



# Guided ultrasound acquisition for nonrigid image registration using reinforcement learning

Shaheer U. Saeed <sup>a</sup>,<sup>\*</sup> João Ramalhinho <sup>a</sup>, Nina Montaña-Brown <sup>a</sup>, Ester Bonmati <sup>a,b</sup>,  
Stephen P. Pereira <sup>c</sup>, Brian Davidson <sup>d</sup>, Matthew J. Clarkson <sup>a</sup>, Yipeng Hu <sup>a</sup>

<sup>a</sup> UCL Hawkes Institute, and Department of Medical Physics & Biomedical Engineering, University College London, London, UK

<sup>b</sup> School of Computer Science and Engineering, University of Westminster, London, UK

<sup>c</sup> Institute for Liver and Digestive Health, University College London, London, UK

<sup>d</sup> Division of Surgery & Interventional Science, University College London, London, UK

## ARTICLE INFO

### MSC:

41A05

41A10

65D05

65D17

### Keywords:

Image registration

Reinforcement learning

Surgical guidance

Liver

## ABSTRACT

We propose a guided registration method for spatially aligning a fixed preoperative image and untracked ultrasound image slices. We exploit the unique interactive and spatially heterogeneous nature of this application to develop a registration algorithm that interactively suggests and acquires ultrasound images at optimised locations (with respect to registration performance). Our framework is based on two trainable functions: (1) a deep hyper-network-based registration function, which is generalisable over varying location and deformation, and adaptable at test-time; (2) a reinforcement learning function for producing test-time estimates of image acquisition locations and adapted deformation regularisation (the latter is required due to varying acquisition locations). We evaluate our proposed method with real preoperative patient data, and simulated intraoperative data with variable field-of-view. In addition to simulation of intraoperative data, we simulate global alignment based on previous work for efficient training, and investigate probe-level guidance towards an improved deformable registration. The evaluation in a simulated environment shows statistically significant improvements in overall registration performance across a variety of metrics for our proposed method, compared to registration without acquisition guidance or adaptable deformation regularisation, and to commonly used classical iterative methods and learning-based registration. For the first time, efficacy of proactive image acquisition is demonstrated in a simulated surgical interventional registration, in contrast to most existing work addressing registration post-data-acquisition, one of the reasons we argue may have led to previously under-constrained nonrigid registration in such applications. Code: [https://github.com/s-sd/rl\\_guided\\_registration](https://github.com/s-sd/rl_guided_registration).

## 1. Introduction

Ultrasound imaging is routinely used for guiding surgery and other medical procedures, where 2-dimensional (2D) ultrasound images are continuously manually acquired by clinicians. These ultrasound images, when acquired intraoperatively, are often sparse in three aspects: (1) they are not sampled on regular grid locations as other volumetric medical images such as magnetic resonance (MR) and computed tomography (CT), with variable intervals between slices, (2) the relative positions between these images are unknown, without external spatial tracking, and (3) they only cover limited regions and with likely incomplete regions of interest (ROIs), such as missing features that are not visible using the imaging modality, or showing only a part of a structure e.g., an organ.

In this work, we focus on the development of a machine learning method that overcomes these above-discussed challenges, in the specific context of registration between ultrasound images and a 3-dimensional (3D) volumetric preoperative CT image. Current state-of-the-art methods may often suffer in terms of registration accuracy due to the sparsity of the ultrasound images and limited information available for the alignment. In particular, we argue that the interactive nature of ultrasound acquisition may be leveraged to address the above-discussed challenges. In this context then, enabling the ultrasound sampling to become an integral part of the learned registration algorithm may be the key to improve this type of registration but thus far has received limited investigation or formulation.

<sup>\*</sup> Corresponding author.

E-mail address: [shaheer.saeed.17@ucl.ac.uk](mailto:shaheer.saeed.17@ucl.ac.uk) (S.U. Saeed).

<https://doi.org/10.1016/j.media.2025.103555>

Received 1 June 2023; Received in revised form 26 June 2024; Accepted 14 March 2025

Available online 29 March 2025

1361-8415/© 2025 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

### 1.1. Volume-to-stack registration

Spatial alignment between a stack of ultrasound images and a volumetric preoperative image (such as a 3D CT image) is useful in several clinical scenarios. Examples include tracking tumour locations identified in the preoperative image, over the course of an interventional procedure guided by ultrasound images (Drakopoulos and Chrisochoides, 2016; Gooya et al., 2010; Elhawary et al., 2010), or to localise ROIs in the typically smaller volumes during intraoperative ultrasound imaging with respect to the larger field-of-view (FOV) acquired preoperatively, for intraoperative navigation (Hu et al., 2012; Amin et al., 2003; Wein et al., 2008; Penney et al., 2004).

A volumetric image is defined to be a 3D image with voxel locations defined in a single image coordinate system, here, the preoperative CT image, with its intensity values sampled from regular grid locations. In this case, if the ultrasound images can be reconstructed into volumetric images, the registration becomes a *volume-to-volume* registration problem. The 3D reconstruction of freehand ultrasound images, with or without external trackers (Tomaževič et al., 2005; Heldmann et al., 2010; Heiselman et al., 2020), may suffer from problems such as the need to optimise reconstruction parameters or to deal with incomplete information. Existing image registration methods have focused on volume-to-volume registration (Villalon et al., 2011; Thomson et al., 2020), lately with machine learning-based algorithms for improved efficiency and accuracy (Chen et al., 2021). Most methods remain specific to the application of interest and a general solution applicable to a wider array of problems is still an active area of research.

The *volume-to-slice* registration problem has also been investigated, due in part to the unknown relative locations between 2D images, where each 2D ultrasound image is aligned with the volumetric image individually and often independently (Ferrante and Paragios, 2017). The ultrasound image coordinates and the volume-to-slice transformation are both defined in 3D spatial coordinate systems, thus the 2D ultrasound images are referred to as *image slices*.

Formulating as a machine learning method, the volume-to-slice registration may benefit from multiple neighbouring images as input, as well as jointly predicting transformations between multiple slices (2D images) and the volumetric image (Ferrante and Paragios, 2017). This is termed *volume-to-stack* registration in this paper for clarity, where a number of consecutively acquired images is referred to as a *stack* of slices. The difference between the volume-to-stack registration and the 3D-reconstruction-first methods is whether the relative locations between 2D images are independently estimated beforehand (reconstruction-first) or are implicitly handled (volume-to-stack).

### 1.2. Under-constrained estimation of deformation

Whilst the volume-to-volume registration is in general an ill-posed problem for estimating nonrigid deformation, it is the added need to estimate 3D transformation between individual slices and the volumetric image that makes the volume-to-stack registration arguably even less constrained. This lack of constraint has motivated most existing studies to employ highly constrained transformation models, such as the rigid transformation assumption with six degrees-of-freedom (DOF), which is commonly employed for registration of intraoperative to preoperative data. A typical example is our application of liver ultrasound to CT registration in this work, where rigid registration has been first applied for aligning intraoperative ultrasound slices and a preoperative CT image (Wei et al., 2021b; Wein et al., 2008; Sun et al., 2018; Ramalhinho et al., 2020).

In addition to the coordinate system conversion between slices and the volumetric image, which can be modelled as a constrained rigid or affine transformation, organ motion and other soft tissue deformation needs to be compensated for using higher order transformation models. Nonrigid registration has been proposed for this application (Elhawary et al., 2010; Lee et al., 2010; Lange et al., 2003). Partly due to the

under-constrained nature, an important type of *registration parameter* is used in non-rigid methods: the weighting on a deformation regularisation (such as a non-smoothness penalty term  $L^2$ -norm of displacement gradient or bending energy (Haskins et al., 2020; Fu et al., 2020b; Chen et al., 2021), encouraging smooth (thus constrained) local tissue deformation). However, this weight has often been fixed and assumed the same everywhere, despite the fact that the material properties such as stiffness vary at different sections of the liver, and between different anatomical (such as vessels) and pathological (such as tumours) structures. Similar to this, the early attempts at using bio-mechanical models as a prior-constrained transformation model only investigate homogeneous elasticity (Özgür et al., 2018).

It is however clear in many reported results that the registration performance varies substantially between different regions within the liver (Ramalhinho et al., 2020). Spatially adaptive deformation regularisation may allow a more flexible transformation model for this application, albeit potentially worsen its ill-posedness.

Outside of the volume-to-slice registration, spatially adaptive deformation regularisation may be inferred through data-driven approaches (Simpson et al., 2015; Papież et al., 2013; Risholm et al., 2010; Hermosillo et al., 2002). For learning-based registration algorithms, these registration parameters become hyperparameters and can be reparameterised by hyper-networks (Hoopes et al., 2021) for efficient test-time adaptation. Other adaptable registration methods have also been proposed in forms of meta learning (Park et al., 2022; Baum et al., 2022). These require a trained registration model to be fine-tuned at test-time, based on a few examples. For instance, Baum et al. (2022) formulate the problem in an interactive setup where during fine-tuning, the registration may be manually adjusted by user-added data.

### 1.3. Unequal ultrasound sampling locations

In addition to local deformation, imaging content including anatomical features also varies substantially at different imaging locations. It is another likely cause for the spatial dependency of registration performance. For example, in a vessel-based registration algorithm, the complexity and richness of the vascular structures contained in the acquired slices is an important positive contributor towards a successful registration (Ramalhinho et al., 2020, 2018, 2022).

In this essay, we consider the sampling locations of all slices to be parameterised by (1) the starting location of the stack (with respect to certain reference) and (2) the physical length of the stack (e.g. in mm), which assumes a linear acquisition path along the organ surface. This approximately linear acquisition path is applicable for a wide range of contact-required ultrasound imaging in clinical applications, such as invasive liver surgery, transrectal prostate intervention, fetal anomaly emanation and musculo-skeletal imaging. The stack length can therefore be considered as a registration parameter which controls the balance between spatial coverage and accurate initial localisation (to be followed by local non-rigid registration), i.e. a longer stack should cover more characteristic features for registration, but may require a more complex nonrigid transformation to be estimated, while as the stack length approaches zero, between-slice transformation is no longer needed.

The above observations have motivated one of the main contributions of this work: development of an interactive registration that allows sequential suggestions of ultrasound sampling locations and, from these suggested locations, acquiring subsequent ultrasound slices. The sequence of image sampling actions should in turn provide image slice data, for a better-posed, easier-to-register volume-to-stack registration.

#### 1.4. Guiding image acquisition for registration

In summary, this work proposes a new framework that optimises a sequence of *a*) ultrasound image stack sampling locations and *b*) spatially-varying deformation regularisation, for challenging volume-to-stack, intraoperative registration tasks.

The registration task is also a part-to-whole registration due to its sparsity in corresponding features and dependency in ultrasound imaging orientation and depth. It may be interesting to highlight the similarity between our proposed work, which involves both registration and localisation, and methods in simultaneous localisation and mapping (SLAM) literature (Stachniss et al., 2016; Julier and Uhlmann, 2001; Davison, 2003; Salas-Moreno et al., 2013). In this work, localising and visualising the ultrasound probe with respect to reference anatomical landmarks (optionally as a visual intraoperative aid, identifiable in the volumetric CT), may be considered as the mapping component in a SLAM formulation. This perspective of a SLAM algorithm is not formalised further, but consistent notations are adopted when possible.

We propose a guided intraoperative image acquisition framework for acquiring US sweeps suitable for registration with preoperative CT. As opposed to previous works that usually refer to guidance that fine-tunes registration for already-acquired images (Baum et al., 2022; Park et al., 2022), our work guides image acquisition towards regions which improve registration performance. Furthermore, to enable this novel application we derive robust reward signals from hyper-networks, which enables RL controller training informed by task-performance, an under-explored area in RL.

The overall contributions are summarised as follows.

- (1) We formulate a new interactive volume-to-stack registration algorithm, between sequentially acquired stacks of ultrasound slices and preoperative CT.
- (2) We propose a hyper-network-based registration model which is generalisable over different regions and adaptable at test-time with respect to three conditioning variables i.e., a parameter that controls deformation regularisation, a stack length parameter, and a parameter to indicate the starting location of the stack (details in Section 2.1);
- (3) We propose a reinforcement learning-based optimisation for producing test-time estimates of these conditioning variables, for guiding ultrasound stack acquisition for registration purpose;
- (4) We evaluate our proposed framework using clinical preoperative data derived from CT scans of patients undergoing laparoscopic liver staging and resection, and simulated intraoperative ultrasound data, and compare it to other registration approaches common in the literature.

In Section 3, we argue that the use of simulated intraoperative data provides practically feasible ground-truth for validation of deformable models in surgery, currently not possible with real clinical data. It also serves as necessary first experimental evidence supporting further interactive data acquisition, due to the nature of the surgical application and proposed guided acquisition. Additionally, the use of simulated intraoperative data also allows us to report informative validation measures such as target registration error for corresponding landmarks in intra- and preoperative images, where this may not be possible on real intraoperative data. The evaluation in a simulated environment, with simulated intraoperative data and global alignment based on error estimates from previous work, allows us to demonstrate a novel image-acquisition guidance framework to guide acquisition towards improved registration.

