# Metric-Driven Learning of Correspondence Weighting for 2-D/3-D Image Registration

Roman Schaffert<sup>1</sup>, Jian Wang<sup>2</sup>, Peter Fischer<sup>2</sup>, Anja Borsdorf<sup>2</sup>, and Andreas Maier<sup>1,3</sup>

 <sup>1</sup> Pattern Recognition Lab, Friedrich-Alexander Universität Erlangen-Nürnberg, Erlangen, Germany
roman.schrom.schaffert@fau.de
<sup>2</sup> Siemens Healthineers AG, Forchheim, Germany
<sup>3</sup> Graduate School in Advanced Optical Technologies (SAOT), Erlangen, Germany

Abstract. Registration of pre-operative 3-D volumes to intra-operative 2-D X-ray images is important in minimally invasive medical procedures. Rigid registration can be performed by estimating a global rigid motion that optimizes the alignment of local correspondences. However, inaccurate correspondences challenge the registration performance. To minimize their influence, we estimate optimal weights for correspondences using PointNet. We train the network directly with the criterion to minimize the registration error. We propose an objective function which includes point-to-plane correspondence-based motion estimation and projection error computation, thereby enabling the learning of a weighting strategy that optimally fits the underlying formulation of the registration task in an end-to-end fashion. For single-vertebra registration, we achieve an accuracy of  $0.74\pm0.26$  mm and highly improved robustness. The success rate is increased from 79.3 % to 94.3 % and the capture range from 3 mm to 13 mm.

Keywords: medical image registration  $\cdot$  2-D/3-D registration  $\cdot$  deep learning  $\cdot$  point-to-plane correspondence model

# 1 Introduction

Image fusion is frequently involved in modern image-guided medical interventions, typically augmenting intra-operatively acquired 2-D X-ray images with pre-operative 3-D CT or MRI images. Accurate alignment between the fused images is essential for clinical applications and can be achieved using 2-D/3-D rigid registration, which aims at finding the pose of a 3-D volume in order to align its projections to 2-D X-ray images. Most commonly, intensity-based methods are employed [8], where a similarity measure between the 2-D image and the projection of the 3-D image is defined and optimized as e.g. described by Kubias et al. [6]. Despite decades of investigations, 2-D/3-D registration remains challenging. The difference in dimensionality of the input images results in an ill-posed problem. In addition, content mismatch between the pre-operative and

intra-operative images, poor image quality and a limited field of view challenge the robustness and accuracy of registration algorithms. Miao et al. [9] propose a learning-based registration method that is build upon the intensity-based approach. While they achieve a high robustness, registration accuracy remains challenging.

The intuition of 2-D/3-D rigid registration is to globally minimize the visual misalignment between 2-D images and the projections of the 3-D image. Based on this intuition, Schmid and Chênes [13] decompose the target structure to local shape patches and model image forces using Hooke's law of a spring from image block matching. Wang et al. [15] propose a point-to-plane correspondence (PPC) model for 2-D/3-D registration, which linearly constrains the global differential motion update using local correspondences. Registration is performed by iteratively establishing correspondences and performing the motion estimation. During the intervention, devices and implants, as well as locally similar anatomies, can introduce outliers for local correspondence search (see Fig. 3a and 3b). Weighting of local correspondences, in order to emphasize the correct correspondences, directly influences the accuracy and robustness of the registration. An iterative reweighted scheme is suggested by Wang et al. [15] to enhance the robustness against outliers. However, this scheme only works when outliers are a minority of the measurements.

Recently, Qi et al. [11] proposed the PointNet, a type of neural network directly processing point clouds. PointNet is capable of internally extracting global features of the cloud and relating them to local features of individual points. Thus, it is well suited for correspondence weighting in 2-D/3-D registration. Yi et al. [16] propose to learn the selection of correct correspondences for wide-baseline stereo images. As a basis, candidates are established, e.g. using SIFT features. Ground truth labels are generated by exploiting the epipolar constraint. This way, an outlier label is generated. Additionally, a regression loss is introduced, which is based on the error in the estimation of a known essential matrix between two images. Both losses are combined during training. While including the regression loss improves the results, the classification loss is shown to be important to find highly accurate correspondences. The performance of iterative correspondence-based registration algorithms (e.g. [13], [15]) can be improved by learning a weighting strategy for the correspondences. However, automatic labeling of the correspondences is not practical for iterative methods as even correct correspondences may have large errors in the first few iterations. This means that labeling cannot be performed by applying a simple rule such as a threshold based on the ground truth position of a point.

