

Dilated FCN for Multi-Agent 2D/3D Medical Image Registration

Shun Miao,^{1*} Sebastien Piat,¹ Peter Fischer,² Ahmet Tuysuzoglu,¹
Philip Mewes² Tommaso Mansi,¹ Rui Liao¹

¹Siemens Healthineers, Medical Imaging Technologies, Princeton, NJ, USA

²Siemens Healthineers, Forchheim, Germany

Abstract

2D/3D image registration to align a 3D volume and 2D X-ray images is a challenging problem due to its ill-posed nature and various artifacts presented in 2D X-ray images. In this paper, we propose a multi-agent system with an auto attention mechanism for robust and efficient 2D/3D image registration. Specifically, an individual agent is trained with dilated Fully Convolutional Network (FCN) to perform registration in a Markov Decision Process (MDP) by observing a local region, and the final action is then taken based on the proposals from multiple agents and weighted by their corresponding confidence levels. The contributions of this paper are threefold. First, we formulate 2D/3D registration as a MDP with observations, actions, and rewards properly defined with respect to X-ray imaging systems. Second, to handle various artifacts in 2D X-ray images, multiple local agents are employed efficiently via FCN-based structures, and an auto attention mechanism is proposed to favor the proposals from regions with more reliable visual cues. Third, a dilated FCN-based training mechanism is proposed to significantly reduce the Degree of Freedom in the simulation of registration environment, and drastically improve training efficiency by an order of magnitude compared to standard CNN-based training method. We demonstrate that the proposed method achieves high robustness on both spine cone beam Computed Tomography data with a low signal-to-noise ratio and data from minimally invasive spine surgery where severe image artifacts and occlusions are presented due to metal screws and guide wires, outperforming other state-of-the-art methods (single agent-based and optimization-based) by a large margin.

Introduction

The goal of 2D/3D medical image registration is to find the 6 Degree of Freedom (DoF) pose of a 3D volume (e.g. Computed Tomography (CT), Magnetic Resonance Imaging (MRI) etc), to align its projections with given 2D X-ray images. Reliable 2D/3D registration is a key enabler for image-guided surgeries in modern operating rooms. It brings measurement and plannings done on the pre-operative data into the operating room, and fuse it with intra-operative live 2D X-ray images. It can be used to provide augmented reality image guidance for the surgery, or provide navigation for robotic surgery. Despite that 2D/3D registration has been

actively researched for decades (Markelj et al. 2012), it remains a very challenging and unsolved problem, especially in the complex environment of hybrid operating rooms.

A Markov Decision Process (MDP) formulation for 3D/3D image registration was recently introduced and resulted in a significant robustness improvement (Liao et al. 2017). The original formulation however has three major limitations that make it ineffective for 2D/3D registration in a clinical setup. First, it requires the rough location and size of the target object to be registered as a prior, in order to extract a local region around the target as the agent's observation. However, in 2D/3D registration, the location and size of the target object can vary significantly in 2D X-ray images, due to variations in C-arm geometry and imaging protocols such as collimation factors. Second, there could be various artifacts / interference coming from medical devices in 2D X-ray images and simulation of all of them in the training samples is not practical. Therefore an auto attention mechanism is needed in order to be able to inherently detect the regions with more reliable visual cues to drive registration, and this mechanism is not provided in (Liao et al. 2017). Third, training data need to be sampled extensively from the registration environment with a high DoF (i.e., environmental DoFs include location of agent's observation and pose of the 3D volume), which is associated with an high computational cost. In fact, five million samples are needed as reported in (Liao et al. 2017), even after using location prior knowledge to reduce the DoF by 3. Since data sampling grows exponentially with the DoF, without the location prior knowledge, computational cost would be prohibitively high for our 2D/3D registration problem.

In this paper, we propose a multi-agent system with an auto attention mechanism to address the above three limitations of agent-based registration, and apply it to a challenging 2D/3D registration application for minimally invasive spine surgery. In particular, multiple agents are employed to observe multiple regions of the image, and the system adaptively favors proposals from regions with more distinct visual cues for registration. We furthermore propose a policy network architecture with separate encoding of the fixed and moving images, and a dilated Fully Convolutional Network

Disclaimer: This feature is based on research, and is not commercially available. Due to regulatory reasons its future availability cannot be guaranteed.




Figure 1: Example X-ray images (top row) and corresponding Digitally Reconstructed Radiographs (DRRs) (bottom row). The first four X-ray images are clinical data from spine surgery, which contain various highly opaque metal objects and have very different Field-of-VIEWS (FoVs). The last three X-ray images are from CBCT data, which have a relatively low SNR due to a very small dynamic range.

(FCN) based training strategy to train all observation regions in each back propagation. This strategy significantly reduces the DoFs of the registration environment, and as a result training efficiency is improved by an order of magnitude compared to standard CNN-based training method. The proposed FCN structure also naturally supports our multi-agent system in the application phase for efficient 2D/3D registration. We demonstrate that when there are severe image artifacts and occlusions, e.g. due to metal screws and guide wires, the proposed method significantly outperforms single agent-based method and other state-of-the-art optimization-based methods by reducing the gross failure rate (GRF) by 10 fold.