## 2. Methods

The proposed guidance framework relies on two functions: (1) a hyper-network-based registration function to allow adaptive registration; and (2) a reinforcement learning function to guide the operator




Fig. 1. An overview of the proposed registration and guidance scheme being used intraoperatively. The guided image stack sampling is the effector  $e$  described in Section 2.3.1. The transformations between slices as well as between slices and the preoperative volume for this challenging volume-to-stack registration, are unknown.

towards regions that aid registration. The hyper-network based registration function allows varying conditioning variables at test-time i.e., varying stack length or regularisation weight and can accommodate different sampling locations within the liver. These conditioning variables and sampling locations are suggested by the reinforcement learning function, at test-time. The reinforcement learning function may be trained using a reward signal derived only from the performance gain/loss of the registration network subsequent to each suggestion. Once trained, the RL function guides the operator in terms of the conditioning variables (including a prescribed stack length, regularisation weight and sampling location) towards an optimal registration and the hyper-network based registration function provides the registration using these conditioning variables.

### 2.1. Overview

Our proposed framework involves two trainable functions, (1) the hyper-network (hypernet) based registration function  $f_{reg}$ ; and (2) the reinforcement learning (RL) function  $f_{rl}$ , both parameterised by neural networks with network parameters  $\theta^{reg}$  and  $\theta^{rl}$ , respectively.

At a given time step  $t$ , a stack of  $n$  ultrasound slices, with  $m$  number of pixels in each slice,  $x_t \in \mathcal{X} \subset \mathbb{R}^{m \times n}$  is available, together with a pre-operative volumetric CT image  $y \in \mathcal{Y}$ , acquired prior to the procedure. Whilst  $x_t$  and  $y$  are in general vectors of per-pixel image features, such as intensity values or binary valued morphological features (i.e. ROI segmentation masks), in the respective spaces  $\mathcal{X}$  and  $\mathcal{Y}$  (that may be considered as representing their respective image coordinate systems).

To localise the slices, the registration function  $f_{reg} : \mathcal{X} \times \mathcal{Y} \rightarrow \mathcal{D}$  computes a set of displacement vectors  $d_t \in \mathcal{D} \subset \mathbb{R}^{m \times n \times 3}$  for all slices in the stack:

$$d_t = f_{reg}(x_t, y; \sim) \quad (1)$$

such that  $\hat{y}_t = x_t \otimes d_t$  denotes the warped (or registered)  $x_t$  in  $\mathcal{Y}$  coordinate system using displacements  $d_t$ . Eq. (1) represents a static registration solution at time step  $t$  for the acquired stack.

As discussed in Sections 1.2 and 1.3, the registration may benefit from sequentially optimising three conditioning variables, at time  $t$ :

- A starting location  $a_t^{loc} \in \mathbb{R}^3$  of the stack, indicates the relative orientation and distance, from the last slice in the previous stack to the first slice in the current stack, for acquisition of the stack.
- A stack length  $a_t^{len}$ , a scalar in physical length and independent of the number of slices sampled in a stack.
- A positive weight  $a_t^{def}$  on the deformation regularisation, displacement gradient norm in this work; controlling the spatial regularity of deformation.

The first two collectively specify the ultrasound image sampling location and all three are denoted as  $a_t = \{a_t^{loc}, a_t^{len}, a_t^{def}\} \in \mathcal{A}$ . Thus, the action-conditioned registration function becomes:

$$d_t = f_{\theta^{reg}}(x_t, y; a_t) \quad (2)$$

The RL function  $f_{\theta^{rl}} : \mathcal{X} \times \mathcal{Y} \rightarrow \mathcal{A}$  is trained to suggest values for these conditioning variables, to be used at the next time step.

$$a_{t+1} = f_{\theta^{rl}}(x_t, y) \quad (3)$$

where  $a_t$  can be considered as actions suggested by a RL agent, which then interacts with a Markov decision process (MDP) environment, which then determines the ultrasound image stack  $x_{t+1}$  sampled in the next time step.

An overview is presented in Fig. 1. In the following sections, we provide details on how each of the registration function and the RL function can be formulated and optimised in this application.

## 2.2. The registration function

In our formulation the registration consists of two parts, the global alignment and the non-rigid registration. The global alignment followed by non-rigid registration provide a deformable registration, in our work. The non-rigid registration is made adaptable at test-time, with respect to three variables (sampling location for the intraoperative image, stack length of an acquired intraoperative image sweep, and the deformation regularisation weight). This is done by means of a hypernetwork, which takes as input the controlling/conditioning variables and modifies the parameters of the non-rigid registration network, in order to modify it such that it can accommodate the new conditioning variable setting. In this section, we describe each of the components involved in the registration in our framework.

### 2.2.1. Registering with rigid global pre-alignment

In our proposed method, the registration function consists of two main components, a global alignment algorithm and a nonrigid registration network for predicting nonrigid local displacements. The global alignment algorithm is not implemented in this work and is simulated based on errors reported in Ramalhinho et al. (2020).

The global alignment algorithm serves two purposes: (1) to obtain an initial rigid alignment between  $x_t$  and  $y$  and (2) to resample image features  $z_t \in \mathcal{Z}$  from a latent (more compact)  $\mathcal{Y}$  - an intermediate coordinate system:

$$z_t = y \otimes (d_t^{global})^{-1} \quad (4)$$

where

$$d_t^{global} = g(x_t, y) \quad (5)$$

and  $g : \mathcal{X} \times \mathcal{Y} \rightarrow \mathcal{D}$  is a rigid registration that computes displacement vectors  $d_t^{global} \in \mathcal{D}$ . The global alignment function with its inputs and outputs is visualised in Fig. 2. While the rigidly-registered  $x_t$  is denoted as  $x_t \otimes d_t^{global}$ , the inverse rigid transformation  $(d_t^{global})^{-1}$  is analytically computed for resampling  $y$  image features in  $\mathcal{Z}$  (Eq. (4)). Here, the latent  $z_t$  offers a substantially compact representation and a focused FOV section of  $y$ , which is localised with respect to  $x_t$ , thus allowing non-rigid registration between  $x_t$  and  $z_t$  much more efficiently in training as well as during test-time.

The nonrigid registration network aligns the  $x_t$  and  $z_t$  in the latent  $\mathcal{Z}$  coordinates, using a neural network  $f_{\theta^{local}}$  with network parameters  $\theta^{local}$ :

$$d_t^{local} = f_{\theta^{local}}(x_t, z_t; a_t) \quad (6)$$

where  $d_t^{local} \in \mathcal{D}$  is a set of displacement vectors describing local nonrigid transformation, such that the registered  $x_t$  in  $\mathcal{Z}$  is denoted as:

$$\hat{z}_t = x_t \otimes d_t^{local} \quad (7)$$




Fig. 2. An overview of the global alignment function, which has been simulated in this work based on errors reported in previous works (Ramalhinho et al., 2022, 2020). Grey: non-parametric fixed function; Dotted bubble: expanded visualisation.

This two-step registration in Eqs. (5) and (6) is equivalent to the problem described in Eq. (2).

Once registered,  $x_t$  can be aligned with  $y$ :

$$\begin{aligned} \hat{y}_t &= x_t \otimes d_t \\ \Rightarrow \hat{y}_t &= x_t \otimes d_t^{local} \otimes d_t^{global} \end{aligned} \quad (8)$$

In general, the global registration algorithm  $g$  can also be a learnable machine learning model, while several global registration methods (Robu et al., 2018; Ramalhinho et al., 2020; Lange et al., 2009) have been proposed for this application. In this work one of our previously-proposed methods (Ramalhinho et al., 2020) was simulated as an example (i.e., reported errors of  $\pm 12$  mm and  $\pm 40^\circ$  used to obtain the alignment in simulation).

### 2.2.2. Training nonrigid local registration networks

In addition to most existing learning-based registration methods, the proposed local nonrigid registration network  $f_{\theta^{local}}(x_t, z_t; a_t)$  aims to adapt the three types of conditioning variables  $a_t$  at test time. A hypernet  $\theta^{local} = h_{\theta^{hyp}}(a_t^{len}, a_t^{def})$  can re-parameterise Eq. (6) into a new network  $f_{\theta^{hyp}}$  (visualised within the whole pipeline in Fig. 1 and on its own, with inputs, outputs and training scheme in Fig. 3). It combines  $h_{\theta^{hyp}}$  and  $f_{\theta^{local}}$ , with the only learnable parameters being  $\theta^{hyp}$ :

$$\begin{aligned} d_t^{local} &= f_{\theta^{hyp}}(x_t, z_t; a_t) \\ &= f_{\theta^{local}}(x_t, z_t; a_t), \text{ where } \theta^{local} = h_{\theta^{hyp}}(a_t^{len}, a_t^{def}) \end{aligned} \quad (9)$$

This enables efficient inference with registration parameters, variables  $a_t^{len}$  and  $a_t^{def}$ . Together with the other conditioning variable, the stack starting location  $a_t^{loc}$ , they are optimised, with respect to the reward, in the subsequent RL function training (Section 2.3), in which  $a_t^{len}$  and  $a_t^{loc}$  collectively determine the sampling location of  $x_t$  (therefore conditioning the registration).

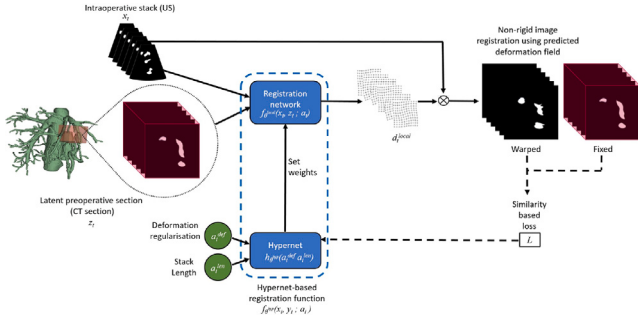
In this application, the hypernet-based local registration network is optimised in the latent coordinate space  $\mathcal{Z}$ , by minimising the distance between the resampled image features  $z_t$  and the warped image  $\hat{z}_t$  in Eq. (7):

$$\begin{aligned} L_{\theta^{hyp}}(z_t, x_t, d_t^{local}) \\ = (1 - \lambda)L^{dice}(z_t, x_t \otimes d_t^{local}) + \lambda L^{def}(d_t^{local}) \end{aligned} \quad (10)$$

where  $L^{dice}$  and  $L^{def}$  are the Dice loss between two binary masks and the  $L^2$ -norm of the displacement gradient in 3D, with the additional parameter  $\lambda = a_t^{def}$  controlling the weight between the two terms. Training the registration networks is therefore a hypernet optimisation:

$$\hat{\theta}^{hyp} = \arg \min_{\theta^{hyp}} \mathbb{E}_t [L_{\theta^{hyp}}(z_t, x_t, d_t^{local})] \quad (11)$$

Although the RL function uses an episodic training methodology which reflects the intended sequential interactions, the registration network



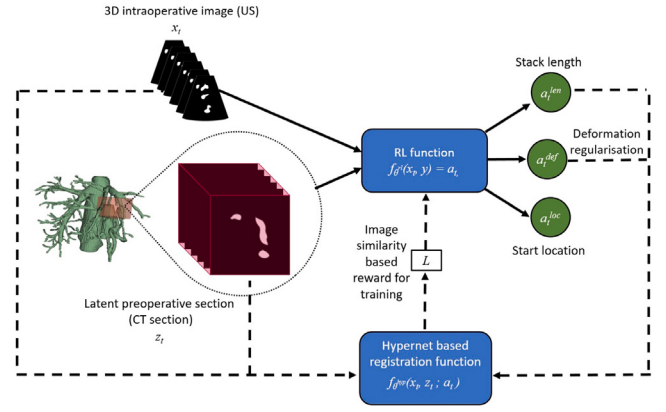
**Fig. 3.** Hypernet based registration network. Blue: parametric trainable function; Blue dashed bubble: combination of functions; Dotted bubble: expanded visualisation; Green: variables; Dashed line: used for function training. The hypernet re-parameterises the weights of the registration network based on the input conditioning variables, this weight-modification adapts the registration network to represent the new conditioning variable setting.

was trained by randomly sampling  $t$  with respect to the initial global alignment. The trained registration network was found sufficiently accurate to be used in the RL training and further details are discussed in Section 2.3. Although the stack length is not explicitly included in the objective, it is implicitly included as the 3D intraoperative image  $x_t$  varies with varying stack length. Its explicit inclusion as an input to the hypernet ensures that this dependence is explicitly learned for allowing test-time variability of the registration network based on this variable. In practice, stack length can be discretised at test-time (e.g., short, medium, long) as exact measurements are not required for the functioning of the hypernet, which can accommodate deviations from RL-suggested condition variables as described in Section 3.2.

