR)* (pp. 3582–3589). IEEE.
- Hartley, R., & Zisserman, A. (2003). *Multiple view geometry in computer vision*. Cambridge: Cambridge University Press.
- Islam, M. J., Ho, M., & Sattar, J. (2018). Understanding human motion and gestures for underwater human–robot collaboration. *Journal of Field Robotics (JFR)*, 1–23.
- Islam, M. J., Hong, J., & Sattar, J. (2019). Person-following by autonomous robots: A categorical overview. *International Journal of Robotics Research (IJRR)*, 38(14), 1581–1618.
- Janabi-Sharifi, F., & Marey, M. (2010). A Kalman-filter-based method for pose estimation in visual servoing. *Transactions on Robotics (TRO)*, 26(5), 939–947.
- Johnson, S., & Everingham, M. (2011). Learning effective human Pose estimation from inaccurate annotation. In *Conference on computer vision and pattern recognition (CVPR)* (pp. 1465–1472). IEEE.
- Johnson-Roberson, M., Bryson, M., Friedman, A., Pizarro, O., Troni, G., Ozog, P., & Henderson, J. C. (2017). High-resolution underwater robotic vision-based mapping and three-dimensional reconstruction for archaeology. *Journal of Field Robotics (JFR)*, 34(4), 625–643.
- Kalaitzakis, M., Cain, B., Vitzilaios, N., Rekleitis, I., & Moulton, J. (2020). A marsupial robotic system for surveying and inspection of freshwater ecosystems. *Journal of Field Robotics (JFR)*.
- Kim, A., & Eustice, R. M. (2013). Real-time visual SLAM for autonomous underwater hull inspection using visual saliency. *IEEE Transactions on Robotics (TRO)*, 29(3), 719–733.
- Kümmerle, R., Ruhnke, M., Steder, B., Stachniss, C., & Burgard, W. (2013). A navigation system for robots operating in crowded urban environments. In *International conference on robotics and automation (ICRA)* (pp. 3225–3232). IEEE.
- Landa-Torres, I., Manjarres, D., Bilbao, S., & Del Ser, J. (2017). Underwater robot task planning using multi-objective meta-heuristics. *Sensors*, 17(4), 762.
- Lei, J., Song, M., Li, Z.-N., & Chen, C. (2015). Whole-body humanoid robot imitation with pose similarity evaluation. *Signal Processing*, 108, 136–146.
- Li, W., Zhao, R., Xiao, T., & Wang, X. (2014). Deepreid: Deep filter pairing neural network for person re-identification. In *Conference on computer vision and pattern recognition (CVPR)* (pp. 152–159). IEEE.
- Mainprice, J. & Berenson, D. (2013). Human–robot collaborative manipulation planning using early prediction of human motion. In *International conference on intelligent robots and systems (IROS)* (pp. 299–306). IEEE/RSJ.
- Manderson, T., Higuera, J. C. G., Cheng, R., & Dudek, G. (2018). Vision-based autonomous underwater swimming in dense coral for combined collision avoidance and target selection. In *IEEE/RSJ international conference on intelligent robots and systems (IROS)* (pp. 1885–1891). IEEE.
- Mead, R., & Matarić, M. J. (2017). Autonomous human–robot proxemics: Socially aware navigation based on interaction potential. *Autonomous Robots*, 41(5), 1189–1201.
- Montemerlo, M., Thrun, S., & Whittaker, W. (2002). Conditional particle filters for simultaneous mobile robot localization and people-tracking. In *International conference on robotics and automation (ICRA)* (Vol. 1, pp. 695–701). IEEE.
- NVIDIA™ (2014). Embedded Computing Boards. developer.nvidia.com/embedded/jetson-tx2. Accessed 2 August 2019.
- OpenCV. (2018). Fast library for approximate nearest neighbors (FLANN)-based 2D feature matching algorithm. https://docs.opencv.org/3.4/d5/d6f/tutorial_feature_flann_matcher.html. Accessed 20 June 2020.
- Otero, D. & Vrscay, E. R. (2014). Solving optimization problems that employ structural similarity as the fidelity measure. In *International conference on image processing, computer vision, and pattern recognition (IPCV)* (p. 1).
- Pishchulin, L., Andriluka, M., Gehler, P., & Schiele, B. (2013). Poselet conditioned pictorial structures. In *Conference on computer vision and pattern recognition (CVPR)* (pp. 588–595). IEEE.
- Pishchulin, L., Insafutdinov, E., Tang, S., Andres, B., Andriluka, M., Gehler, P. V., & Schiele, B. (2016). DeepCut: Joint subset partition and labeling for multi person Pose estimation. In *Conference on computer vision and pattern recognition (CVPR)* (pp. 4929–4937). IEEE.