Related Work

Optimization-based 2D/3D Registration

The most widely adopted formulation for 2D/3D registration is to solve it as an optimization problem to minimize a cost function that quantifies the quality of alignment between the images being registered (Markelj et al. 2012)(Gendrin et al. 2012)(Schmid and Chênes 2014)(Miao, Liao, and Zheng 2011). The success of these approaches depends on the global optimization of the cost function. Existing image-based cost functions are mostly based on low level metrics, e.g., Mutual Information (MI) (Maes et al. 1997), Cross Correlation (CC) (Knaan and Joskowicz 2003), Gradient Correlation (GC) (De Silva et al. 2016) and etc. These metrics compare images directly at the pixel level without understanding the higher level structures in the image. As a result, on images with a low signal-to-noise ratio (SNR) and/or severe image artifacts like examples shown in Fig. 1, they often have numerous local minima, which makes it extremely challenging to locate the global solution using mathematical optimization strategy, e.g. Powell’s method, Nelder-Mead, BFGS, CMA-ES. Although there are global optimization methods, e.g., simulated annealing (Van Laarhoven and Aarts 1987) and genetic algorithm (Deb et al. 2002), they require comprehensive sampling of the parameter space, which leads to a prohibitively high computa-

tional cost. Few attempts have been made to seek heuristic semi-global optimization to seek a proper balance between robustness and computational cost (Uneri et al. 2017)(Lau and Chung 2006).

Learning-based Image Registration

Several attempts have been made recently to tackle 2D/3D registration problems using learning-based approaches. Hand crafted features and simple regression modules like Multi-Layer Perceptron and linear regression were proposed to regress 2D/3D registration parameters (Gouveia et al. 2015) (Chou et al. 2013). A CNN-based regression approach was introduced in (Miao, Wang, and Liao 2016) to solve 2D/3D registration for 6 DoF device pose estimation from X-ray images, and later on a domain adaptation scheme was proposed (Zheng, Miao, and Liao 2017) to improve the performance generalization of the CNN model on unseen real data. In (Wohlhart and Lepetit 2015), a range image descriptor was learned via CNN to differentiate 3D pose of the object, and was used for 6 DoF pose estimation. While the CNN-based methods improve robustness of 2D/3D registration, they are limited to registering only highly opaque objects with a fixed shape described by a CAD model (e.g., rigid metal devices), because they aimed at explicitly modeling the appearance-pose relationship. Therefore, it cannot solve 2D/3D registration of anatomical structures with varying shapes across patients.

Learning-based approaches have also been explored for 2D/2D and 3D/3D image registration, where the goal is to find the spatial transformation to align images of the same dimension. Unsupervised learning using CNN was proposed in (Wu, Qi, and Shen 2006) to extract features for deformable registration. These features however were extracted separately from the image pairs and therefore may not be optimal for registration purpose. Optical flow estimation between 2D RGB images has been proposed using CNN via supervised learning in (Weinzaepfel et al. 2013)(Dosovitskiy et al. 2015)(Ilg et al. 2016). A learnable module, called spatial transformer network (STN), was introduced

in (Jaderberg et al. 2015). The focus of STN was not an accurate alignment of two images, but a rough transformation of a single input image to a canonical form. In a recent work, rigid-body 3D/3D image registration is formulated as a MDP and a policy network is trained to perform image registration (Liao et al. 2017). Motivated by this work, we design a 2D/3D registration system by utilizing multiple agents coupled with an auto attention mechanism for high robustness, and train the agent using dilated FCN-based strategy for high efficiency.

Background

2D/3D Registration

Given a 3D CT volume $J : \mathbb{R}^3 \rightarrow \mathbb{R}$, a projection image can be calculated following the X-ray imaging model (Bushberg and Boone 2011):

$$H_T(\mathbf{p}) = \int J(T^{-1} \circ \mathbf{L}(\mathbf{p}, r)) dr, \quad (1)$$

where $I(\mathbf{p})$ is the intensity of the synthetic X-ray image at point \mathbf{p} , $\mathbf{L}(\mathbf{p}, r)$ is the line connecting the X-ray source and the point \mathbf{p} , parameterized by r , and $T : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ is the geometric transformation of the 3D volume. Such projection image is referred to as DRR, and can be computed using the Ray-Casting algorithm (Kruger and Westermann 2003).

In 2D/3D registration problems, a 3D volume $J(\cdot)$, a 2D X-ray image $I(\cdot)$ and the camera model of the X-ray images $\mathbf{L}(\cdot)$ are given. The goal is to find the transformation T that aligns the projection of the 3D volume $H_T(\cdot)$ with the X-ray image $I(\cdot)$. Due to the ambiguity in matching a 3D volume with a single projected 2D image, multiple X-ray images from different projection angles are often employed in 2D/3D registration. In such cases, the goal is to find the transformation T that aligns all DRR and X-ray image pairs, denoted as $H_{i,T}(\cdot)$ and I_i , where i denotes the index of the X-ray image.