Segmentation masks are proposed here as the image features for  $x_t$  and  $z_t$ , from both intraoperative ultrasound and preoperative CT images, for two practical considerations: (1) there have been previously proposed segmentation methods for both CT (Saeed et al., 2024; Wei et al., 2021a; Ciecholewski and Kassjański, 2021; Zeng et al., 2016; Gao et al., 1996) and intraoperative ultrasound slices (Montaña-Brown et al., 2021; Porter et al., 2001; Song et al., 2015; Nam et al., 2012) for automating this step; and (2) segmentation provides relatively consistent features between those from images with simulated deformation and those from interactively-acquired data (e.g. from a future prospective clinical validation study), such that the simulated data may be used as training data for or assisting the training of the proposed models. However, other image features or the unprocessed intensity values, using an unsupervised loss, could also be useful for training the registration network (Hoffmann et al., 2021). As an example, learned label maps could guide registration as opposed to pre-selected label maps (Hoffmann et al., 2021), however, this is beyond the scope of this work where we use liver vessels as features for registration. This is because vessels are one of the only visible features in intraoperative imaging in this application, with most previous works adopting similar approaches for feature-based registration (Ramalhinho et al., 2020; Wei et al., 2020; Nam et al., 2012), as vessels are arguably the only corresponding features for the purpose of robust registration. Furthermore, previous work has explored the features required for a successful registration in liver intraoperative to preoperative imaging and concluded that vessels and their bifurcations (especially  $>1$ ) offer unique descriptions of the regions for registration (Ramalhinho et al., 2018; Song et al., 2015; Wei et al., 2020; Nam et al., 2012).

### 2.3. The reinforcement learning function for test-time estimation of optimal conditioning variables

The RL function in our formulation provides suggestions for conditioning variables (stack sampling location, deformation regularisation, and stack length) that lead towards an optimal registration. This



**Fig. 4.** Reinforcement learning function. Blue: parametric trainable function; Dotted bubble: expanded visualisation; Green: variables; Dashed lines: used for function training. The RL function suggests conditioning variable settings that guide the operator towards an optimal registration. The function is trained using a reward signal derived from the registration network, quantifying performance gains/losses subsequent to any conditioning variables suggestions by the RL function.

function is trained solely using a reward derived from performance gains/losses quantified for the registration network subsequent to a suggestion made by the RL function.

The RL function  $f_{\theta^{rl}}$  in Eq. (3) is so-called because it is trained using a RL reward signal based on a task performance directly generated using the fixed and pre-trained registration function. This function is visualised within the entire framework in Fig. 1 and showing its inputs, outputs and training scheme in Fig. 4. Conceptually similar to previous task-based reward signals in RL (Saeed et al., 2022, 2021; Yoon et al., 2020), the three types of conditioning variables  $a_t$  are predicted by the RL function and optimised during registering samples representative of the kinds of deformation cases encountered during surgery (see Section 3.1 and Appendix for details). The time-step  $t$  indicates each time a new image stack  $x_t$  is sampled and a new feature  $z_t$  is resampled, during a sequential  $a_t$ -predicting process modelled within a MDP environment, as described as follows.

#### 2.3.1. The interactive MDP environment

**The policy and RL functions.** In a policy gradient RL algorithm, we wish to learn a parametric RL policy function  $\pi_{\theta^{rl}}(a_t|s_t) : S \times \mathcal{A} \in [0, 1]$  that predicts the probability of action  $a_t$  given observed state  $s_t \in S$ . The actions with highest probability can then be suggested by the RL function, at time step  $t$ :

$$\hat{a}_t = f_{\theta^{rl}}(x_t, y) = \arg \max_{a_t \in \mathcal{A}} \pi_{\theta^{rl}}(a_t|s_t) \quad (12)$$

It should also be noted that it may be practically beneficial to have the previous sampling location  $a_{t-1}^{loc}$  as an additional input to this function, in spite of the MDP assumption, which is omitted here for notational convenience.

**Observed states, actions and the effector.** The observable state at time step  $t$  in the MDP environment includes the intraoperative image stack  $x_t$  and the preoperative image  $y$  (represented with observed  $z_t$ ).

In this application, the three types of conditioning variables  $a_t$  are considered as actions.  $a_t^{len}$  and  $a_t^{def}$  are modelled directly using the hypernet described in Section 2.2.2.  $a_t^{len}$  and  $a_t^{loc}$  determine where to sample the intraoperative ultrasound image stack  $x_t$  in the next time step.

The  $x_t$  sampling is performed by an effector  $e : \mathcal{A} \rightarrow \mathcal{X}$

$$x_{t+1} = e(a_t) \quad (13)$$

where the effector  $e$  is commonly denoted in RL algorithms for an “external agent”. In offline training strategies, the effector generates the




Fig. 5. Effector. Acquisition of the next intraoperative stack  $x_{t+1}$ , given conditioning variables. Grey: non-parametric fixed function; Green: variables. The effector takes in the guidance sampled from the RL function and uses it to acquire a new stack/sweep of US. This may be a human following the RL guidance intraoperatively, where the guidance is delivered to the human operator via a constructed map (see Section 2.4).

retrospectively recorded or simulated interaction data. In other cases, such as the use of the proposed method, a human observer or a robotic system becomes the effector following RL-function-suggested actions (although this guidance may be ignored or not followed exactly, as described in Section 3.2). The effector is visualised showing inputs and outputs in Fig. 5. In general, the effector executes an action to influence the environment and this makes the system non-differentiable. This also motivates the proposed RL function optimisation, as opposed to for example a direct gradient descent optimisation.

**The partially observable environment.** To summarise, the proposed MDP environment as a whole encompasses the following observed data:

- The actions  $\{a_t^{loc}, a_t^{len}, a_t^{def}\}$ , suggested by the RL function  $f_{\theta^{rl}}$ .
- The intraoperative image stack  $x_t$ , sampled by the effector  $e$ , using the suggested  $\{a_t^{loc}, a_t^{len}\}$ .
- The latent features  $z_t$ , resampled by the fixed global alignment algorithm  $g$ , from  $y$  with respect to  $x_t$ .
- The transformation  $\hat{d}_t^{local}$  between  $x_t$  and  $z_t$ , estimated by the hypernet-based local registration network  $f_{\theta^{local}}(x_t, z_t; a_t)$ , using the suggested  $a_t^{def}$ .

### 2.3.2. Training with a registration performance reward

After the registration, the registration performance in Eq. (10) can be used to construct the reward function  $r : S \times \mathcal{A} \rightarrow \mathcal{R}$ , to train the RL function  $f_{\theta^{rl}}$ , i.e. optimising parameters  $\theta^{rl}$ . The scalar reward  $R_t \in \mathcal{R} \subset \mathbb{R}$  represents the improvement in registration performance between consecutive time steps:

$$R_t = r(s_t, a_t) = L_{\theta^{hyp}}(z_t, x_t, \hat{d}_t^{local}) - L_{\theta^{hyp}}(z_{t-1}, x_{t-1}, \hat{d}_{t-1}^{local}) \quad (14)$$

Given a sequence of observed states, actions and rewards  $\{s_1, a_1, R_1, s_2, a_2, R_2, \dots, s_t, a_t, R_t, \dots\}$ , the optimal RL function parameters  $\hat{\theta}^{rl}$  may be obtained by maximising the accumulated reward:

$$\hat{\theta}^{rl} = \arg \max_{\theta^{rl}} \mathbb{E}_{\pi_{\theta^{rl}}} \left[ \sum_k \gamma^k R_{t+k} \right] \quad (15)$$

where  $\gamma \in [0, 1]$  is the discount factor for future rewards. Further details of the interactive training procedure are summarised in Algo. 1 and Fig. 1.

### 2.4. Optional map construction as an intraoperative aid

The sequence of multiple intraoperative image stacks  $x_t$  in previous sections are samples representing patient anatomy (also termed ‘environment’ in SLAM literature), which contains anatomical landmarks.

Some of these landmarks are of limited spatial variation, such that constructing a map of them may be useful for a visual reference to aid intraoperative actions, especially when the human effector is trying to follow the suggested actions, e.g. where to sample the next image stack and how far to move the ultrasound probe. This serves as a mechanism to deliver the intraoperative RL-sampled guidance in a human-readable format to the operator.

In this work, we propose to construct an optional map using point landmarks for aiding the intraoperative image acquisition in this interactive registration application. The landmark map, represented by a point cloud  $p^y \in P$ , is reconstructed from a set of  $U$  image stacks  $[x_u]_{u=1, \dots, U}$ . The map construction  $l : \mathcal{X}^U \rightarrow P$  is based on a reference map, a predefined landmark point cloud  $\tilde{p}^y \in P$  in preoperative image  $y$ , including the following steps:

1. Obtain globally-aligned image stacks  $z_u$  (Eq. (4)).
2. Detect available landmarks  $p_u^z$  in each image stack.
3. Estimate  $\hat{p}_u^y$  with respect to the reference map, by registering each  $p_u^z$  to  $\tilde{p}^y$  using an ICP (iterative closest point) algorithm.
4. Update the map  $\hat{p}^y = \frac{1}{U} \sum_u \hat{p}_u^y$ , by averaging.

The landmark detection may be either done manually e.g., by conducting a binary classification of landmark presence within an acquired stack or automatically using previously proposed automatic segmentation or localisation methods (Menteş and Bilge, 2017; Montaña-Brown et al., 2021). Thus, the effector defined in Eq. (13) can be conditioned:  $x_{t+1} = e(a_t | \hat{p}^y)$ .

Other visual cues, such as the ultrasound probe location, known in the original image stacks, can also be localised in the updated map to aid navigation. We also propose to use a 2D map for the navigation purpose, by projecting the computed landmark map  $\hat{p}^y$  onto plane overlaid with the preoperative volume rendering. A coronal plane example is shown in Fig. 6.

There are alternative mechanisms to sample point clouds from multiple image stacks in simultaneous localisation and mapping literature, such as landmark localisation using a nonrigid mapping and updating the registration using the localised landmarks. Here, we adopt a simple approach as a proof of concept i.e. using the global registration function  $g$  to localise the landmarks based on the images independently. Other intraoperative visualisation techniques may be used for this registration framework, which together with how intraoperative guidance visualisation impacts surgical procedures, may be beyond the scope of this work.

## 3. Experiments

In general, developing the registration network (described in Section 2.2) requires paired CT and ultrasound images, while sufficiently sampled ultrasound stacks from individual patients may also be required to allow the registration function to generalise to various acquisition locations, as well as for training the interactive RL function (Section 2.3). Both data requirements are limited by the prevalence of the surgical procedures and available clinical data are generally considered insufficient to train large interactive networks. Therefore, this work tested the networks that are trained using preoperative images from clinical patients, with simulated intraoperative ultrasound image masks (Ramalhinho et al., 2023). Using ROI segmentation masks as image-representing features for  $x_t$  and  $z_t$ , proposed in Section 2.2.2, was also motivated for maximising the generalisation ability for the partial-simulation-trained networks. It would be interesting to investigate methods to utilise available paired data for improving the training, but they are not considered in this work.

Validation of the proposed methods, however, cannot sidestep prospective studies, in which interactions should take place, i.e. the trained neural networks are applied on the ultrasound images which are acquired online or in real-time, while both the ultrasound acquisition




Fig. 6. A 2D point cloud map of landmarks constructed as an intraoperative aid (projected into 2D coronal view and overlaid onto a 2D projection of the preoperative volume). Intraoperative points sampled using global alignment function and correction done as described in Section 2.4. hv: hepatic vein (1, 2, 3 corresponds to right, middle and left); pv: portal vein branching point; gb: gall bladder; ao: aorta; vc: vena cava. Sampling guidance arrow comprises of  $a_i^{loc}$  and shows the direction and magnitude of suggested direction.

and registration are guided by the network-suggested actions. In this work, we report a set of validation experiments, based on the simulated intraoperative data. It is necessary for two reasons: (1) Testing the proposed machine learning methodology using consistent training and test data, to demonstrate the efficacy in controlled experiments for further methodological development; and (2) Providing an estimate of the effect size in registration improvement due to the added interactive ultrasound acquisition, for designing and planning a future clinical validation study.

### 3.1. Data and simulated deformation

We use patient data to evaluate our framework, where the preoperative 3D volume used is the CT volume acquired prior to a laparoscopic liver resection or laparoscopic liver staging procedure, (with liver vessels segmented to create binary masks using a commercial service (Anon, 2023)). All data involving human participants was acquired in accordance with the ethical standards of the institutional and/or national research committee and with the 1964 Helsinki declaration and its later amendments or comparable ethical standards, under the NIHR grant [II-LA-1116-20005] with codes REC = 14/LO/1264, IRAS = 158321. We used 10 CT volumes in total, divided into the train, validation and holdout sets each with 4, 2 and 4 CT volumes. In this work the size of  $y$ , the preoperative volume, is  $512 \times 512 \times 512$  voxels, where the 20 voxels correspond to 1 cm (i.e., 0.5 mm voxel size in all dimensions).