In this paper, we propose a method to learn an optimal weighting strategy for the local correspondences for rigid 2-D/3-D registration directly with the criterion to minimize the registration error, without the need of per-correspondence ground truth annotations. We treat the correspondences as a point cloud with extended per-point features and use a modified PointNet architecture to learn global interdependencies of local correspondences according to the PPC registration metric. We choose to use the PPC model as it was shown to enable a high registration accuracy as well as robustness [15]. Furthermore, it is differentiable and therefore lends itself to the use in our training objective function. To train the network, we propose a novel training objective function, which is composed of the motion estimation according to the PPC model and the registration error computation steps. It allows us to learn a correspondence weighting strategy by minimizing the registration error. We demonstrate the effectiveness of the learned weighting strategy by evaluating our method on single-vertebra registration, where we show a highly improved robustness compared to the original PPC registration.

## 2 Registration and Learned Correspondence Weighting

In the following section, we begin with an overview of the registration method using the PPC model. Then, further details on motion estimation (see Sec. 2.2) and registration error computation (see Sec. 2.3) are given, as these two steps play a crucial role in our objective function. The architecture of our network is discussed in Sec. 2.4, followed by the introduction of our objective function in Sec. 2.5. At last, important details regarding the training procedure are given in Sec. 2.6.

## 2.1 Registration Using Point-to-Plane Correspondences

Wang et al. [15] measure the local misalignment between the projection of a 3-D volume V and the 2-D fluoroscopic (live X-ray) image  $I^{\text{FL}}$  and compute a motion which compensates for this misalignment. Surface points are extracted from Vusing the 3-D Canny detector [1]. A set of contour generator points [4]  $\{\mathbf{w}_i\}$ , i.e. surface points  $\mathbf{w}_i \in \mathbb{R}^3$  which correspond to contours in the projection of V, are projected onto the image as  $\{\mathbf{p}_i\}$ , i.e. a set of points  $\mathbf{p}_i \in \mathbb{R}^3$  on the image plane. Additionally, gradient projection images of V are generated and used to perform local patch matching to find correspondences for  $\mathbf{p}_i$  in  $I^{\text{FL}}$ . Assuming that the motion along contours is not detectable, the patch matching is only performed in the orthogonal direction to the contour. Therefore, the displacement of  $\mathbf{w}_i$  along the contour is not known, as well as the displacement along the viewing direction. These unknown directions span the plane  $\Pi_i$  with the normal  $\mathbf{n}_i \in \mathbb{R}^3$ . After the registration, a point  $\mathbf{w}_i$  should be located on the plane  $\Pi_i$ . To minimize the point-to-plane distances  $d(\mathbf{w}_i, \Pi_i)$ , a linear equation is defined for each correspondence under the small angle assumption. The resulting system of equations is solved for the differential motion  $\delta \mathbf{v} \in \mathbb{R}^6$ , which contains both rotational components in the axis-angle representation  $\delta \omega \in \mathbb{R}^3$  and translational components  $\delta \boldsymbol{\nu} \in \mathbb{R}^3$ , i.e.  $\delta \mathbf{v} = (\delta \boldsymbol{\omega}^{\intercal}, \delta \boldsymbol{\nu}^{\intercal})^{\intercal}$ . The correspondence search and motion estimation steps are applied iteratively over multiple resolution levels. To increase the robustness of the motion estimation, the maximum correntropy criterion for regression (MCCR) [3] is used to solve the system of linear equations [15]. The motion estimation is extended to coordinate systems related to the camera coordinates by a rigid transformation by Schaffert et al. [12].

The PPC model sets up a linear relationship between the local point-to-plane correspondences and the differential transformation, i.e. a linear misalignment metric based on the found correspondences. In this paper, we introduce a learning method for correspondence weighting, where the PPC metric is used during training to optimize the weighting strategy for the used correspondences with respect to the registration error.