Special Euclidean Group SE(3)

Special Euclidean group SE(3) is the set of 4×4 matrices corresponding to translations and rotations. The tangent space of SE(3) is described using the Lie algebra $\mathfrak{se}(3)$, which has six generators corresponding to the derivatives of translation and rotation along/around each of the standard axes. An element of $\mathfrak{se}(3)$ is then represented by multiples of the generators

$$\delta = (\mathbf{u}, \mathbf{v}) \in \mathbb{R}^6 \quad (2)$$

$$\delta_{\times} = u_1 G_1 + u_2 G_2 + u_3 G_3 + v_1 G_4 + v_2 G_5 + v_3 G_6 \in \mathfrak{se}(3), \quad (3)$$

where (G_1, G_2, G_3) are the translation generators, and (G_4, G_5, G_6) are the rotation generators. Matrix exponential and logarithm can be taken to convert elements between SE(3) and $\mathfrak{se}(3)$.

Proposed Method

Markov Decision Process

We cast the problem of finding T to align the DRR $H_T(\cdot)$ with a X-ray image $I(\cdot)$ as a MDP, which is defined by a




Figure 2: Policy network with encoder CNN and decoder neural network (NN). An ROI centered on the projection of the agent’s location is extracted from the X-ray image and DRR. The extracted regions are encoded using CNN to obtain feature vectors, which are concatenated and decoded by a NN to obtain reward estimation.

5-tuple $\{\mathcal{T}, \mathcal{A}, P(\cdot), R(\cdot), \gamma\}$, where \mathcal{T} is the set of possible states (i.e., transformations in SE(3)), \mathcal{A} is the set of actions (i.e., modification of the transformation), $P_A(T)$ is the state obtained by taking action A in state T , $R_A(T)$ is the reward received by taking action A in state T , and γ is the discount factor that controls the importance of future rewards. With the action space \mathcal{A} and the reward scheme $R(\cdot)$ defined (details are provided in the next sections), the core problem of MDP is to find a policy $\pi(\cdot)$ that specifies the optimal action $\pi(T_t)$ to be taken at state T_t to maximize the long term reward:

$$\sum_{t=0}^{\infty} \gamma^t R_{A_t}(T_t), \quad \text{where we choose } A_t = \pi(T_t). \quad (4)$$

Action Space

To make the policy learnable, we define the action space based on the X-ray projection geometry such that each action is correlated with a specific appearance change of the DRR that is largely independent of C-arm geometry. Specifically, the transformation T is described in the image coordinate system, which has its origin at the upper left corner of the image, its (x, y) axes along the image edges, and its z axes perpendicular to the image (illustrated in Fig. 2). We further define an agent coordinate system with the same orientation as the image coordinate system and an origin that can be selected individually (details on the selection of the origin will be discussed later). The coordinate transform from the image coordinate system to the agent coordinate system is denoted as E . The transformation T_t is then describe in the agent coordinate system, written as $E \circ T_t$, and the action space is defined as small movements in the tangent space of SE(3) at $E \circ T_t$, parameterized by $\mathfrak{se}(3)$. Specifically, the action space contains 12 actions of positive and negative movements along the 6 generators of $\mathfrak{se}(3)$:

$$\mathcal{A} = \{-\lambda_1 G_1, \lambda_1 G_1, \dots, -\lambda_6 G_6, \lambda_6 G_6\}, \quad (5)$$

where λ_i is the step size for the action along the generator G_i . Application of an action $A \in \mathcal{A}$ is represented as

$$T_{t+1} = E^{-1} \circ \exp(A) \circ E \circ T_t. \quad (6)$$

Since the actions need relatively small step size in order to achieve high accuracy, we set $\lambda_{1,2,3}$ to be 1 to get a step size of 1 mm in translation, and $\lambda_{4,5,6}$ to be $\pi/180 = 0.0174$ to get a step size of 1 degree in rotation. Our definition of the action space ensures that the translation actions causes 2D shift and zooming of the DRR, and rotation actions causes rotation of object in the DRR around the agent’s origin. Therefore the image appearance change of the DRR for a given action is largely independent of the underlying C-arm geometry. Note there is no action for terminating the MDP. Instead, we run the agent for a fixed number of steps (i.e., 50 in our experiments).

Reward System

In standard MPD, the optimization target is a long term reward, i.e., an accumulation of discounted future reward, due to the difficulty of forging a reward system that directly associate the immediate reward with the long term goal. For 2D/3D registration, however, we can define a distance-based reward system such that the immediate reward is tied with the improvement of the registration. The reward scheme is defined as the reduction of distance to the ground truth transformation in the agent coordinate system:

$$R_A(T) = D(E \circ T, E \circ T_g) - D(E \circ T', E \circ T_g), \quad (7)$$

where T and T' are transformations before and after the action, T_g is the ground truth transformation. The distance metric $D(\cdot, \cdot)$ is defined as the geodesic distance of two transformations on $SE(3)$ (Hartley, Aftab, and Trumpp 2011):

$$\begin{aligned} D(T_1, T_2) &= \|\log(T_2 \circ T_1^{-1})\|_F \\ &= (2\|\mathbf{u}\|_2^2 + \|\mathbf{v}\|_2^2)^{\frac{1}{2}}, \end{aligned} \quad (8)$$

where $\log(\cdot)$ takes $T_2 \circ T_1^{-1} \in SE(3)$ into $se(3)$, \mathbf{u} and \mathbf{v} are rotation and translation coefficients of $\log(T_2 \circ T_1^{-1})$ as described in Eqn. (2). Because the units for rotation and translation are radian and mm, the distance impact of rotation is too small. Therefore, we scale the rotation coefficients \mathbf{v} by $180/\pi$ to balance the impacts of rotation and translation.