The intraoperative image stacks are simulated by first deforming the 3D CT volumes and sampling simulated US masks from these using US simulation as in Ramalhinho et al. (2023, 2020). To generate intraoperative volumes from which we sample intraoperative image stacks, we deform the 3D CT volumes using random deformation simulating insufflation by applying compression to the volumes (in the coronal plane) and breathing by applying compression or expansion to the volumes (in the coronal or transverse/axial planes) and postural deformations by small random deformations in all three planes. Insufflation, breathing-based deformations and postural deformation compared to the preoperative volume are common in laparoscopic surgery which is why we simulate these deformations. We generate 1000 such deformed cases for each of the patients (further details in the Appendix).

Random deformations (as described above, with further details in the Appendix) are applied to this volume to obtain an intraoperative

volume of the same size as the preoperative volume, for sampling intraoperative simulated US images. The intraoperative US stacks,  $x_i$ , are formed of each slice being  $72 \times 72$  pixels (20 pixels correspond to 1 cm). The stack length is between 2 to 10 cm. Each US stack is formed over 24 slices over this stack length. The  $z_i$  sampled preoperative section after global alignment consists of stacks formed of slices with  $96 \times 96$  pixels, with 40 slices in each stack. The stack length was 2 cm larger than the corresponding US stack  $x_i$ . It is noteworthy that the sampled latent CT sections were designed to be larger than the US stacks in order to accommodate any likely insufficient overlap following the global alignment. The dimensions for the images were configured empirically with realistic sizes (sampled using 24 slices) over representative stack lengths (2-10 cm) where these were set with input from expert clinicians with approximately 20 years experience with laparoscopic and endoscopic US.

### 3.2. Interactive data sampling for training and validation

During training, the effector was modelled in simulation by following the stack start location guidance  $a_i^{loc}$  and the stack length parameter  $a_i^{len}$ . In the simulator, the ultrasound image acquisition path is determined by a) the start location - a vector specifying the relative spatial location from the slice centre from the last stack  $x_{i-1}$ , b) a spline fit to the deformed liver surface from the starting location towards the tangential direction  $\frac{a_i^{loc}}{\|a_i^{loc}\|_2}$ , and c) the image slices are acquired along the spline until the suggested stack length is reached, defined as the distance between slice centres.

It is interesting to note that the above-described effector is one example of how such sequential suggestions are followed and actions are executed. Large variance in its behaviour would increase the variance of the sampled ultrasound stacks. It is indeed expected and taken into account both in the registration function  $f_{ghyp}$  training, with data sampled at different regions from the entire liver, and in the RL function training, which only takes input data from one previous step in the Markovian environment. Therefore, it is not a requirement to follow the exact acquisition guidance to generate the next guidance.

After training, upon deployment of the framework in clinical practice, this becomes an interactive registration framework, where the RL function produces guidance of where to sample to obtain a better registration (in terms of registration accuracy) while the sampled images are registered with sections resampled from the preoperative volume. For receiving the stack start location guidance intraoperatively we propose to use a simple 2D map with the probe position localised with respect to a few key landmarks together with an indication of the direction and magnitude of movement for the probe based on the stack start location guidance as described in Section 2.3, but investigation of how this guidance visualisation impacts surgical procedures is beyond the scope of this work. At inference it is not required or indeed feasible to follow the guidance exactly and relatively large deviations were therefore reflected in the simulated intraoperative data. The “noisy” actions consisted of actions with noise added, in the ranges  $\pm 0.5$  cm for  $a_i^{len}$  and  $\pm 1.0$  cm for  $a_i^{loc}$ .

To summarise, we envision that within clinical practice the operator that controls the intraoperative imaging would get an overlay of the registered intraoperative image onto the preoperative image alongside our proposed map as a visual aid. The RL-suggested actions would then be delivered to the operator in a human-readable format via the map, and the operator may choose to follow this guidance if the overlay registration is to be improved.

### 3.3. Network implementation and training

The hypernet-based registration function consists of a displacement-predicting network based on a 3D U-Net (Çiçek et al., 2016; Hoopes et al., 2021) with a fully-connected hypernet for predicting the registration network parameters, as in previous works (Hoopes et al., 2021).

The RL function is based on a convolutional neural network (Schulman et al., 2017; Mnih et al., 2016), trained using proximal policy optimisation (Schulman et al., 2017). The specific network architectures and hyperparameters used are outlined in the Appendix. The hypernet-based registration function and RL function were trained for approximately 72 h and 192 h, respectively, on a single Nvidia Tesla V100 GPU. TensorFlow with python was used to implement the networks and their training procedures.

### 3.4. Evaluation metrics

We use the training set to train both functions independently, the validation set to monitor performance and rewards during training, and the holdout set (i.e., 4000 deformation cases simulated from 4 patients) to report performance metrics. Four metrics are reported in this work, Dice Similarity Coefficient (DSC), Global Target Registration Error (GTRE), Local Target Registration Error (LTRE) and a Geometric Distance Measure (GDM). The GTRE is computed for intraoperative stacks where the landmarks are visible, averaged over all patients and deformation cases in the holdout set; the landmarks used are the branching points for the right, middle and left hepatic vein, centre of the vena cava and aorta, the branching point for the portal vein and the gall bladder, as illustrated in Fig. 6. The LTRE is computed based on features randomly sampled from vessel boundaries, averaged over all patients and deformation cases in the holdout set, since deformation between intraoperative volume and preoperative volume is known from simulations. The GDM is the distance between the centre of the 3D section extracted from the preoperative volume using the global alignment function and the post-registration intraoperative image stack, averaged over all patients and deformation cases in the holdout set. A perfect DSC of 1 would indicate perfect overlap between the warped intraoperative stack and the latent preoperative section, where a DSC of 0 would indicate no overlap. The other three metrics are distance-based and zero would indicate a perfect registration. For context the metrics before registration were 0.21DSC, 41.8 mm GTRE, 36.7 mm LTRE, 29.4 mm GDM matching scales reported by Vijayan et al. (2014). The metrics after the simulated global alignment but before any deformable registration or guidance were 0.31DSC, 23.9 mm GTRE, 22.1 mm LTRE, 24.2 mm GDM matching errors scales from Ramalhinho et al. (2022).

### 3.5. Evaluating the registration function

The registration function is evaluated by comparing with other common registration methods, i.e., an iterative nonrigid registration, implemented using GPU-enabled TensorFlow (Fu et al., 2020a), and a non-hypernet deep learning based nonrigid registration, i.e. equivalent to a re-implementation of VoxelMorph (Balakrishnan et al., 2019). DRAMMS (Machado et al., 2018) and DEEDS (Heinrich, 2018) are also used for comparison due to observed performance improvements in other tasks (Xiao et al., 2019). Additionally, we have also tested a variant of our proposed registration function which takes inputs of images rather than segmentation maps, called  $f_{\text{ghyp}}$  (image-input). All reported variants were trained separately, and the architecture of the registration function is the same between networks (Section 3.3).

These comparisons are made with locations  $a_t^{\text{loc}}$  being randomly sampled within the liver, while the other two conditioning variables  $a_t^{\text{def}}$  and  $a_t^{\text{len}}$  are fixed and optimised over the holdout set. The optimal  $a_t^{\text{def}}$  and  $a_t^{\text{len}}$  for the hypernet were estimated using a grid search over an empirically defined range. For  $a_t^{\text{len}}$  the range was 2 cm to 10 cm, as suggested by clinicians with over 20 years of experience with intraoperative ultrasound. For  $a_t^{\text{def}}$ , this was 0 to 1 as in Hoopes et al. (2021) to be used as the regularisation weight as in Section 2.2.2. Similarly, for the deep-learning-based registration, optimal  $a_t^{\text{def}}$  and  $a_t^{\text{len}}$  were estimated using a grid search over this empirically defined range, where 20,000 training iterations of the registration network were used for this search and the one with optimal performance was selected for comparison.

### 3.6. Evaluating the test-time optimisation using the hypernet

The test-time optimisation enabled by the hypernet is evaluated by comparing, for locations  $a_t^{\text{loc}}$  being randomly sampled within the liver, the hypernet with optimal  $a_t^{\text{def}}$  and  $a_t^{\text{len}}$  selected by the RL function with:

- (1) the same hypernet with  $a_t^{\text{def}}$  and  $a_t^{\text{len}}$  randomly selected from the empirically defined range (Section 3.5), per sample;
- (2) the same hypernet with  $a_t^{\text{def}}$  and  $a_t^{\text{len}}$  randomly selected from an empirically defined range, fixed over the holdout set;
- (3) the same hypernet with  $a_t^{\text{def}}$  fixed and optimised over the holdout set and  $a_t^{\text{len}}$  sampled randomly per intraoperative stack;
- (4) a non-hypernet registration with  $a_t^{\text{def}}$  and  $a_t^{\text{len}}$  fixed and optimised over the holdout set as described above (see Section 3.5);
- (5) a non-hypernet registration with  $a_t^{\text{def}}$  and  $a_t^{\text{len}}$  fixed, selected randomly from an empirically defined range. For all of these comparisons, the locations are sampled randomly i.e.,  $a_t^{\text{loc}}$  is randomly sampled.

### 3.7. Evaluating the reinforcement learning function

To the best of our knowledge, no such interactively guided registration methods have been proposed yet. Next we evaluate the efficacy of the actions sampled from the RL function by conducting ablation studies to investigate the impact of RL suggested conditioning variables (ablation refers to the removal of a methodological component within the framework).

We compare our implementation (i.e., hypernet with all  $a_t^{\text{len}}$ ,  $a_t^{\text{def}}$  and  $a_t^{\text{loc}}$  all optimised per acquired image using the RL function) with the following:

- (1) the same hypernet with  $a_t^{\text{len}}$  and  $a_t^{\text{def}}$  optimally selected by the RL function per intraoperative stack sampled while locations  $a_t^{\text{loc}}$  are sampled randomly (see Section 3.6);
- (2) the same hypernet with fixed  $a_t^{\text{len}}$  and  $a_t^{\text{def}}$  (optimal over the set) and RL selected locations  $a_t^{\text{loc}}$ ;
- (3) the same hypernet with fixed  $a_t^{\text{len}}$  and  $a_t^{\text{def}}$  (optimal over the set) and locations  $a_t^{\text{loc}}$  randomly sampled (see Section 3.6);
- (4) a non-hypernet registration with conditioning variables  $a_t^{\text{loc}}$  and  $a_t^{\text{def}}$  fixed and optimised over the holdout set and locations  $a_t^{\text{loc}}$  sampled using RL function;
- (5) a non-hypernet registration with conditioning variables  $a_t^{\text{loc}}$  and  $a_t^{\text{def}}$  fixed and optimised over the holdout set as described above and locations  $a_t^{\text{loc}}$  sampled randomly (see Section 3.5).

### 3.8. Evaluating the impact of conditioning variables on registration performance

#### 3.8.1. Evaluating varying spatial regularity of deformation

Our proposed registration function allows for the spatial regularity of deformation to be controlled at test-time individually for each acquired US sweep using the  $a_t^{\text{def}}$ . We investigate how varying this for samples impacts the registration performance in terms of DSC. We sample various different locations within the liver and observe the impact of varying the spatial regularity of deformation on the registration performance at these locations.

#### 3.8.2. Evaluating varying US stack length

Similar to the spatial regularity of deformation, our proposed registration function also allows for the stack length of individual US sweeps to be controlled at test time. We investigate the effect of varying this across the liver. We sample different locations in the liver and vary the stack length parameter to investigate its effect on registration performance.

**Data:** Set of preoperative volumetric CT from different patient cases  $\{y\}_{1...N}$  and corresponding intraoperative volumes (deformed compared to preoperative volume) from which to sample intraoperative stacks.