#### 2.2 Weighted Motion Estimation

Motion estimation according to the PPC model is performed by solving a linear system of equations defined by  $\mathbf{A} \in \mathbb{R}^{N \times 6}$  and  $\mathbf{b} \in \mathbb{R}^N$ , where each equation corresponds to one point-to-plane correspondence and N is the number of used correspondences. We perform the motion estimation in the camera coordinate system with the origin shifted to the centroid of  $\{\mathbf{w}_i\}$ . This allows us to use the regularized least-squares estimation

$$\delta \mathbf{v} = \operatorname*{arg\,min}_{\delta \mathbf{v}'} \left( \frac{1}{N} \| \mathbf{A}_s \delta \mathbf{v}' - \mathbf{b}_s \|_2^2 + \lambda \| \delta \mathbf{v}' \|_2^2 \right) \tag{1}$$

in order to improve the robustness of the estimation. Here,  $\mathbf{A}_s = \mathbf{S} \cdot \mathbf{A}$ ,  $\mathbf{b}_s = \mathbf{S} \cdot \mathbf{b}$ and  $\lambda$  is the regularizer weight. The diagonal matrix  $\mathbf{S} = \text{diag}(\mathbf{s})$  contains weights  $\mathbf{s} \in \mathbb{R}^N$  for all correspondences. As Eq. (1) is differentiable w.r.t.  $\delta \mathbf{v}'$ , we obtain

$$\delta \mathbf{v} = \operatorname{RegPPC}(\mathbf{A}, \mathbf{b}, \mathbf{s}) = (\mathbf{A}_s^{\mathsf{T}} \mathbf{A}_s + N \cdot \lambda \mathbf{I})^{-1} \mathbf{A}_s^{\mathsf{T}} \mathbf{b}_s \quad , \tag{2}$$

where  $\mathbf{I} \in \mathbb{R}^{6 \times 6}$  is the identity matrix. After each iteration, the registration  $\mathbf{T} \in \mathbb{R}^{4 \times 4}$  is updated as

$$\mathbf{T} = \begin{pmatrix} \cos(\alpha)\mathbf{I} + (1 - \cos(\alpha)\mathbf{r}\mathbf{r}^{\mathsf{T}}) + \sin(\alpha)[\mathbf{r}]_{\times} \ \delta\boldsymbol{\nu} \\ 0 & 1 \end{pmatrix} \cdot \hat{\mathbf{T}} \quad (3)$$

where  $\alpha = \|\delta \boldsymbol{\omega}\|, \mathbf{r} = \delta \boldsymbol{\omega}/\|\delta \boldsymbol{\omega}\|, [\mathbf{r}]_{\times} \in \mathbb{R}^{3 \times 3}$  is a skew matrix which expresses the cross product with  $\mathbf{r}$  as a matrix multiplication and  $\hat{\mathbf{T}} \in \mathbb{R}^{4 \times 4}$  is the registration after the previous iteration [15].

#### 2.3 Registration Error Computation

In the training phase, the registration error is measured and minimized via our training objective function. Different error metrics, such as the mean target registration error (mTRE) or the mean re-projection distance (mRPD) can be used. For more details on these metrics, see Sec. 3.3. In this work, we choose the projection error (PE) [14], as it directly corresponds to the visible misalignment in the images and therefore roughly correlates to the difficulty to find correspondences by patch matching for the next iteration of the registration method. The PE is computed as

$$e = \operatorname{PE}(\mathbf{T}, \mathbf{T}^{\operatorname{GT}}) = \frac{1}{M} \sum_{j=1}^{M} \|P_{\mathbf{T}}(\mathbf{q}_j) - P_{\mathbf{T}^{\operatorname{GT}}}(\mathbf{q}_j)\| , \qquad (4)$$



Fig. 1: Modified PointNet [11] architecture used for correspondence weighting. Rectangles with dashed outlines indicate feature vectors (orange for local features, i. e. containing information from single correspondences, and red for global features, i. e. containing information from the entire set of correspondences). Sets of feature vectors (one feature vector per correspondence) are depicted as a column of feature vectors (three correspondences shown here). MLP denotes a multi-layer perceptron, which is applied to each feature vector individually.

where a set of M target points  $\{\mathbf{q}_j\}$  is used and j is the point index.  $P_{\mathbf{T}}(\cdot)$  is the projection onto the image plane under the currently estimated registration and  $P_{\mathbf{T}^{\mathrm{GT}}}(\cdot)$  the projection under the ground-truth registration matrix  $\mathbf{T}^{\mathrm{GT}} \in \mathbb{R}^{4 \times 4}$ . Corners of the bounding box of the point set  $\{\mathbf{w}_i\}$  are used as  $\{\mathbf{q}_j\}$ .