Since the distance $D(T, T_g)$ measures the distance to the ground truth, a greedy policy that maximizes the immediate reward (i.e., minimizes the distance to the ground truth) would lead to the correct registration.

$$\pi(T) = \arg \max_A R_A(T). \quad (9)$$

It can also be seen as a special case of MPD with the discount factor $\gamma = 0$.

Policy Learning via Dilated FCN

Since we seek a greedy policy that maximizes the immediate reward, we model the reward function as a deep neural network, as shown in Fig. 2, which can be learned via supervised learning. The input of the network is an observation of the current state, which consists of an observed region of fixed size (i.e. 61×61 pixels with 1.5×1.5 mm pixel spacing in our experiment) from both X-ray and DRR, referred

Table 1: Layer configurations for encoder/decoder CNNs and their equivalent dilated FCNs. Parameters for convolutional layers are written as $m \times n \times f$, where $n \times m$ is the convolution kernel size, f is the number of feature maps. sk indicates that the layer has input stride k , and dk indicates that the filter kernel is dilated k times. All convolutional layers have zero padding. SELU activation function is applied after all layers except for the input and output layers. The column "output size" specifies the output sizes for CNN.

	Layer Name	Output Size	Original CNN	Equivalent FCN
Encoder	input	61x61	-	-
	conv1	59x59	3x3x32	3x3x32,
	conv2	57x57	3x3x32	3x3x32,
	conv3	27x27	3x3x64,s2	3x3x64,d2
	conv4	25x25	3x3x64	3x3x64,d2
	conv5	11x11	3x3x128,s2	3x3x128,d4
	conv6	9x9	3x3x128	3x3x128,d4
	conv7	3x3	3x3x256,s2	3x3x256,d8
Decoder	fc1	1x1	1024	3x3x1024,d8
	fc2	1x1	1024	1x1x1024,d8
	output	1x1	128	1x1x128,d8
	input	1x1	-	-
	fc1	1x1	1024	3x3x1024,d8
	output	1x1	6	1x1x6,d8

to as Region of Interest (ROI). The ROI of an agent is centered on the projection of its origin, so that rotation actions cause rotations with respect to the center of the ROI. Such ROI center ensures that the same action always cause the same visual change in the observation, leading to a consistent mapping between the observation and rewards of different actions that can be modeled by a neural network. We propose a network architecture that encodes X-ray and DRR separately and decodes them jointly to obtain the estimated reward, as shown in Fig. 2 and Table 1.

Such CNN can be trained by supervised learning on data randomly sampled from the registration environment, similar to (Liao et al. 2017). In particular, given one pair of 2D and 3D data, training data of the CNN can be generated with 8 DoFs, including 2 DoFs for (x, y) of the agent’s origin (z is set to be the mid-point of the 3D volume) corresponding to the location of the agent, and 6 DoFs for the transformation T corresponding to the pose of the 3D volume. The ground truth rewards can be calculated following Eqn. (7). To make the learned reward function generalizable to unseen data, all 8 DoFs need to be sufficiently sampled during training, which has a prohibitively high computational cost. In (Liao et al. 2017), five million samples are needed to cover training data of 6 DoFs, and the sampling requirement grows exponentially with the DoF.

We propose a dilated FCN-based training mechanism to reduce the DoF of training data from 8 to 4. First, the encoder and decoder CNNs are converted to equivalent dilated FCNs. The filter dilation technique introduced in (Long,