**Result:** Trained RL function  $f_{\theta^{rl}}$  with optimised weights  $\theta^{rl}$ .

**while not converged do**

Sample a new preoperative volumetric CT  $y$  (and a corresponding deformed intraoperative volume to sample moving intraoperative image stacks from)

Start at  $t = 0$ ;

Sample an intraoperative stack  $x_t$  at a random location;

Compute  $z_t$  using  $d_t^{global} = g(x_t, y)$  followed by  $z_t = y \otimes (d_t^{global})^{-1}$  (Eq. (5) followed by Eq. (4));

Sample action from RL function  $a_t = f_{\theta^{rl}}(x_t, y)$  (Eq. (12));

Register  $x_t$  and  $z_t$  using  $\hat{d}_t^{local} = f_{\hat{\theta}^{hyp}}(x_t, z_t; a_t)$  followed by  $x_t \otimes \hat{d}_t^{local} \otimes d_t^{global}$  (i.e., Eq. (9) followed by Eq. (8));

Generate the next intraoperative image stack i.e.  $x_{t+1} = e(a_t)$  (Eq. (13));

Compute  $z_{t+1}$  using  $d_{t+1}^{global} = g(x_{t+1}, y)$  followed by  $z_{t+1} = y \otimes (d_{t+1}^{global})^{-1}$  (Eq. (5) followed by Eq. (4));

Sample action from RL function  $a_{t+1} = f_{\theta^{rl}}(x_{t+1}, y)$  (Eq. (12));

Register  $x_{t+1}$  and  $z_{t+1}$  using  $\hat{d}_{t+1}^{local} = f_{\hat{\theta}^{hyp}}(x_{t+1}, z_{t+1}; a_{t+1})$  followed by  $x_{t+1} \otimes \hat{d}_{t+1}^{local} \otimes d_{t+1}^{global}$  (i.e., Eq. (9) followed by Eq. (8));

Compute the first reward  $R_{t+1} = L_{\theta^{hyp}}(z_{t+1}, x_{t+1}, \hat{d}_{t+1}^{local}) - L_{\theta^{hyp}}(z_t, x_t, \hat{d}_t^{local})$  (Eq. (14));

**for**  $t \leftarrow 2$  **to**  $T$  **do**

Note:  $t$  is now iterating starting at  $t = 2$ ;

Generate the next intraoperative image stack i.e.  $x_t = e(a_{t-1})$  (Eq. (13));

Compute  $z_t$  using  $d_t^{global} = g(x_t, y)$  followed by  $z_t = y \otimes (d_t^{global})^{-1}$  (Eq. (5) followed by Eq. (4));

Sample action from RL function  $a_t = f_{\theta^{rl}}(x_t, y)$  (Eq. (12));

Register  $x_t$  and  $z_t$  using  $\hat{d}_t^{local} = f_{\hat{\theta}^{hyp}}(x_t, z_t; a_t)$  followed by  $x_t \otimes \hat{d}_t^{local} \otimes d_t^{global}$  (i.e., Eq. (9) followed by Eq. (8));

Compute the reward  $R_t = L_{\theta^{hyp}}(z_t, x_t, \hat{d}_t^{local}) - L_{\theta^{hyp}}(z_{t-1}, x_{t-1}, \hat{d}_{t-1}^{local})$  (Eq. (14));

**end**

Once  $R_{t=1:T}$  collected, update RL function using gradient ascent based on Eq. (15)

**end**

**Algorithm 1:** Interactions of the RL function and MDP environment

### 3.8.3. Evaluating registration performance at varying sampling locations

Our proposed registration function is generalisable to varying locations within the liver. To explore the effect of varying starting location  $a_t^{loc}$  across the liver, we compute registration performance at different locations within the liver, in terms of DSC to investigate which areas within the liver may be better registered.

## 4. Results

### 4.1. Registration performance

For evaluating the registration function, we turn the reader's attention to the top block of results presented in Table 1. When comparing the hypernet, with fixed optimal registration parameter values over the entire set, to a non-hypernet registration network or to an empirically configured iterative registration, we expect little difference between the two for supporting a test-time optimisable alternative without compromising registration capability. Indeed, as detailed in Table 1, statistical significance was not observed for all tested metrics (p-values > 0.050 for all), except for the GDM metric for comparison with the iterative registration variant (p-values = 0.008 and 0.006 for the two comparisons). The standard deviation of the Jacobian determinant was between 0.58-1.97 (95th percentile). For comparisons to DRAMMS and DEEDS, while some metrics are improved compared to the hypernet approach, these improvements are within the margin of error. Moreover, DEEDS or DRAMMS would require substantial modifications within our hypernet-based guidance framework, which allows test-time adaptation of conditioning variables. Note that these methods require geometrically reconstructed images on a regular grid, with substantial overlapping common anatomical features already present, which makes

their use in our application challenging and uncommon within current literature. Thus, we use our hypernet approach in subsequent experiments. It should also be noted that image-only registration network still manages to substantially reduce registration error compared to un-registered cases. We use the best-performing hypernet approach for subsequent experiments.

### 4.2. Test-time optimisation performance

For evaluating the test-time optimisation of conditioning variables using the hypernet we present results in the middle block of Table 1. Here, we observed higher performance of the proposed hypernet with stack length  $a_t^{len}$  and deformation regularisation  $a_t^{def}$  being optimised locally (estimates predicted by RL function) i.e. for each intraoperative image stack that is registered (first row in the middle block of the table). Statistical significance can be found when compared to all other tested methods for all tested metrics (all p-values < 0.010). For evaluating the test-time optimisation of conditioning variables using the hypernet we present results in the middle block of Table 1. Here, we observed higher performance of the proposed hypernet with stack length  $a_t^{len}$  and deformation regularisation  $a_t^{def}$  being optimised locally (estimates predicted by RL function) i.e. for each intraoperative image stack that is registered (first row in the middle block of the table). Statistical significance can be found when compared to all other tested methods for all tested metrics (all p-values < 0.010).

It is also noteworthy that all variants with stack length  $a_t^{len}$  and deformation regularisation  $a_t^{def}$  randomly selected form an empirical range performed worse than their counterparts with optimised stack length  $a_t^{len}$  and deformation regularisation  $a_t^{def}$  for all tested metrics (all p-values < 0.010). One exception to this was the GDM metric for

**Table 1**

Table of results on the holdout set. The conditioning variable settings are: ‘fixed optimal’: fixed for the entire holdout set and optimal selected by averaging over the whole set; ‘fixed random’: fixed for the entire holdout set and randomly selected from an empirically defined range; ‘random’: randomly selected per acquired intraoperative stack; ‘RL optimal’: optimal selected by RL function per acquired intraoperative stack. The networks are:  $f_{\theta^{hyp}}$ : hypernet;  $f_{\theta^{local}}$ : non-hypernet (VoxelMorph) (Balakrishnan et al., 2019); Iterative Registration: (Fu et al., 2020a). Bold indicates best performing method in each block.

| Network                          | $a_t^{len}$          | $a_t^{def}$          | $a_t^{loc}$       | DSC             | GTRE (mm)      | LTRE (mm)      | GDM (mm)       |
|----------------------------------|----------------------|----------------------|-------------------|-----------------|----------------|----------------|----------------|
| Un-registered                    | N/A                  | N/A                  | N/A               | $0.21 \pm 0.04$ | $41.8 \pm 0.7$ | $36.7 \pm 0.4$ | $29.4 \pm 0.5$ |
| Globally-registered              | N/A                  | N/A                  | N/A               | $0.31 \pm 0.06$ | $23.9 \pm 0.6$ | $22.1 \pm 0.6$ | $24.2 \pm 0.5$ |
| $f_{\theta^{hyp}}$               | <b>fixed optimal</b> | <b>fixed optimal</b> | <b>random</b>     | $0.68 \pm 0.05$ | $7.6 \pm 0.6$  | $7.9 \pm 0.4$  | $7.1 \pm 0.4$  |
| $f_{\theta^{hyp}}$ (image-input) | <b>fixed optimal</b> | <b>fixed optimal</b> | <b>random</b>     | $0.59 \pm 0.07$ | $8.9 \pm 0.7$  | $9.4 \pm 0.5$  | $8.3 \pm 0.4$  |
| $f_{\theta^{local}}$             | fixed optimal        | fixed optimal        | random            | $0.63 \pm 0.05$ | $8.6 \pm 0.3$  | $9.0 \pm 0.3$  | $7.0 \pm 0.2$  |
| Iterative                        | fixed optimal        | fixed optimal        | random            | $0.61 \pm 0.04$ | $8.7 \pm 0.4$  | $9.1 \pm 0.3$  | $8.0 \pm 0.2$  |
| DRAMMS                           | fixed optimal        | fixed optimal        | random            | $0.67 \pm 0.07$ | $8.1 \pm 0.7$  | $8.4 \pm 0.5$  | $7.4 \pm 0.5$  |
| DEEDS                            | fixed optimal        | fixed optimal        | random            | $0.69 \pm 0.05$ | $8.2 \pm 0.6$  | $7.9 \pm 0.3$  | $7.2 \pm 0.4$  |
| $f_{\theta^{hyp}}$               | <b>RL optimal</b>    | <b>RL optimal</b>    | <b>random</b>     | $0.74 \pm 0.05$ | $6.7 \pm 0.4$  | $7.1 \pm 0.6$  | $6.4 \pm 0.3$  |
| $f_{\theta^{hyp}}$               | random               | random               | random            | $0.62 \pm 0.04$ | $8.8 \pm 0.5$  | $9.8 \pm 0.4$  | $7.3 \pm 0.3$  |
| $f_{\theta^{hyp}}$               | fixed random;        | fixed random;        | random            | $0.60 \pm 0.06$ | $9.1 \pm 0.4$  | $9.9 \pm 0.6$  | $8.1 \pm 0.2$  |
| $f_{\theta^{hyp}}$               | random               | fixed optimal        | random            | $0.64 \pm 0.06$ | $7.9 \pm 0.4$  | $8.3 \pm 0.3$  | $7.0 \pm 0.3$  |
| $f_{\theta^{local}}$             | fixed optimal        | fixed optimal        | random            | $0.63 \pm 0.05$ | $8.6 \pm 0.3$  | $9.0 \pm 0.3$  | $7.0 \pm 0.2$  |
| $f_{\theta^{local}}$             | fixed random         | fixed random         | random            | $0.54 \pm 0.07$ | $9.9 \pm 0.6$  | $10.7 \pm 0.5$ | $7.5 \pm 0.3$  |
| $f_{\theta^{hyp}}$               | <b>RL optimal</b>    | <b>RL optimal</b>    | <b>RL optimal</b> | $0.77 \pm 0.05$ | $6.5 \pm 0.6$  | $7.0 \pm 0.3$  | $6.5 \pm 0.4$  |
| $f_{\theta^{hyp}}$               | RL optimal           | RL optimal           | random            | $0.74 \pm 0.05$ | $6.7 \pm 0.4$  | $7.1 \pm 0.6$  | $6.4 \pm 0.3$  |
| $f_{\theta^{hyp}}$               | fixed optimal        | fixed optimal        | RL optimal        | $0.71 \pm 0.05$ | $7.2 \pm 0.3$  | $7.5 \pm 0.4$  | $7.0 \pm 0.2$  |
| $f_{\theta^{hyp}}$               | fixed optimal        | fixed optimal        | random            | $0.68 \pm 0.05$ | $7.6 \pm 0.6$  | $7.9 \pm 0.4$  | $7.1 \pm 0.4$  |
| $f_{\theta^{local}}$             | fixed optimal        | fixed optimal        | RL optimal        | $0.67 \pm 0.04$ | $8.2 \pm 0.4$  | $8.4 \pm 0.4$  | $7.2 \pm 0.3$  |
| $f_{\theta^{local}}$             | fixed optimal        | fixed optimal        | random            | $0.63 \pm 0.05$ | $8.6 \pm 0.3$  | $9.0 \pm 0.3$  | $7.0 \pm 0.2$  |

the comparison between  $f_{\theta^{local}}$  with parameters randomly selected vs optimised ( $p$ -value = 0.052).

Moreover, comparing our method to three methods from existing works for random starting location  $a_t^{loc}$ , we found that our proposed method with a hypernet-based registration with locally optimised stack length  $a_t^{len}$  and deformation regularisation  $a_t^{def}$  performed better than the tested alternatives, with statistical significance (all  $p$ -values < 0.010). The tested alternatives were: (1) iterative registration (with fixed optimal  $a_t^{len}$  and  $a_t^{def}$ ); (2) non-hypernet registration with fixed optimal  $a_t^{len}$  and  $a_t^{def}$ ; (3) non-hypernet registration with fixed optimal  $a_t^{def}$  and random  $a_t^{len}$  (most commonly used in practice).

#### 4.3. Guidance performance

For this set of evaluations, in the bottom block of results presented in Table 1, we observed higher performance for our proposed solution, with a hypernet based registration function with RL estimated optimal stack length  $a_t^{len}$ , deformation regularisation  $a_t^{def}$  and starting location  $a_t^{loc}$  (first row in the bottom block of the Table), compared to all other tested methods (all  $p$ -values < 0.010).

It is noteworthy that all variants with locations sampled using the RL optimal estimates performed better than their counterparts with randomly sampled locations, with statistical significance for all tested metrics except GDM (all  $p$ -values < 0.010). For GDM significance was not found for these comparisons for any of the corresponding comparisons (all  $p$ -values > 0.050).

The hypernet-based registration function took approximately 120 ms at inference and the RL function took approximately 60 ms, making the combined inference time approximately 180 ms per sampled sweep.

#### 4.4. Exploring the effect of conditioning variables

We conduct some small exploratory studies to investigate the effects of the stack length  $a_t^{len}$  and deformation regularisation  $a_t^{def}$  and starting location  $a_t^{loc}$ . While these studies are not exhaustive in investigating the impacts at every possible location within the liver, they allow us to explore the additional flexible aspects of our framework. Results from these studies are presented in the following Sections 4.4.1, 4.4.2 and 4.4.3.




Fig. 7. Effect of varying  $a_t^{def}$ , the spatial regularity of deformation at various starting locations  $a_t^{loc}$ . Left: Linear scale; Right: Logarithmic scale, where each colour is a different sampling location and dotted lines indicate optimal values.




Fig. 8. Effect of  $a_t^{def}$ , the spatial regularity of deformation, averaged over 1000 locations.

##### 4.4.1. Effect of varying spatial regularity of deformation

We observe varying optimal values for the spatial regularity of deformation across varying starting location  $a_t^{loc}$  within the liver as shown in Fig. 7.