#### 2.4 Network Architecture

We want to weight individual correspondences based on their geometrical properties as well as the image similarity, taking into account the global properties of the correspondence set. For every correspondence, we define the features

$$\mathbf{f}_{i} = \left(\mathbf{w}_{i}^{\mathsf{T}} \mathbf{n}_{i}^{\mathsf{T}} d(\mathbf{w}_{i}, \Pi_{i}) \operatorname{NGC}_{i}\right)^{\mathsf{T}} , \qquad (5)$$

where  $NGC_i$  denotes the normalized gradient correlation for the correspondences, which is obtained in the patch matching step.

The goal is to learn the mapping from a set of feature vectors  $\{\mathbf{f}_i\}$  representing all correspondences to the weight vector  $\mathbf{s}$  containing weights for all correspondences, i. e. the mapping

$$\mathbf{M}_{\boldsymbol{\theta}}: \{\mathbf{f}_i\} \mapsto \mathbf{s} \quad , \tag{6}$$

where  $M_{\theta}$  is our network, and  $\theta$  the network parameters.

To learn directly on correspondence sets, we use the PointNet [11] architecture and modify it to fit our task (see Fig. 1). The basic idea behind PointNet is to process points individually and obtain global information by combining the points in a symmetric way, i. e. independent of order in which the points appear in the input [11]. In the simplest variant, the PointNet consists of a multi-layer perceptron (MLP) which is applied for each point, transforming the respective  $\mathbf{f}_i$  into a higher-dimensional feature space and thereby obtaining a local point descriptor. To describe the global properties of the point set, the resulting local descriptors are combined by max pooling over all points, i. e. for each feature, the

maximum activation over all points in the set is retained. To obtain per-point outputs, the resulting global descriptor is concatenated to the local descriptors of each point. The resulting descriptors, containing global as well as local information, are further processed for each point independently by a second MLP. For our network, we choose MLPs with the size of  $8 \times 64 \times 128$  and  $256 \times 64 \times 1$ , which are smaller than in the original network [11]. We enforce the output to be in the range of [0; 1] by using a softsign activation function [2] in the last layer of the second MLP and modify it to re-scale the output range from (-1; 1) to (0; 1). Our modified softsign activation function  $f(\cdot)$  is defined as

$$f(x) = \left(\frac{x}{1+|x|} + 1\right) \cdot 0.5 \quad , \tag{7}$$

where x is the state of the neuron. Additionally, we introduce a global trainable weighting factor which is applied to all correspondences. This allows for an automatic adjustment of the strength of the regularization in the motion estimation step. Note that the network is able to process correspondence sets of variable size so that no fixed amount of correspondences is needed and all extracted correspondences can be utilized.

## 2.5 Training Objective

We now combine the motion estimation, PE computation and the modified PointNet to obtain the training objective function as

$$\boldsymbol{\theta} = \underset{\boldsymbol{\theta}'}{\operatorname{arg\,min}} \frac{1}{K} \sum_{k=1}^{K} \operatorname{PE}(\operatorname{RegPPC}(\mathbf{A}_{k}, \mathbf{b}_{k}, \operatorname{M}_{\boldsymbol{\theta}'}(\{\mathbf{f}_{i}\}_{k})), \mathbf{T}_{k}^{\operatorname{GT}}) , \qquad (8)$$

where k is the training sample index and K the overall number of samples. Equation (2) is differentiable with respect to s, Eq. (3) with respect to  $\delta \mathbf{v}$  and Eq. (4) with respect to **T**. Therefore, gradient-based optimization can be performed on Eq. (8).

Note that using Eq. (8), we learn directly with the objective to minimize the registration error and no per-correspondence ground-truth weights are needed. Instead, the PPC metric is used to implicitly assess the quality of the correspondences during the back-propagation step of the training and the weights are adjusted accordingly. In other words, the optimization of the weights is driven by the PPC metric.