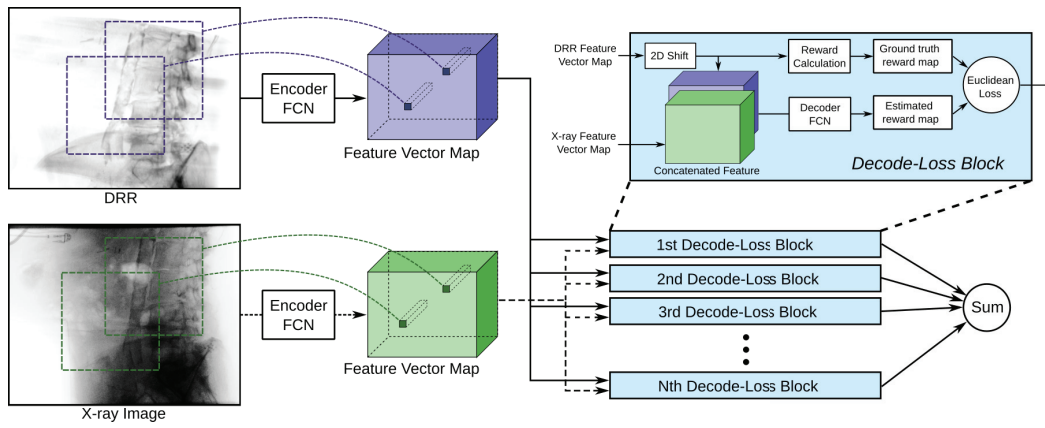


Figure 3: Training of the reward network using dilated FCN. The encoder CNN and decoder NN are converted to equivalent dilated FCNs. Densely overlapping ROIs are encoded to a dense feature vector map. The feature vector map of DRR is then randomly shifted N ($N=4$ in our experiments) times to simulate translations of the 3D volume in the imaging plane. A ground truth dense reward map is calculated from each 2D shift combined with the 3D transformation. Euclidean loss comparing the estimated and ground truth rewards is used for training.

Shelhamer, and Darrell 2015) is used to make each pixel in the FCN output exactly the same as that produced using the corresponding CNN on its receptive field in the input images. Layer to layer correspondence of CNN and dilated FCN is shown in Table 1. Training using dilated FCN is illustrated in Fig. 3. Since the encoder FCN encodes all ROIs in the input image, it can replace sampling (x, y) of the agent’s origin. The (x, y) translation part of T creates 2D shift of the DRR, which can be replaced by shifting the encoded feature map of the DRR during training. Therefore, 4 DoFs of training data can be avoided by using FCN-based training, reducing the number of total DoFs from 8 to 4. Since the complexity of the environment grows exponentially with its DoF, reducing the DoF by 4 can significantly improve the training efficiency.

In the FCN-based training, the ground truth reward map needs to be calculated during training, following Eqn. (7) and Eqn. (8). Given a X-ray/DRR pair with known T_g and T , the agent coordinate system $\{E_i\}$ is calculated for all ROIs within the input image by tracing the pixel back to the 3D volume. For each ROI, (u_i, v_i) of $\log(E_i \circ T \circ (E_i \circ T_g)^{-1})$ is pre-calculated. The effect of 2D shift of the DRR feature map is essentially adding the translation to T , which is equivalent to updating $u_i \leftarrow u_i + \Delta u$, where Δu is the 2D shift represented in mm. The distance before action is then calculated using (u_i, v_i) following Eqn. (8) for each ROI. Similarly, the distance after action is calculated following the same steps by replacing T with $T' = E_i^{-1} \circ \exp(A) \circ E_i \circ T$. Note that the expensive computation of matrix logarithms can be pre-calculated for each X-ray/DRR pair, as it does not depend on the 2D shift.

Multi-agent System

Since X-ray images during surgery and interventions can have very different Field-of-View (FoV) and contain many structures that do not match with the 3D image (e.g., medical devices), there can be many ROIs that do not contain reli-




Figure 4: Confidence map of densely overlapping ROIs within the image. Color of each pixel indicates the confidence value from the ROI centered on this pixel.

able visual cues for registration (as shown in Fig. 1). Therefore we propose a multi-agent system to provide auto attention and adaptively choose most reliable ROIs during registration. Specifically, the FCN policy network is applied on the X-ray and DRR to produce a dense reward map, which contains estimated rewards for agents with all possible ROIs from the input images, denoted as $R_i(A)$, where i is the index of the agent and $A \in \mathcal{A}$ is the action. For every agent, the maximum reward is calculated and the action associated with it is selected:

$$\begin{aligned} \hat{R}_i &= \max_{A \in \mathcal{A}} R_i(A), \\ A_i &= \arg \max_{A \in \mathcal{A}} R_i(A). \end{aligned} \quad (10)$$

Since \hat{R}_i is the anticipated reward of its selected action A_i , it represents the agent’s confidence on the action. Fig. 4 shows a strong correlation between the confidence score \hat{R}_i and the quality of the corresponding ROI: for image with a large FoV, the confidence score is high on spine (i.e., good visual cue for registration) and low on soft tissue; when severe oc-




Figure 5: *Left*: Confidence distribution of agents producing correct and wrong actions. *Right*: Relation between the percentage of correct action among selected action and the confidence threshold used for selecting action.

clusion is presented due to medical devices, the occluded area has low confidence scores.

We propose to use the confidence score \hat{R}_i to derive an attention mechanism that only selects agents with \hat{R}_i above a certain confidence threshold. To determine the threshold, the behavior of the confidence score was analyzed on validation data. Actions are categorized into correct and wrong, based on their impact on the registration (i.e., increase or decrease the distance to ground truth). Fig. 5 shows the distribution of the confidence scores from agents producing correct and wrong actions, and relation between the percentage of correction action among selected actions and the confidence threshold used for selecting actions. Based on the analysis, we selected a confidence threshold (i.e., 0.67 in our experiment) such that the correct rate of selected actions is above 95%. To avoid the scenario that too few agents are selected for a given test image, if less than 10% of the agents have a confidence score above this threshold, the top 10% agents will then be selected. After the agents are selected, denoted as \mathcal{I} , actions from the selected agents are further aggregated by L2 chordal mean to obtain the final actions:

$$\hat{A} = \arg \min_{A \in SE(3)} \sum_{i \in \mathcal{I}} \|A_i - A\|_F^2. \quad (11)$$

The L2 chordal mean can be solved globally in close form as described in (Hartley, Aftab, and Trunpf 2011).