Fig. 8 shows that no consistent trend between the compared registration methods exists, when studying impact of spatial regularity of deformation i.e.,  $a_t^{len}$  in our framework. This, coupled with, no consistent performance improvement and other difficulties with using non-adaptive methods such as DEEDS and DRAMMS, as discussed in Section 4.1, lead us to the use of hyper-nets in our proposed framework.

##### 4.4.2. Effect of varying US stack length

We observe two trends across the liver as we vary the stack length at different locations, as shown in Fig. 9: (1) registration performance




Fig. 9. Effect of varying stack length at various sampling locations.




Fig. 10. Registration performance (DSC) overlaid onto a patient volume with fixed stack length and spatial regularity of deformation. Each column showing a different view for the same patient (legend on the top-left represents the view presented).

showing a constant upward trend or plateau with increasing stack length, for regions with long vessels (vessels that do not show branching over more than 4 cm); (2) registration performance declining after high stack length for vessel-rich regions (three or more branches appearing in a 4 cm by 4 cm by 4 cm region), respectively.

#### 4.4.3. Effect of varying starting location $a_i^{loc}$ with fixed stack length $a_i^{len}$ and deformation regularisation $a_i^{def}$

We compute registration performance at 8000 different randomly sampled locations within the liver, with other conditioning variables being fixed. We interpolate and smooth these values and overlay them onto a patient volume as show in Fig. 10. High DSC is observed for larger vessels and low DSC is observed in peripheral regions, where vessel sizes are smaller. Regions where either the moving or fixed image or both did not contain any vessels were omitted from the visualisation.

#### 4.5. Impact of segmentation on registration performance

To explore the impact of segmentation variability on registration performance, we add varying levels of corruption to our segmentation labels mimicking realistic segmentation noise in the form of missing or eroded (morphological filtering operator) vessels (Montaña-Brown et al., 2021). The corruption strength denotes the amount of corruption

Table 2

Impact of segmentation noise on registration performance. The registration model used was  $f_{ghyp}$  from the first row of Table 1.

| Corruption strength | Registration performance (DSC) |
|---------------------|--------------------------------|
| 0.0                 | $0.68 \pm 0.05$                |
| 0.1                 | $0.67 \pm 0.04$                |
| 0.2                 | $0.67 \pm 0.07$                |
| 0.3                 | $0.68 \pm 0.06$                |
| 0.4                 | $0.67 \pm 0.06$                |
| 0.5                 | $0.68 \pm 0.07$                |
| 0.6                 | $0.68 \pm 0.05$                |
| 0.7                 | $0.67 \pm 0.04$                |
| 0.8                 | $0.64 \pm 0.04$                |
| 0.9                 | $0.64 \pm 0.05$                |
| 1.0                 | $0.63 \pm 0.08$                |




Fig. 11. Registration results presented for 2D slices. Top: latent slice from CT section (obtained from global alignment), 2nd row: slice from warped US sweep (warped using hypernet-based registration function-predicted deformation), 3rd row: flow field showing warps, Bottom: slice from un-warped US sweep. Slices sampled using Ramalhinho et al. (2023).

applied where 1.0 indicates erosion applied to all visible vessels in the intraoperative volume and 0.0 indicates erosion applied to no vessels. Where vessels were smaller than 2.5 mm, these were removed instead of applying the erosion filter, if corruption was applied to the vessel. The results are summarised in Table 2.

We observe consistent performance for corruption strengths from 0.0 to 0.7 (statistical significance not found, all p-values > 0.074). Although a small reduction from corruption strength 0.7 to 0.8, this difference is statistically significant ( $p$ -value = 0.043). These results shows that our method is robust even to large amounts of segmentation noise and inaccuracies up-to corruption strength 0.7.

#### 4.6. Qualitative comparisons

Examples and qualitative results are presented in Figs. 11–13. We present registration results in both 2D and 3D. Additionally, noteworthy within these figures are the vessel sizes in the sampled US sweeps, which have narrower lumens due to insufflation and intraoperative deformation and nonrigid deformations compared to the ground truth CT data.

It should also be noted that these qualitative samples show that the locally nonrigid components are largely driven by the available vessel morphology, as expected. Locations where the support data is sparse are visibly rigid and global, owing to the global initialisation and deformation regularisation.




Fig. 12. Registration results presented for equally spaced 2D slices from a single 3D volume. Top: latent slice from CT section (obtained from global alignment), 2nd row: slice from warped US sweep (warped using hypernet-based registration function-predicted deformation), 3rd row: flow field showing warps, Bottom: slice from un-warped US sweep. Slices sampled using Ramalhinho et al. (2023).




Fig. 13. Registration results presented in 3D using interpolation of ultrasound frames. Red: Sampled US (moving image), Green: latent CT section (obtained from global alignment) (fixed image), Purple: Warped US (warped moving image).

## 5. Discussion

It is interesting to note that based on the results presented in Section 4, our proposed method with RL function-suggested conditioning variables performs statistically significantly better than the

variants from all other conducted experiments. This indicates, that the components of the proposed framework add value to the overall registration framework and therefore increase the registration accuracy. Furthermore, the conducted ablation experiments quantify the efficacy and their impact from individual proposed components, as removing each component in isolation always leads to a worse performance. We observe this for RL locations and the hypernet with locally optimised stack length  $a_i^{len}$  and deformation regularisation  $a_i^{def}$ , where these models consistently perform better than the ablated counterparts, for a variety of metrics. It is interesting to note an inconsistent trend in the GDM metric, which may rely more on global rigid alignment as opposed to local nonrigid alignment.

Further exploratory studies demonstrate the flexibility of the proposed framework, allowing us to investigate optimality in terms of stack length  $a_i^{len}$  and deformation regularisation  $a_i^{def}$  across the liver, and perhaps more importantly, the need for such flexibility in the real-world clinical hepatic application. Through these studies we found variable optimal settings for the stack length  $a_i^{len}$  and deformation regularisation  $a_i^{def}$  across the entire liver. This is also an evidence that registration performance can be improved by adaptive algorithms that allows overall inhomogeneity of registration performance, across large organs such as the liver, due to its varying appearances of vessels across the organ and varying optimal spatial regularity of deformation (partially based on vessel distribution).

The qualitative results further validate our framework and show that not only are the registrations physically plausible, the warped intraoperative volumes simulated from CT are also physically plausible, showing nonrigid deformation and narrower vessel lumens compared to the preoperative volumes (e.g., due to insufflation). The plausibility of registration is demonstrated via a 3D visualisation presented in Fig. 13, where we observe correction of intraoperative lumen narrowing and other deformations, while the overall vessel tree shape is preserved. Specifically, we can see for example, lumen narrowing being corrected in column 1, row 1 and column 2, row 4; translation and rotation being corrected in column 1, row 3 and column 2 row 2; as well as other corrections taking place, while preserving the overall shape of the vessels.

The use of simulated data and global alignment in this work, based on error estimates from previous works, serves as a first set of experimental results demonstrating feasibility of image acquisition guidance towards improved registration, in intraoperative settings. Despite the best effort in considering these plausible types of intraoperative variability, the limitations imposed by the fact that the intraoperative data are simulated in this study warrant follow-up study to test the generalisation to real-time-acquired interactive data using the proposed methods. Nonetheless, we argue that this study using the partially-simulated data is ethically and practically feasible, for reporting an informative first set of results enabling future prospective validation.

From a technical aspect, future work will also explore further registration parameters from two aspects of the image registration process, i.e., from image generation e.g., the tested stack length which controls density of sampling and from registration e.g., the tested spatial regularity of deformation. Additionally, we will explore others from other aspects of the registration pipeline, e.g., controlling size of the US volume, or feature extraction for registration. There is also an opportunity to further investigate guidance signals generated by the RL function, other than stack start location guidance and optimal registration parameter estimates, e.g., US sweep types such as a rotation, curved sweep, or sweep along the liver surface which lead to improved performance.

## 6. Conclusion

In this work we propose a flexible guided nonrigid registration framework based on two functions central to the framework. The first is a hyper-network-based registration function, which is adaptable

at test-time with respect to registration parameters. The second is a reinforcement learning function which learns to estimate optimal registration parameter settings as well as sampling locations to improve nonrigid registration performance for a common preoperative to intraoperative registration task, stack-to-volume registration. The guidance in terms of sampling location for intraoperative imaging and the settings for registration parameters, which control the image acquisition as well as registration, allows for a nonrigid registration to be obtained which leads to lower error compared to counterparts without these components for the tested application of liver ultrasound to CT registration. We have demonstrated the applicability of the proposed approach using real preoperative CT data and simulated intraoperative US data from patients undergoing laparoscopic liver surgery, and have shown meaningful performance improvements together with physically plausible registrations. Although, evaluated on simulated data, this work serves as a base for future work that can evaluate guidance towards improved registration within realistic intraoperative settings, where interaction between the RL controller and data acquisition are considered.

### CRedit authorship contribution statement

**Shaheer U. Saeed:** Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **João Ramalhinho:** Writing – review & editing, Software, Resources, Data curation, Conceptualization. **Nina Montaña-Brown:** Writing – review & editing. **Ester Bonmati:** Writing – review & editing. **Stephen P. Pereira:** Writing – review & editing, Resources, Project administration, Funding acquisition, Data curation, Conceptualization. **Brian Davidson:** Writing – review & editing, Supervision, Resources, Project administration, Funding acquisition, Data curation, Conceptualization. **Matthew J. Clarkson:** Writing – review & editing, Supervision, Resources, Project administration, Funding acquisition, Data curation, Conceptualization. **Yipeng Hu:** Writing – review & editing, Writing – original draft, Visualization, Supervision, Resources, Project administration, Methodology, Investigation, Funding acquisition, Data curation, Conceptualization.

### Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

### Acknowledgements

This work is supported by the EPSRC grant [EP/T029404/1]; Wellcome/EPSRC Centre for Interventional and Surgical Sciences [203145Z/16/Z]; the NIHR i4i grant [II-LA-1116-20005]; the EPSRC CDT in i4health [EP/S021930/1]; the International Alliance for Cancer Early Detection, an alliance between Cancer Research UK [EDDAPA-2024/100014] & [C73666/A31378], Canary Center at Stanford University, the University of Cambridge, OHSU Knight Cancer Institute, University College London and the University of Manchester; and the National Institute for Health Research University College London Hospitals Biomedical Research Centre.

### Appendix

#### A.1. Glossary of terminology and notations

Terminology and notation used to describe the data at time-step  $t$ .

- Intraoperative image stack ( $x_t$ ): in this work the US image stack (stack of multiple 2D slices) sampled from the domain  $\mathcal{X}$

- Preoperative image volume ( $y$ ): in this work a volumetric CT image acquired before the operation, sampled from the domain  $\mathcal{Y}$
- Conditioning variables ( $a_t$ ): variables used to control intraoperative stack sampling and registration with the preoperative volume; formed of  $\{a_t^{loc}, a_t^{en}, a_t^{def}\}$ 
  - $a_t^{loc}$ : start location for sampling an intraoperative stack
  - $a_t^{en}$ : length of an intraoperative stack
  - $a_t^{def}$ : deformation regularisation (spatial regularity of deformation)

Terminology and notation relating to the registration function.

- Image transform or warp operator ( $\otimes$ ): used to denote registration/warping between an image and displacement vectors
- Registration function ( $f_{\theta^{reg}}$ ): a registration function with parameters  $\theta^{reg}$
- Displacement vectors ( $d_t$ ): used to register the intraoperative stack
- Global alignment function ( $g$ ): used for rigid registration of the intraoperative stack and preoperative volume
- Latent preoperative section ( $z_t$ ): post-global alignment latent image features resampled from the preoperative volume  $y$
- Global displacement vectors ( $d_t^{global}$ ): the rigid global alignment displacement vector used to sample the latent preoperative section and produce a rigid registration
- Non-rigid registration function ( $f_{\theta^{local}}$ ): a non-rigid registration function modelled as a deep neural network with parameters  $\theta^{local}$
- Local displacement vectors ( $d_t^{local}$ ): used for non-rigid registration
- Hyper-network ( $h_{\theta^{hyp}}$ ): a hyper-network with parameters  $\theta^{hyp}$
- Hyper-network-based registration function ( $f_{\theta^{hyp}}$ ): a combination of the non-rigid registration function and hyper-network, with parameters  $\theta^{hyp}$
- Hyper-network loss function ( $L_{\theta^{hyp}}$ ): loss function used to measure image similarity and train the hyper-network-based registration function
- Dice loss:  $L^{dice}$ : image similarity loss between two binary masks
- $L^2$ -norm of displacement gradient:  $L^{def}$ : deformation regularisation controlling spatial regularity of deformation

Terminology and notation relating to the reinforcement learning (RL) function.

- Policy function ( $\pi_{\theta^{rl}}$ ): policy parameterised by  $\theta^{rl}$  to predict the probability of an action
- The observed state:  $s$ : a partial observable state of the Markov decision process environment
- RL function ( $f_{\theta^{rl}}$ ): the reinforcement learning function with parameters  $\theta^{rl}$ , modelled as a deep neural network
- Effector ( $e$ ): a non-parametric effector function within the Markov decision process environment which executes an action to sample the new observed state
- Reward ( $R_t$ ): The reward at time-step  $t$  used for training the RL function

Terminology and notation relating to the optional map construction.