## 2.6 Training Procedure

To obtain training data, a set of volumes  $\{V\}$  is used, each with one or more 2-D images  $\{I^{\text{FL}}\}$  and a known  $\mathbf{T}^{\text{GT}}$  (see Sec. 3.1). For each pair of images, 60 random initial transformations with an uniformly distributed mTRE are generated [5]. For details on the computation of the mTRE and start positions, see Sec. 3.3.



Fig. 2: Examples of 2-D images used as  $I^{\text{FL}}$  (top row) and the corresponding 3-D images used as V (bottom row) in the registration evaluation. Evaluated vertebrae are marked by a yellow cross in the top row.

Estimation of correspondences at training time is computationally expensive. Instead, the correspondence search is performed once and the precomputed correspondences are used during training. Training is performed for one iteration of the registration method and start positions with a small initial error are assumed to be representative for subsequent registration iterations at test time. For training, the number of correspondences is fixed to 1024 to enable efficient batch-wise computations. The subset of used correspondences is selected randomly for every training step. Data augmentation is performed on the correspondence sets by applying translations, in-plane rotations and horizontal flipping, i.e. reflection over the plane spanned by the vertical axis of the 2-D image and the principal direction. For each resolution level, a separate model is trained.

# **3** Experiments and Results

# 3.1 Data

We perform experiments for single-view registration of individual vertebrae. Note that single-vertebra registration is challenging due to the small size of the target structure and the presence of neighbor vertebrae. Therefore, achieving a high robustness is challenging. We use clinical C-arm CT acquisitions from the thoracic and pelvic regions of the spine for training and evaluation. Each acquisition consists of a sequence of 2-D images acquired with a rotating C-arm. These images are used to reconstruct the 3-D volume. To enable reconstruction, the C-arm geometry has to be calibrated with a high accuracy (the accuracy is  $\leq 0.16$  mm

for the projection error at the iso-center in our case). We register the acquired 2-D images to the respective reconstructed volume and therefore the ground truth registration is known within the accuracy of the calibration. Vertebra are defined by an axis-aligned volume of interest (VOI) containing the whole vertebra. Only surface points inside the VOI are used for registration. We register the projection images (resolution of  $616 \times 480$  pixels, pixel size of 0.62 mm) to the reconstructed volumes (containing around 390 slices with slice resolution of  $512 \times 512$  voxels and voxel size of 0.49 mm). To simulate realistic conditions, we add Poisson noise to all 2-D images and rescale the intensities to better match fluoroscopic images.

The training set consists of 19 acquisitions with a total of 77 vertebrae. For each vertebra, 8 different 2-D images are used. An additional validation set of 23 vertebrae from 6 acquisitions is used to monitor the training process. The registration is performed on a test set of 6 acquisitions. For each acquisition, 2 vertebrae are evaluated and registration is performed independently for both the anterior-posterior and the lateral views. Each set contains data from different patients, i. e. no patient appears in two different sets. The sets were defined so that all sets are representative to the overall quality of the available images, i. e. contain both pelvic and thoracic vertebrae, as well as images with more or less clearly visible vertebrae. Examples of images used in the test set are shown in Fig. 2.

## 3.2 Compared Methods

We evaluate the performance of the registration using the PPC model in combination with the learned correspondence weighting strategy (PPC-L), which was trained using our proposed metric-driven learning method. To show the effectiveness of the correspondence weighting, we compare PPC-L to the original PPC method. The compared methods differ in the computation of the correspondence weights s and the regularizer weight  $\lambda$ . For PPC-L, the correspondence weights  $\mathbf{s}^{\mathrm{L}} = \mathrm{M}_{\boldsymbol{\theta}}(\{\mathbf{f}\})$  and  $\lambda = 0.01$  are used. For PPC, we set  $\lambda = 0$  and the used correspondence weights  $\mathbf{s}^{PPC}$  are the NGC<sub>i</sub> values of the found correspondences, where any value below 0.1 is set to 0, i.e. the correspondence is rejected. Additionally, the MCCR is used in the PPC method only. The minimum resolution level has a scaling of 0.25 and the highest a scaling of 1.0. For the PPC method, registration is performed on the lowest resolution level without allowing motion in depth first, as this showed to increase the robustness of the method. To differentiate between the effect of the correspondence weighting and the regularized motion estimation, we also consider registration using regularized motion estimation. We use a variant where the global weighting factor, which is applied to all points, is matched to the regularizer weight automatically by using our objective function (PPC-R). For the different resolution levels, we obtained a data weight in the range of [2.0; 2.1]. Therefore, we use  $\lambda = 0.01$  and  $\mathbf{s}^{\mathrm{R}} = 2.0 \cdot \mathbf{s}^{\mathrm{PPC}}$ . Additionally, we empirically set the correspondence weight to  $\mathbf{s}^{\text{RM}} = 0.25 \cdot \mathbf{s}^{\text{PPC}}$ , which increases the robustness of the registration while still allowing for a reasonable amount of motion (PPC-RM).