- Pishchulin, L., Jain, A., Andriluka, M., Thormählen, T., & Schiele, B. (2012). Articulated people detection and pose estimation: Reshaping the future. In *Conference on computer vision and pattern recognition (CVPR)* (pp. 3178–3185). IEEE.
- Ramakrishna, V., Munoz, D., Hebert, M., Bagnell, J. A., & Sheikh, Y. (2014). Pose machines: Articulated Pose estimation via inference machines. In *European conference on computer vision (ECCV)* (pp. 33–47). Springer.
- Rekleitis, I., Meger, D., & Dudek, G. (2006). Simultaneous planning, localization, and mapping in a camera sensor network. *Robotics and Autonomous Systems*, 54(11), 921–932.
- Rekleitis, I. M., Dudek, G., & Milios, E. E. (2002). Multi-robot cooperative localization: A study of trade-offs between efficiency and accuracy. In *International conference on intelligent robots and systems (IROS)* (Vol. 3, pp. 2690–2695). IEEE/RSJ.
- ROS.org. (2018). ROS time synchronizer. http://wiki.ros.org/message_filters. Accessed 20 June 2020.
- Sattar, J., Dudek, G., Chiu, O., Rekleitis, I., Giguere, P., Mills, A., Plamondon, N., Prahacs, C., Girdhar, Y., & Nahon, M., et al. (2008). Enabling autonomous capabilities in underwater robotics. In *International conference on intelligent robots and systems (IROS)* (pp. 3628–3634). IEEE/RSJ.
- Se, S., Lowe, D. G., & Little, J. J. (2005). Vision-based global localization and mapping for mobile robots. *Transactions on Robotics (TRO)*, 21(3), 364–375.
- Shkurti, F., Chang, W.-D., Henderson, P., Islam, M. J., Higuera, J. C. G., Li, J., Manderson, T., Xu, A., Dudek, G., & Sattar, J. (2017). Underwater multi-robot convoying using visual tracking by detection. In *International conference on intelligent robots and systems (IROS)*. IEEE/RSJ.
- Toshev, A. & Szegedy, C. (2014). DeepPose: Human Pose estimation via deep neural networks. In *Conference on computer vision and pattern recognition (CVPR)* (pp. 1653–1660). IEEE.
- Trawny, N. & Roumeliotis, S. I. (2010). On the global optimum of planar, range-based robot-to-robot relative pose estimation. In *International conference on robotics and automation (ICRA)* (pp. 3200–3206). IEEE.
- Trawny, N., Zhou, X. S., Zhou, K., & Roumeliotis, S. I. (2010). Inter-robot transformations in 3D. *Transactions on Robotics (TRO)*, 26(2), 226–243.
- Valgren, C., & Lilienthal, A. J. (2010). SIFT, SURF & seasons: Appearance-based long-term localization in outdoor environments. *Robotics and Autonomous Systems*, 58(2), 149–156.
- Wang, J. & Wilson, W. J. (1992). 3D relative position and orientation estimation using Kalman filter for robot control. In *International conference on robotics and automation (ICRA)* (pp. 2638–2645). IEEE.
- Wang, Z., Bovik, A. C., Sheikh, H. R., Simoncelli, E. P., et al. (2004). Image quality assessment: From error visibility to structural similarity. *Transactions on Image Processing (TIP)*, 13(4), 600–612.
- Zhao, L., Li, X., Zhuang, Y., & Wang, J. (2017). Deeply-learned part-aligned representations for person Re-identification. *IEEE International conference on computer vision (ICCV)* (pp. 3219–3228).
- Zheng, W.-S., Gong, S., & Xiang, T. (2011). Person Re-identification by probabilistic relative distance comparison. In *Conference on computer vision and pattern recognition (CVPR)* (pp. 649–656). IEEE.
- Zheng, Y., Kuang, Y., Sugimoto, S., Astrom, K., & Okutomi, M. (2013). Revisiting the PnP problem: A fast, general and optimal solution. In *International conference on computer vision (ICCV)* (pp. 2344–2351). IEEE.
- Zhou, X. S., & Roumeliotis, S. I. (2008). Robot-to-robot relative Pose estimation from range measurements. *Transactions on Robotics (TRO)*, 24(6), 1379–1393.

- Zhou, X. S. & Roumeliotis, S. I. (2011). Determining the robot-to-robot 3D relative Pose using combinations of range and bearing measurements (Part II). In *International conference on robotics and automation (ICRA)* (pp. 4736–4743). IEEE.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Bangladesh University of Engineering and Technology (BUET), Dhaka.

Md Jahidul Islam is an IEEE student member, and currently a Ph.D. candidate at the Computer Science and Engineering (CSE) Department at the University of Minnesota Twin Cities, supervised by Dr. Junaed Sattar. His research work focuses on the design and development of visual perception techniques to understand human motions, gestures, body poses, and activities for human–robot cooperation. Before starting his Ph.D., he received his B.Sc. (Engg.) and M.Sc. (Engg.) degrees in CSE from



Jiawei Mo is a Ph.D. candidate at the Computer Science and Engineering (CSE) Department at the University of Minnesota Twin Cities. He received his Bachelor's degree from the same department and then started his Ph.D. program at the Interactive Robotics and Vision Laboratory under the supervision of Dr. Junaed Sattar. His research interests include 3D state estimation and visual SLAM.



Junaed Sattar is an IEEE member, and an assistant professor at the Department of Computer Science and Engineering at the University of Minnesota Twin Cities, and a MnDrive (Minnesota Discovery, Research, and Innovation Economy) faculty. Junaed is the founding director of the Interactive Robotics and Vision Lab, where he and his students investigate problems in field robotics, robot vision, human–robot communication, assisted driving and applied (deep) machine learning, not to mention developing rugged robotic systems. Before coming to the UoM, Junaed was a post-doctoral fellow at the University of British Columbia working on service and assistive robotics, and at Clarkson University in upstate New York as an Assistant Professor. Find him at junaedsattar.org, and the IRV Lab at irvlab.cs.umn.edu or [@irvlab](https://twitter.com/irvlab) on Twitter.