Experiment and Result

We applied the proposed method on a clinical application of 2D/3D registration during minimally invasive spine surgery, which aims at registering spine in 3D cone beam CT (CBCT) and two X-ray images acquired from different angles. This is a challenging problem because surgical objects like screws and guide wires can be presented separately in the 3D and 2D images, creating severe image artifacts and occlusion of the target object (examples are shown in Fig. 1).

Training

During minimally invasive spine surgery, the initial pose offset between the CBCT and the X-ray images can be up to 20 mm in translation and 10 degrees in rotation. Therefore, we train the agents to perform registration starting from within




Figure 6: Comparison of training speed using CNN-based training and dilated FCN-based training.

this range. In particular, the X-ray/DRR pairs used for training have random rotation offset up to 10° , and the DRR feature map is randomly shifted up to 20 mm during training. The training data was generated from 77 CBCT data sets, where each CBCT data set consists of one CBCT and ~ 350 X-ray images used for reconstructing the CBCT. From each CBCT data set, we extracted 350 X-ray/DRR pairs. Since the number of CBCTs is limited, we also generated pairs of synthetic X-ray image and DRR from 160 CTs as additional training data, where 200 pairs were generated from each CT. In total, our training data consist of 58,950 data, i.e., 26,950 CBCT data and 32,000 synthetic data. Training was performed on a Nvidia Titan Xp GPU using pyTorch.

We compared training using CNN and the proposed dilated FCN. In the training of CNN, random ROIs are extracted from the X-ray and DRR as the CNN input, and ground truth rewards are calculated and used as supervision. Curves for the training loss and correct action rate using CNN-based and dilated FCN-based training are shown in Fig. 6. FCN-based training finished in 17 hours, with a testing loss of ~ 0.13 and a testing correct action rate of $\sim 90\%$. In comparison, CNN-based training after 17 hours reached a test loss of ~ 0.22 and a testing correct action rate of $\sim 80\%$, which is close to the performance of FCN-based method after 2 hours of training.

Testing

To evaluate the contribution of the multi-agent strategy, we tested the agent-based method in two modes: 1) using a single agent with an ROI at the center of the image, referred to as *agt-s*, and 2) using the proposed multi-agent strategy, referred to as *agt-m*. To apply the proposed method on 2D/3D registration with two X-ray images, in every step, one action is obtained from each X-ray image, and the obtained actions are applied sequentially. We also tested a combination of *agt-m* and an optimization-based method, referred to as *agt-m-opt*, where optimization of GC using BOBYQA optimizer was applied starting from the result of *agt-m*. We compared the agent-based method with state-of-the-art optimization-based methods. Multiple similarity measures were evaluated in (De Silva et al. 2016) using CMA-ES optimizer for spine 2D/3D registration, and GC and GO were reported to achieve the best performance. Therefore, we evaluated CMA-ES optimization of GO and GC using the same parameters reported in (De Silva et al. 2016), referred to as *ES-GO*

and *ES-GC*, respectively. Registration error is measured by Target Registration Error (TRE), which is calculated as the Root Mean Square Error of the locations of seven anatomical landmarks located on spine vertebrae.

CBCT Data Testing was first performed on 116 CBCT data sets via three-fold cross validation (77 used for training and 39 used for testing). The typical size of the CBCT data is $512 \times 512 \times 389$ with a pixel spacing of 0.486 mm. On each data set, 10 pairs of X-ray images that are $>60^\circ$ apart (common practice for spine surgery) were randomly selected, and 2D/3D registration was performed on each pair, starting from a perturbation of the ground truth transformation within 20 mm translation and 10° rotation, leading to 1,160 test cases. Note that X-ray images in CBCT data have a relatively low SNR with a faint spine as shown in Fig. 1.

Experiment results are summarized in Table 2. The two optimization-based methods, ES-GO and ES-GC, resulted in relatively high gross failure rate (account for $TRE > 10$ mm, which is about 1/2 of a vertebrae and considered to be grossly off by our clinical collaborating partners). This is mainly due to the low image quality (e.g., low SNR, image artifacts in CBCT and etc.), which leads to a highly non-convex optimization problem using low level similarity measures like GC and GC. In comparison, *agt-m* achieved a much lower gross failure rate, demonstrating the robustness advantage of the agent-based method. Comparison of *agt-s* and *agt-m* shows that the multi-agent strategy can noticeably improve robustness by aggregating information from most confident agents. The comparison of median TRE shows that while the agent-based method provides low failure rate, its accuracy is lower than that of optimization-based methods. This is primarily due to the discrete actions of 1 mm and 1 degree, and location information loss during stride in the CNN. By applying *opt-local* to refine the result of *agt-m*, *agt-m-opt* achieved both low failure rate and high accuracy.