## 3.3 Evaluation Metrics

To evaluate the registration, we follow the standardized evaluation methodology [5, 10]. The following metrics are defined by van de Kraats et al. [5]:

- Mean Target Registration Error: The mTRE is defined as the mean distance of target points under  $\mathbf{T}^{\text{GT}}$  and the estimated registration  $\mathbf{T}^{\text{est}} \in \mathbb{R}^{4 \times 4}$ .
- Mean Re-Projection Distance (mRPD): The mRPD is defined as the mean distance of target points under  $\mathbf{T}^{\text{GT}}$  and the re-projection rays of the points as projected under  $\mathbf{T}^{\text{est}}$ .
- Success Rate (SR): The SR is the number of registrations with with a registration error below a given threshold. As we are concerned with single-view registration, we define the success criterion as a mRPD  $\leq 2 \text{ mm}$ .
- Capture Range (CR): The CR is defined as the maximum initial mTRE for which at least 95% of registrations are successful.

Additionally, we compute the gross success rate (GSR) [9] as well as a gross capture range (GCR) with a success criterion of a mRPD  $\leq 10$  mm in order to further assess the robustness of the methods in case of a low accuracy. We define target points as uniformly distributed points inside the VOI of the registered vertebra. For the evaluation, we generate 600 random start transformations for each vertebra in a range of 0 mm - 30 mm initial mTRE using the methodology described by van de Kraats et al. [5]. We evaluate the accuracy using the mRPD and the robustness using the SR, CR GSR and GCR.

## 3.4 Results and Discussion

Accuracy and Robustness The evaluation results for the compared methods are summarized in Tab. 1. We observe that PPC-L achieves the best SR of 94.3% and CR of 13 mm. Compared to PPC (SR of 79.3% and CR of 3 mm), PPC-R also achieves a higher SR of 88.1% and CR of 6 mm. For the regularized motion estimation, the accuracy decreases for increasing regularizer influence  $(0.79\pm0.22 \text{ mm} \text{ for PPC-R} \text{ and } 1.18\pm0.42 \text{ mm} \text{ for PPC-RM})$ , compared to PPC  $(0.75\pm0.21 \text{ mm})$  and PPC-L  $(0.74\pm0.26 \text{ mm})$ . A sample registration result using PPC-L is shown in Fig. 3d.

For strongly regularized motion estimation, we observe a large difference between the GSR and the SR. While for PPC-R, the difference is relatively

Table 1: Evaluation results for the compared methods. The mRPD is computed for the 2 mm success criterion and is shown as mean  $\pm$  standard deviation.

Method	mRPD [mm]	SR [%]	CR [mm]	GSR $[\%]$	GCR [mm
PPC	$0.75 \pm 0.21$	79.3	3	81.8	3
PPC-R	$0.79 {\pm} 0.22$	88.1	6	90.72	6
PPC-RM	$1.18{\pm}0.42$	59.6	4	95.1	20
PPC-L	$0.74{\pm}0.26$	94.3	13	96.3	22



Fig. 3: Registration example: (a) shows  $I^{\rm FL}$  with one marked vertebra to register. Red dots depict initially extracted (b, c) and final aligned (d) contour points. Green lines depict the same randomly selected subset of correspondences, whose intensities are determined by NGC<sub>i</sub> (b) and learned weights (c). Final PPC-L registration result overlaid in yellow (d). Also see video in the supplementary material.

small (88.1% vs. 90.7%), it is very high for PPC-RM. Here a GSR of 95.1% is achieved, while the SR is 59.6%. This indicates that while the method is robust, the accuracy is low. Compared to the CR, the GCR is increased for PPC-L (22 mm vs. 13 mm) and especially for PPC-RM (20 mm vs. 4 mm). Overall, this shows that while some inaccurate registrations are present in PPC-L, they are very common for PPC-RM.