- Pre-defined landmark point cloud ( $\bar{p}^v$ ): landmark point cloud extracted from preoperative image
- intraoperative landmark point cloud ( $p^s$ ): landmark point cloud extracted post-global alignment of intraoperative stacks
- Registered/corrected landmark point cloud ( $\hat{p}$ ): landmark point cloud extracted post-ICP between the pre-defined landmark point cloud and intraoperative landmark point cloud

**Table 3**Hyperparameters for  $f_{\theta^{hyp}}$ .

| Hyperparameter                  | Value            |
|---------------------------------|------------------|
| $f_{\theta^{loc}}$ architecture | U-Net            |
| Encoding Channels               | (16, 32, 32, 32) |
| Decoding Channels               | (32, 32, 32, 32) |
| $h_{\theta^{hyp}}$ architecture | Fully connected  |
| Units per layer                 | (64, 64, 64, 64) |
| Optimiser                       | Adam             |
| Batch size                      | 128              |

**Table 4**Hyperparameters for  $f_{\theta^{rl}}$ .

| Hyperparameter                 | Value                         |
|--------------------------------|-------------------------------|
| $f_{\theta^{rl}}$ architecture | CNN (Krizhevsky et al., 2017) |
| Conv Layers Channels           | (96, 256, 256, 96)            |
| FC Layers Units                | (32, 32, 32)                  |
| RL Algorithm                   | PPO (Schulman et al., 2017)   |
| Optimiser                      | Adam                          |
| Batch size                     | 256                           |
| Gamma $\gamma$                 | 0.99                          |
| Horizon (T)                    | 1024                          |
| GAE                            | 0.95                          |

### A.2. Hyperparameters for the hypernet-based registration function

Network architectures and hyperparameters for the hypernet-based registration function  $f_{\theta^{hyp}}$  are in Table 3:

### A.3. Hyperparameters for the reinforcement learning function

Network architectures and hyperparameters for the reinforcement learning function  $f_{\theta^{rl}}$  are in Table 4:

### A.4. Details of intraoperative deformation simulation

A  $4 \times 4 \times 4$  grid of points was overlaid onto the preoperative volume and deformation vectors were assigned to each point. The details for sampling deformation vectors are described below and in Table 5. Values were randomly sampled from the ranges specified in the table for each type of deformation and were assigned to each point based on the deformation weight as described below.

For coronal deformations, 16 points on the posterior-most plane were assigned zero displacement and the remaining points were assigned displacements based on the type of deformation applied, according to the deformation weights. As an example, referring to row 1 in the table, if deformation magnitude is  $-6$ , then all the anterior-most 16 points have displacements  $-6$ , the second row  $-6 \times 0.5$ , the third  $-6 \times 0.25$  and the fourth 0.

For transverse or axial deformations simulating breathing we can again follow an example, where if the deformation magnitude is 2, the superior-most 16 points have displacement 2, the second row  $2 \times 0.5$ , the third  $2 \times -0.5$  and the fourth  $2 \times -1$ . The values for the ranges for breathing were approximated based on Jagsi et al. (2007).

For postural deformations, random values were sampled for each point in the grid.

For each grid-point, the final displacement is a sum of the different displacement vectors.

Finally, after computing displacement vectors for the entire grid, these were used to interpolate a dense displacement field which was used to warp the preoperative volume.

This new volume was used for US simulation.

**Table 5**

Details of intraoperative deformations.

| Type of deformation                            | Simulation mechanism                                                         |
|------------------------------------------------|------------------------------------------------------------------------------|
| Coronal compression (insufflation)             | Range: $-8$ to $-2$ cm<br>Anterior to Posterior weight: 1, 0.5, 0.25, 0      |
| Coronal plane deformation (breathing)          | Range: $-3$ to $+3$ cm<br>Anterior to Posterior weight: 1, 0.5, 0.25, 0      |
| Transverse/axial plane deformation (breathing) | Range: $-3$ to $+3$ cm<br>Superior to Inferior weight: 1, 0.5, $-0.5$ , $-1$ |
| Deformation in all planes (postural)           | Range: $-1$ to $+1$ cm                                                       |

### A.5. Using the proposed framework to guide intraoperative prostate registration

The proposed framework is readily generalisable to other applications with a few modifications. We demonstrate this for an application of guiding intraoperative ultrasound (US) for registering US to preoperative MR images. For this purpose, we utilise the muRegPro open-source dataset (Baum et al., 2023), formed of real US and MR images from prostate cancer patients. The 3D US images in the dataset are already assumed globally-registered, thus a global alignment function is not required for this application, eliminating the function  $g$ . Additionally, instead of the free movement of the probe captured as  $a_i^{loc} \in \mathbb{R}^3$  allowed in the liver imaging application, we assume a digital or manual transperineal stepper which allows left or right movement in the coronal plane, captured as  $a_i^{loc} \in \{0, 1, -1\}$ , where 1 indicates movement of the US imaging probe to the right by one step of the stepper,  $-1$  one step to the left and 0 indicating no movement. Additionally, for simplicity, for prostate intraoperative imaging we assume a fixed sweep length of 4 slices, as in real procedures where either 2D slices are used or limited-field-of-view fixed-size 3D volumes are used, this removes the  $a_i^{len}$  variable from the framework.

These simple modifications allow us to utilise the proposed framework for an application of prostate US to MR registration guidance (see Table 6). The muRegPro dataset is formed of paired MR-US globally aligned images with 65, 20 and 20 images in the train, validation and holdout sets. The MR images are  $120 \times 128 \times 128$  voxels and the US images have between 57-112 2D slices of US, acquired using a digital bi-plane transperineal stepper. The US were reconstructed into volumes with  $81 \times 118 \times 88$  voxels. After re-sampling the US to the size of MR, we assumed 4 US slices to form an intraoperative volume to be non-rigidly registered with 4 corresponding slices of MR. These randomly sampled 4 slices from US volumes were used to train the hyper-network with the only conditioning variable as input to the network being  $a_i^{def}$  as  $a_i^{len}$  is omitted due to fixed volume size in this application. Additionally, different to the liver imaging application, in this application we use 5 binary masks to guide registration, as a 5-channel 3D volume. The 5 labels correspond to the prostate gland, tumours in the gland, urethra, seminal vesicles and fluid-filled cysts.

The sampling guidance  $a_i^{loc} \in \{0, 1, -1\}$ , corresponds to probe movement to acquire the next 4 US slices either to the left or right of the current probe position. This guidance will guide the operator to regions that improve registration performance, as outlined in our proposed framework.

The results for this application, for various settings, are presented below.

The results show that our proposed framework resulted in statistically significant performance improvements (p-values  $< 0.01$ ), compared to all other tested variants. The variants included commonly used strategies for setting the deformation regularisation and a non-guided approach to registration. Performance improvements, thus, demonstrate the efficacy of the learnt guidance and show that adapting our

**Table 6**

Results on, the holdout set, comparing performance for a task of guided prostate US to MR non-rigid registration. The conditioning variable settings are: 'fixed optimal': fixed for the entire holdout set and optimal selected by averaging over the whole set; 'fixed random': fixed for the entire holdout set and randomly selected from an empirically defined range; 'random': randomly selected per acquired intraoperative stack; 'RL optimal': optimal selected by RL function per acquired intraoperative stack.

| Network       | $d_i^{def}$   | $d_i^{loc}$ | DSC             |
|---------------|---------------|-------------|-----------------|
| Un-registered | N/A           | N/A         | $0.24 \pm 0.05$ |
| $f_{ghyp}$    | random        | random      | $0.72 \pm 0.06$ |
| $f_{ghyp}$    | random        | RL optimal  | $0.77 \pm 0.06$ |
| $f_{ghyp}$    | fixed random  | random      | $0.71 \pm 0.07$ |
| $f_{ghyp}$    | fixed random  | RL optimal  | $0.76 \pm 0.05$ |
| $f_{ghyp}$    | fixed optimal | random      | $0.75 \pm 0.07$ |
| $f_{ghyp}$    | fixed optimal | RL optimal  | $0.79 \pm 0.06$ |
| $f_{ghyp}$    | RL optimal    | random      | $0.77 \pm 0.05$ |
| $f_{ghyp}$    | RL optimal    | RL optimal  | $0.82 \pm 0.04$ |

proposed framework to other applications requires small application-specific adjustments.

### Data availability

The authors do not have permission to share data.

### References

Amin, D.V., Kanade, T., Digioia, A.M., Jaramaz, B., 2003. Ultrasound registration of the bone surface for surgical navigation. *Comput. Aided Surg.* 8 (1), 1–16.

Anon, 2023. Visible-Patient, <https://www.visiblepatient.com>, Online; accessed May 2023.

Balakrishnan, G., Zhao, A., Sabuncu, M.R., Guttag, J., Dalca, A.V., 2019. VoxelMorph: A learning framework for deformable medical image registration. *IEEE Trans. Med. Imaging* 38 (8), 1788–1800. <http://dx.doi.org/10.1109/tmi.2019.2897538>.

Baum, Z.M., Hu, Y., Barratt, D.C., 2022. Meta-learning initializations for interactive medical image registration. *IEEE Trans. Med. Imaging*.

Baum, Z., Saeed, S., Min, Z., Hu, Y., Barratt, D., 2023. MR to ultrasound registration for prostate challenge-dataset. In: *Medical Image Computing and Computer Assisted Intervention–MICCAI 2023*. Zenodo, Zenodo.

Chen, X., Diaz-Pinto, A., Ravikumar, N., Frangi, A.F., 2021. Deep learning in medical image registration. *Prog. Biomed. Eng.* 3 (1), 012003.

Çiçek, Ö., Abdulkadir, A., Lienkamp, S.S., Brox, T., Ronneberger, O., 2016. 3D U-net: Learning dense volumetric segmentation from sparse annotation. *arXiv:1606.06650*.

Ciecholewski, M., Kassjański, M., 2021. Computational methods for liver vessel segmentation in medical imaging: A review. *Sensors* 21 (6), 2027.

Davison, A.J., 2003. Real-time simultaneous localisation and mapping with a single camera. In: *Computer Vision, IEEE International Conference on*. Vol. 3, IEEE Computer Society, p. 1403.

Drakopoulos, F., Chrisochoides, N.P., 2016. Accurate and fast deformable medical image registration for brain tumor resection using image-guided neurosurgery. *Comput. Methods Biomech. Biomed. Eng.: Imaging & Vis.* 4 (2), 112–126.

Elhawary, H., Oguro, S., Tuncali, K., Morrison, P.R., Tatli, S., Shyn, P.B., Silverman, S.G., Hata, N., 2010. Multimodality non-rigid image registration for planning, targeting and monitoring during CT-guided percutaneous liver tumor cryoablation. *Academic Radiol.* 17 (11), 1334–1344.

Ferrante, E., Paragios, N., 2017. Slice-to-volume medical image registration: A survey. *Med. Image Anal.* 39, 101–123.

Fu, Y., Brown, N.M., Saeed, S.U., Casamitjana, A., Baum, Z.M.C., Delaunay, R., Yang, Q., Grimwood, A., Min, Z., Blumberg, S.B., Iglesias, J.E., Barratt, D.C., Bonmati, E., Alexander, D.C., Clarkson, M.J., Vercateren, T., Hu, Y., 2020a. DeepReg: a deep learning toolkit for medical image registration. *J. Open Source Softw.* 5 (55), 2705. <http://dx.doi.org/10.21105/joss.02705>, URL DOI: 10.21105/joss.02705.

Fu, Y., Lei, Y., Wang, T., Curran, W.J., Liu, T., Yang, X., 2020b. Deep learning in medical image registration: a review. *Phys. Med. Biol.* 65 (20), 20TR01.

Gao, L., Heath, D.G., Kuszyk, B.S., Fishman, E.K., 1996. Automatic liver segmentation technique for three-dimensional visualization of CT data. *Radiology* 201 (2), 359–364.

Gooya, A., Biros, G., Davatzikos, C., 2010. An EM algorithm for brain tumor image registration: A tumor growth modeling based approach. In: *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition-Workshops*. IEEE, pp. 39–46.

Haskins, G., Kruger, U., Yan, P., 2020. Deep learning in medical image registration: a survey. *Mach. Vis. Appl.* 31, 1–18.

Heinrich, M.P., 2018. Intra-operative ultrasound to MRI fusion with a public multimodal discrete registration tool. In: *Simulation, Image Processing, and Ultrasound Systems for Assisted Diagnosis and Navigation: International Workshops, POCUS 2018, BIVPCS 2018, CuRIOUS 2018, and CPM 2018, Held in Conjunction with MICCAI 2018, Granada, Spain, September 16–20, 2018, Proceedings*. Springer, pp. 159–164.

Heiselman, J.S., Jarnagin, W.R., Miga, M.I., 2020. Intraoperative correction of liver deformation using sparse surface and vascular features via linearized iterative boundary reconstruction. *IEEE Trans. Med. Imaging* 39 (6), 2223–2234.