Clinical Data To evaluate the proposed method in a real clinical setup, we blindly selected one trained model from the three-fold cross validation on CBCT data, and tested it on 28 clinical data sets collected from minimally invasive spine surgery. Each data set contains a CBCT acquired before the surgery and two X-ray images acquired during the surgery. Ground truth registration was manually annotated by experts. On each data set, 20 perturbations of the ground truth transformation were randomly generated as starting positions for 2D/3D registration, leading to 560 test cases.

Experiment results on clinical data are summarized in Table 2. Higher TREs are reported for all methods on clinical data than that on the CBCT data, primarily due to three reasons: 1) The ground truth registration for clinical data was manually annotated, which could bear 1~2 mm error; 2) The complexity of clinical data is much higher than the CBCT data (i.e., artifacts and occlusion caused by surgical devices, varying imaging FoVs and etc.); 3) For agent-based methods, the agent was trained without using any real clinical data from spine surgery. We observed that due to the increased complexity, the heuristically selected ROI used in *agt-s* (i.e., center of the image) become even less reliable. As a result, the robustness of *agt-s* degrades significantly com-

Table 2: Experiment results on bi-plane 2D/3D registration on 1,160 test cases from 116 CBCT data sets, and 560 test cases from 28 clinical data sets. Gross failure rate (GFR) accounts for test cases with $TRE > 10$ mm. Median, 75th percentile and 95th percentile TREs are reported.

	Method	GFR	Median (mm)	Percentile (mm)		Run Time
				75%	95%	
CBCT Data	Start	93.4%	19.4	23.2	27.8	-
	ES-GC	32.2%	1.67	22.1	44.1	18.7 s
	ES-GO	34.3%	1.81	21.0	38.6	33.9 s
	agt-s	17.2%	5.30	8.13	23.3	0.5 s
	agt-m	4.1%	3.59	5.30	8.98	2.1 s
	agt-m-opt	2.1%	1.19	1.76	4.38	2.6 s
Clinical Data	Start	95.4%	20.4	23.1	26.8	-
	ES-GC	49.3%	8.79	26.5	55.6	20.4 s
	ES-GO	42.1%	3.18	29.0	84.6	35.8 s
	agt-s	45.7%	8.42	14.2	26.4	0.6 s
	agt-m	6.8%	4.97	7.36	10.3	2.1 s
	agt-m-opt	6.1%	1.99	2.76	10.8	2.7 s

paring to that on the CBCT data. The multi-agent system, *agt-m*, in contrast achieved a much higher robustness than *agt-s*, even though the individual agent was trained without using any clinical data from spine surgery, demonstrating the effectiveness of the multi-agent strategy in dealing with complex scenarios.

Conclusion and Future Works

In this paper, we formulated 2D/3D registration as a MDP, and proposed an multi-agent system to solve this challenging problem. A multi-agent action aggregation scheme is proposed to drive the registration with inherent attention focus. In addition, a dilated FCN-based training scheme is proposed to reduce the number of DoFs that need to be sampled in training from 8 to 4, which speeds up the training efficiency by an order of magnitude. On our experiments on both CBCT data and clinical data, the proposed method is shown to be able to achieve significantly higher robustness than the state-of-the-art 2D/3D registration methods.

In our future works, we will explore the advantages provided by the proposed dilated FCN-based network structure for other general computer vision problems such as optical flow estimation and metric learning. Specifically, unlike FCN-based optical flow (e.g., FlowNet (Dosovitskiy et al. 2015)), where the pooling and uppooling layers inherently make the output insensitive to small motion, our network structure ensures that local image descriptor strictly follows the input ROI (i.e., 1 pixel shift of the ROI causes exactly 1 pixel shift of the descriptor), and therefore can potentially provide high accuracy for motion estimation. In addition, this property could be utilized for metric learning using triplet networks, i.e. embedding ROIs to feature vectors with Euclidean distances correlated with the physical distances

of the ROIs. The dilated FCN based training scheme could in general be highly efficient by allowing training multiple displacements in each back propagation. These are currently under investigation.