**Single Iteration Evaluation** To better understand the effect of the correspondence weighting and regularization, we investigate the registration results after one iteration on the lowest resolution level. In Fig. 4, the PE in pixels (computed using  $\{\mathbf{q}_j\}$  as target points) is shown for all cases in the validation set. As in training, 1024 correspondences are used per case for all methods. We observe that for PPC, the error has a high spread, where for some cases, it is



Fig. 4: Histograms showing initial and result projection error (PE) in pixels for a single iteration of registration on lowest resolution level (on validation set, 1024 correspondences per case). Motion estimation was performed using least squares for all methods. For PPC, no motion in depth is estimated (see Sec. 3.2).



Fig. 5: Box plots for distribution of resulting mRPD on the lowest resolution level for successful registrations for different initial mTRE intervalls.

decreased considerably, while for other cases, it is increased. For PPC-R, most cases are below the initial error. However, the error is decreased only marginally, as the regularization prevents large motions. For PPC-L, we observe that the error is drastically decreased for most cases. This shows that PPC-L is able to estimate motion efficiently. An example for correspondence weighting in PPC-L is shown in Fig. 3c, where we observe a set of consistent correspondences with high weights, while the remaining correspondences have low weights.

Method Combinations We observed that while the PPC-RM method has a high robustness (GCR and GSR), it leads to low accuracy. For PPC-L, we observed an increased GCR compared to the CR. In both cases, this demonstrates that registrations are present with a mRPD between 2 mm and 10 mm. As the PPC works reliably for small initial errors, we combine these methods with PPC by performing PPC on the highest resolution level instead of the respective method. We denote the resulting methods as PPC-RM+ and PPC-L+. We observe that PPC-RM+ achieves an accuracy of  $0.74\pm0.18$  mm, an SR of 94.6%and a CR of 18 mm, while PPC-L+ achieves an accuracy of  $0.74\pm0.19$  mm, an SR of 96.1% and a CR of 19 mm. While the results are similar, we note that for PPC-RM+ a manual weight selection is necessary. Further investigations are needed to clarify the better performance of PPC compared to PPC-L on the highest resolution level. However, this result may also demonstrate the strength of MCCR for cases where the majority of correspondences are correct. We evaluate the convergence behavior of PPC-L+ and PPC-RM+ by only considering cases which were successful. For these cases, we investigate the error distribution after the first resolution level. The results are shown in Fig. 5. We observe that for PPC-L+, a mRPD of below 10 mm is achieved for all cases, while for PPC-RM+, higher misalignment of around 20 mm mRPD is present. The result for PPC-L+ is achieved after an average of 7.6 iterations, while 11.8 iterations were performed on average for PPC-RM+ using the stop criterion defined in [15]. In combination, this further substantiates our findings from the single iteration evaluation and shows the efficiency of PPC-L and its potential for reducing the computational cost.

# 4 Conclusion

For 2-D/3-D registration, we propose a method to learn the weighting of the local correspondences directly from the global criterion to minimize the registration error. We achieve this by incorporating the motion estimation and error computation steps into our training objective function. A modified Point-Net network is trained to weight correspondences based on their geometrical properties and image similarity. A large improvement in the registration robustness is demonstrated when using the learning-based correspondence weighting, while maintaining the high accuracy. Although a high robustness can also be achieved by regularized motion estimation, registration using learned correspondence weighting has the following advantages: it is more efficient, does not need manual parameter tuning and achieves a high accuracy. One direction of future work is to further improve the weighting strategy, e.g. by including more information into the decision process and optimizing the objective function for robustness and/or accuracy depending on the stage of the registration, such as the current resolution level. By regarding the motion estimation as part of the network and not the objective function, our model can also be understood in the framework of precision learning [7] as a regression model for the motion, where we learn only the unknown component (weighting of correspondences), while employing prior knowledge to the known component (motion estimation). Following the framework of precision learning, replacing further steps of the registration framework with learned counterparts can be investigated. One candidate is the correspondence estimation, as it is challenging to design an optimal correspondence estimation method by hand.

**Disclaimer:** The concept and software presented in this paper are based on research and are not commercially available. Due to regulatory reasons its future availability cannot be guaranteed.

# References

- 1. Canny, J.: A Computational Approach to Edge Detection. IEEE Trans Pattern Anal Mach Intell (6), 679–698 (1986)
- 2. Elliott, D.L.: A Better Activation Function for Artificial Neural Networks. Tech. rep. (1993)
- Feng, Y., Huang, X., Shi, L., Yang, Y., Suykens, J.A.: Learning with the Maximum Correntropy Criterion Induced Losses for Regression. J Mach Learn Res 16, 993– 1034 (2015)
- Hartley, R., Zisserman, A.: Multiple View Geometry in Computer Vision, p. 200. Cambridge University Press, 2 edn. (2003)
- van de Kraats, E.B., Penney, G.P., Tomaževič, D., van Walsum, T., Niessen, W.J.: Standardized Evaluation Methodology for 2-D-3-D Registration. IEEE Trans Med Imag 24(9), 1177–1189 (2005)
- Kubias, A., Deinzer, F., Feldmann, T., Paulus, D., Schreiber, B., Brunner, T.: 2D/3D Image Registration on the GPU. Pattern Recognition and Image Analysis 18(3), 381–389 (2008)

- Maier, A., Schebesch, F., Syben, C., Würfl, T., Steidl, S., Choi, J.H., Fahrig, R.: Precision Learning: Towards Use of Known Operators in Neural Networks. arXiv preprint arXiv:1712.00374v3 (2017)
- Markelj, P., Tomaževič, D., Likar, B., Pernuš, F.: A Review of 3D/2D Registration Methods for Image-Guided Interventions. Med. Image Anal. 16(3), 642–661 (2012)
- Miao, S., Piat, S., Fischer, P., Tuysuzoglu, A., Mewes, P., Mansi, T., Liao, R.: Dilated FCN for Multi-Agent 2D/3D Medical Image Registration. In: AAAI Conference on Artificial Intelligence (AAAI). pp. 4694–4701 (2018)
- Mitrović, U., Špiclin, Ž., Likar, B., Pernuš, F.: 3D-2D Registration of Cerebral Angiograms: A Method and Evaluation on Clinical Images. IEEE Trans Med Imag 32(8), 1550–1563 (2013)
- Qi, C.R., Su, H., Mo, K., Guibas, L.J.: PointNet: Deep Learning on Point Sets for 3D Classification and Segmentation. In: IEEE Conference on Computer Vision and Pattern Recogition (CVPR). pp. 77–85 (2017)
- Schaffert, R., Wang, J., Fischer, P., Borsdorf, A., Maier, A.: Multi-View Depth-Aware Rigid 2-D/3-D Registration. In: IEEE Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC) (2017)
- Schmid, J., Chênes, C.: Segmentation of X-ray Images by 3D-2D Registration based on Multibody Physics. In: Cremers, D., Saito, I.R.H., Yang, M. (eds.) ACCV 2014. LNCS. vol. 9004, pp. 674–687. Springer (2014)
- Wang, J., Borsdorf, A., Heigl, B., Köhler, T., Hornegger, J.: Gradient-Based Differential Approach for 3-D Motion Compensation in Interventional 2-D/3-D Image Fusion. In: International Conference on 3D Vision (3DV). pp. 293–300 (2014)
- Wang, J., Schaffert, R., Borsdorf, A., Heigl, B., Huang, X., Hornegger, J., Maier, A.: Dynamic 2-D/3-D Rigid Registration Framework Using Point-To-Plane Correspondence Model. IEEE Trans Med Imag 36(9), 1939–1954 (2017)
- Yi, K.M., Trulls, E., Ono, Y., Lepetit, V., Salzmann, M., Fua, P.: Learning to Find Good Correspondences. In: IEEE Conference on Computer Vision and Pattern Recogition (CVPR). pp. 2666–2674 (2018)