References

- Bushberg, J. T., and Boone, J. M. 2011. *The essential physics of medical imaging*. Lippincott Williams & Wilkins.
- Chou, C.-R.; Frederick, B.; Mageras, G.; Chang, S.; and Pizer, S. 2013. 2d/3d image registration using regression learning. *Computer Vision and Image Understanding* 117(9):1095–1106.
- De Silva, T.; Uneri, A.; Ketcha, M.; Reaungamornrat, S.; Kleinszig, G.; Vogt, S.; Aygun, N.; Lo, S.; Wolinsky, J.; and Siewerdsen, J. 2016. 3d–2d image registration for target localization in spine surgery: investigation of similarity metrics providing robustness to content mismatch. *Physics in medicine and biology* 61(8):3009.
- Deb, K.; Pratap, A.; Agarwal, S.; and Meyarivan, T. 2002. A fast and elitist multiobjective genetic algorithm: Nsga-ii. *IEEE transactions on evolutionary computation* 6(2):182–197.
- Dosovitskiy, A.; Fischer, P.; Ilg, E.; Hausser, P.; Hazirbas, C.; Golkov, V.; van der Smagt, P.; Cremers, D.; and Brox, T. 2015. FlowNet: Learning optical flow with convolutional networks. In *Proceedings of the IEEE International Conference on Computer Vision*, 2758–2766.
- Gendrin, C.; Furtado, H.; Weber, C.; Bloch, C.; Figl, M.; Pawiro, S. A.; Bergmann, H.; Stock, M.; Fichtinger, G.; Georg, D.; et al. 2012. Monitoring tumor motion by real time 2d/3d registration during radiotherapy. *Radiotherapy and oncology* 102(2):274–280.
- Gouveia, A. R.; Metz, C.; Freire, L.; Almeida, P.; and Klein, S. 2015. Registration-by-regression of coronary cta and x-ray angiography. *Computer Methods in Biomechanics and Biomedical Engineering: Imaging & Visualization* (ahead-of-print):1–13.
- Hartley, R.; Aftab, K.; and Trampf, J. 2011. L1 rotation averaging using the weiszfeld algorithm. In *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*, 3041–3048. IEEE.
- Ilg, E.; Mayer, N.; Saikia, T.; Keuper, M.; Dosovitskiy, A.; and Brox, T. 2016. FlowNet 2.0: Evolution of optical flow estimation with deep networks. *arXiv preprint arXiv:1612.01925*.
- Jaderberg, M.; Simonyan, K.; Zisserman, A.; et al. 2015. Spatial transformer networks. In *Advances in Neural Information Processing Systems*, 2017–2025.
- Knaan, D., and Joskowicz, L. 2003. Effective intensity-based 2d/3d rigid registration between fluoroscopic x-ray and ct. *Medical Image Computing and Computer-Assisted Intervention-MICCAI 2003* 351–358.
- Kruger, J., and Westermann, R. 2003. Acceleration techniques for gpu-based volume rendering. In *Proceedings of the 14th IEEE Visualization 2003 (VIS'03)*, 38. IEEE Computer Society.
- Lau, K. K., and Chung, A. C. 2006. A global optimization strategy for 3d-2d registration of vascular images. In *BMVC*, 489–498.
- Liao, R.; Miao, S.; de Tournemire, P.; Grbic, S.; Kamen, A.; Mansi, T.; and Comaniciu, D. 2017. An artificial agent for robust image registration. In *AAAI*, 4168–4175.
- Long, J.; Shelhamer, E.; and Darrell, T. 2015. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 3431–3440.
- Maes, F.; Collignon, A.; Vandermeulen, D.; Marchal, G.; and Suetens, P. 1997. Multimodality image registration by maximization of mutual information. *IEEE transactions on medical imaging* 16(2):187–198.
- Markelj, P.; Tomaževič, D.; Likar, B.; and Pernuš, F. 2012. A review of 3d/2d registration methods for image-guided interventions. *Medical image analysis* 16(3):642–661.
- Miao, S.; Liao, R.; and Zheng, Y. 2011. A hybrid method for 2-d/3-d registration between 3-d volumes and 2-d angiography for trans-catheter aortic valve implantation (tavi). In *Biomedical Imaging: From Nano to Macro, 2011 IEEE International Symposium on*, 1215–1218. IEEE.
- Miao, S.; Wang, Z. J.; and Liao, R. 2016. A cnn regression approach for real-time 2d/3d registration. *IEEE transactions on medical imaging* 35(5):1352–1363.
- Schmid, J., and Chênes, C. 2014. Segmentation of x-ray images by 3d-2d registration based on multibody physics. In *Computer Vision—ACCV 2014*. Springer. 674–687.
- Uneri, A.; De Silva, T.; Goerres, J.; Jacobson, M.; Ketcha, M.; Reaungamornrat, S.; Kleinszig, G.; Vogt, S.; Khanna, A.; Osgood, G.; et al. 2017. Intraoperative evaluation of device placement in spine surgery using known-component 3d–2d image registration. *Physics in Medicine and Biology* 62(8):3330.
- Van Laarhoven, P. J., and Aarts, E. H. 1987. Simulated annealing. In *Simulated annealing: Theory and applications*. Springer. 7–15.
- Weinzaepfel, P.; Revaud, J.; Harchaoui, Z.; and Schmid, C. 2013. Deepflow: Large displacement optical flow with deep matching. In *Proceedings of the IEEE International Conference on Computer Vision*, 1385–1392.
- Wohlhart, P., and Lepetit, V. 2015. Learning descriptors for object recognition and 3d pose estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 3109–3118.
- Wu, G.; Qi, F.; and Shen, D. 2006. Learning-based deformable registration of mr brain images. *IEEE transactions on medical imaging* 25(9):1145–1157.
- Zheng, J.; Miao, S.; and Liao, R. 2017. Learning cnns with pairwise domain adaption for real-time 6dof ultrasound transducer detection and tracking from x-ray images. *Medical Image Computing and Computer-Assisted Intervention-MICCAI 2017*.