Heldmann, S., Beuthien, B., Olesch, J., Papenberg, N., Fischer, B., 2010. Improved minimal-invasive laparoscopic liver surgery by registra-tion of 3D CT and 2D ultrasound slices. *Biomed. Tech.* 55 (1), 84–86.

Hermosillo, G., Chef'Hotel, C., Faugeras, O., 2002. Variational methods for multimodal image matching. *Int. J. Comput. Vis.* 50 (3), 329–343.

Hoffmann, M., Billot, B., Greve, D.N., Iglesias, J.E., Fischl, B., Dalca, A.V., 2021. SynthMorph: learning contrast-invariant registration without acquired images. *IEEE Trans. Med. Imaging* 41 (3), 543–558.

Hoopes, A., Hoffmann, M., Fischl, B., Guttag, J., Dalca, A.V., 2021. Hypermorph: Amortized hyperparameter learning for image registration. In: *Information Processing in Medical Imaging: 27th International Conference, IPMI 2021, Virtual Event, June 28–June 30, 2021, Proceedings* 27. Springer, pp. 3–17.

Hu, Y., Ahmed, H.U., Taylor, Z., Allen, C., Emberton, M., Hawkes, D., Barratt, D., 2012. MR to ultrasound registration for image-guided prostate interventions. *Med. Image Anal.* 16 (3), 687–703.

Jagsi, R., Moran, J.M., Kessler, M.L., Marsh, R.B., Balter, J.M., Pierce, L.J., 2007. Respiratory motion of the heart and positional reproducibility under active breathing control. *Int. J. Radiat. Oncology\* Biology\* Phys.* 68 (1), 253–258.

Julier, S.J., Uhlmann, J.K., 2001. Building a million beacon map. In: McKee, G.T., Schenker, P.S. (Eds.), *Sensor Fusion and Decentralized Control in Robotic Systems IV*. Vol. 4571, SPIE, International Society for Optics and Photonics, pp. 10–21. <http://dx.doi.org/10.1117/12.444158>.

Krizhevsky, A., Sutskever, I., Hinton, G.E., 2017. Imagenet classification with deep convolutional neural networks. *Commun. ACM* 60 (6), 84–90.

Lange, T., Eulenstein, S., Hünerbein, M., Schlag, P.-M., 2003. Vessel-based non-rigid registration of MR/CT and 3D ultrasound for navigation in liver surgery. *Comput. Aided Surg.* 8 (5), 228–240.

Lange, T., Papenberg, N., Heldmann, S., Modersitzki, J., Fischer, B., Lamecker, H., Schlag, P.M., 2009. 3D ultrasound-CT registration of the liver using combined landmark-intensity information. *Int. J. Comput. Assist. Radiol. Surg.* 4, 79–88.

Lee, D., Nam, W.H., Lee, J.Y., Ra, J.B., 2010. Non-rigid registration between 3D ultrasound and CT images of the liver based on intensity and gradient information. *Phys. Med. Biol.* 56 (1), 117.

Machado, I., Toews, M., Luo, J., Unadkat, P., Essayed, W., George, E., Teodoro, P., Carvalho, H., Martins, J., Golland, P., et al., 2018. Deformable MRI-ultrasound registration via attribute matching and mutual-saliency weighting for image-guided neurosurgery. In: *Simulation, Image Processing, and Ultrasound Systems for Assisted Diagnosis and Navigation: International Workshops, POCUS 2018, BIVPCS 2018, CuRIOUS 2018, and CPM 2018, Held in Conjunction with MICCAI 2018, Granada, Spain, September 16–20, 2018, Proceedings*. Springer, pp. 165–171.

Menteşe, A., Bilge, H.S., 2017. Automatic classification of liver veins in ultrasound images. In: *2017 Medical Technologies National Congress. TIPTEKNO, IEEE*, pp. 1–4.

Mnih, V., Badia, A.P., Mirza, M., Graves, A., Lillicrap, T.P., Harley, T., Silver, D., Kavukcuoglu, K., 2016. Asynchronous methods for deep reinforcement learning. *arXiv:1602.01783*.

Montaña-Brown, N., Ramalhinho, J., Allam, M., Davidson, B., Hu, Y., Clarkson, M.J., 2021. Vessel segmentation for automatic registration of untracked laparoscopic ultrasound to CT of the liver. *Int. J. Comput. Assist. Radiol. Surg.* 16 (7), 1151–1160.

Nam, W.H., Kang, D.-G., Lee, D., Lee, J.Y., Ra, J.B., 2012. Automatic registration between 3D intra-operative ultrasound and pre-operative CT images of the liver based on robust edge matching. *Phys. Med. Biol.* 57 (1), 69–91.

Özgür, E., Koo, B., Le Roy, B., Buc, E., Bartoli, A., 2018. Preoperative liver registration for augmented monocular laparoscopy using backward-forward biomechanical simulation. *Int. J. Comput. Assist. Radiol. Surg.* 13, 1629–1640.

Papież, B.W., Heinrich, M.P., Risser, L., Schnabel, J.A., 2013. Complex lung motion estimation via adaptive bilateral filtering of the deformation field. In: *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2013: 16th International Conference, Nagoya, Japan, September 22–26, 2013, Proceedings, Part III* 16. Springer, pp. 25–32.

Park, H., Lee, G.M., Kim, S., Ryu, G.H., Jeong, A., Sagong, M., Park, S.H., 2022. A meta-learning approach for medical image registration. In: *2022 IEEE 19th International Symposium on Biomedical Imaging. ISBI, IEEE*, pp. 1–5.

Penney, G.P., Blackall, J.M., Hamady, M., Sabharwal, T., Adam, A., Hawkes, D.J., 2004. Registration of freehand 3D ultrasound and magnetic resonance liver images. *Med. Image Anal.* 8 (1), 81–91.

Porter, B.C., Rubens, D.J., Strang, J.G., Smith, J., Totterman, S., Parker, K.J., 2001. Three-dimensional registration and fusion of ultrasound and MRI using major vessels as fiducial markers. *IEEE Trans. Med. Imaging* 20 (4), 354–359.

Ramalhinho, J., Dowrick, T., Bonmati, E., Clarkson, M.J., 2023. Fan-slicer: A pycuda package for fast reslicing of ultrasound shaped planes. *J. Open Res. Softw.* 11 (1), 3.

- Ramalhinho, J., Koo, B., Montaña-Brown, N., Saeed, S.U., Bonmati, E., Gurusamy, K., Pereira, S.P., Davidson, B., Hu, Y., Clarkson, M.J., 2022. Deep hashing for global registration of untracked 2D laparoscopic ultrasound to CT. *Int. J. Comput. Assist. Radiol. Surg.* 17 (8), 1461–1468.
- Ramalhinho, J., Robu, M.R., Thompson, S., Gurusamy, K., Davidson, B., Hawkes, D., Barratt, D., Clarkson, M.J., 2018. A pre-operative planning framework for global registration of laparoscopic ultrasound to CT images. *Int. J. Comput. Assist. Radiol. Surg.* 13, 1177–1186.
- Ramalhinho, J., Tregidgo, H.F., Gurusamy, K., Hawkes, D.J., Davidson, B., Clarkson, M.J., 2020. Registration of untracked 2D laparoscopic ultrasound to CT images of the liver using multi-labelled content-based image retrieval. *IEEE Trans. Med. Imaging* 40 (3), 1042–1054.
- Risholm, P., Samset, E., Wells, W., 2010. Bayesian estimation of deformation and elastic parameters in non-rigid registration. In: *Biomedical Image Registration: 4th International Workshop, WBIR 2010, Lübeck, Germany, July 11–13, 2010. Proceedings 4*. Springer, pp. 104–115.
- Robu, M.R., Ramalhinho, J., Thompson, S., Gurusamy, K., Davidson, B., Hawkes, D., Stoyanov, D., Clarkson, M.J., 2018. Global rigid registration of CT to video in laparoscopic liver surgery. *Int. J. Comput. Assist. Radiol. Surg.* 13, 947–956.
- Saeed, S.U., Fu, Y., Stavrinides, V., Baum, Z.M., Yang, Q., Rusu, M., Fan, R.E., Sonn, G.A., Noble, J.A., Barratt, D.C., et al., 2021. Adaptable image quality assessment using meta-reinforcement learning of task amenability. In: *Simplifying Medical Ultrasound: Second International Workshop, ASMUS 2021, Held in Conjunction with MICCAI 2021, Strasbourg, France, September 27, 2021, Proceedings 2*. Springer, pp. 191–201.
- Saeed, S.U., Fu, Y., Stavrinides, V., Baum, Z.M., Yang, Q., Rusu, M., Fan, R.E., Sonn, G.A., Noble, J.A., Barratt, D.C., et al., 2022. Image quality assessment for machine learning tasks using meta-reinforcement learning. *Med. Image Anal.* 78, 102427.
- Saeed, S.U., Ramalhinho, J., Pinnock, M., Shen, Z., Fu, Y., Montaña-Brown, N., Bonmati, E., Barratt, D.C., Pereira, S.P., Davidson, B., et al., 2024. Active learning using adaptable task-based prioritisation. *Med. Image Anal.* 95, 103181.
- Salas-Moreno, R.F., Newcombe, R.A., Strasdat, H., Kelly, P.H., Davison, A.J., 2013. Slam++: Simultaneous localisation and mapping at the level of objects. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 1352–1359.
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., Klimov, O., 2017. Proximal policy optimization algorithms. *arXiv:1707.06347*.
- Simpson, I.J., Cardoso, M.J., Modat, M., Cash, D.M., Woolrich, M.W., Andersson, J.L., Schnabel, J.A., Ourselin, S., Initiative, A.D.N., et al., 2015. Probabilistic non-linear registration with spatially adaptive regularisation. *Med. Image Anal.* 26 (1), 203–216.
- Song, Y., Totz, J., Thompson, S., Johnsen, S., Barratt, D., Schneider, C., Gurusamy, K., Davidson, B., Ourselin, S., Hawkes, D., et al., 2015. Locally rigid, vessel-based registration for laparoscopic liver surgery. *Int. J. Comput. Assist. Radiol. Surg.* 10, 1951–1961.
- Stachniss, C., Leonard, J.J., Thrun, S., 2016. Simultaneous localization and mapping. In: *Siciliano, B., Khatib, O. (Eds.), Springer Handbook of Robotics*. Springer International Publishing, Cham, pp. 1153–1176. [http://dx.doi.org/10.1007/978-3-319-32552-1\\_46](http://dx.doi.org/10.1007/978-3-319-32552-1_46).
- Sun, Y., Moelker, A., Niessen, W.J., van Walsum, T., 2018. Towards robust ct-ultrasound registration using deep learning methods. In: *Understanding and Interpreting Machine Learning in Medical Image Computing Applications*. Springer, pp. 43–51.
- Thomson, B.R., Smit, J.N., Ivashchenko, O.V., Kok, N.F., Kuhlmann, K.F., Ruers, T.J., Fusaglia, M., 2020. MR-to-US registration using multiclass segmentation of hepatic vasculature with a reduced 3D U-net. In: *Medical Image Computing and Computer Assisted Intervention–MICCAI 2020: 23rd International Conference, Lima, Peru, October 4–8, 2020, Proceedings, Part III*. Springer, pp. 275–284.
- Tomaževič, D., Likar, B., Pernuš, F., 2005. Reconstruction-based 3D/2D image registration. In: *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2005: 8th International Conference, Palm Springs, CA, USA, October 26–29, 2005, Proceedings, Part II 8*. Springer, pp. 231–238.
- Vijayan, S., Reinertsen, I., Hofstad, E.F., Rethy, A., Hernes, T.A.N., Langø, T., 2014. Liver deformation in an animal model due to pneumoperitoneum assessed by a vessel-based deformable registration. *Minim. Invasive Ther. & Allied Technol.* 23 (5), 279–286.
- Villalon, J., Joshi, A.A., Lepore, N., Brun, C., Toga, A.W., Thompson, P.M., 2011. Comparison of volumetric registration algorithms for tensor-based morphometry. In: *2011 IEEE International Symposium on Biomedical Imaging: From Nano To Macro*. IEEE, pp. 1536–1541.
- Wei, X., Chen, X., Lai, C., Zhu, Y., Yang, H., Du, Y., et al., 2021a. Automatic liver segmentation in CT images with enhanced GAN and mask region-based CNN architectures. *BioMed Res. Int.* 2021.
- Wei, W., Haishan, X., Alpers, J., Rak, M., Hansen, C., 2021b. A deep learning approach for 2D ultrasound and 3D CT/MR image registration in liver tumor ablation. *Comput. Methods Programs Biomed.* 206, 106117.
- Wei, W., Rak, M., Alpers, J., Hansen, C., 2020. Towards fully automatic 2D US to 3D CT/MR registration: A novel segmentation-based strategy. In: *2020 IEEE 17th International Symposium on Biomedical Imaging. ISBI, IEEE*, pp. 433–437.
- Wein, W., Brunke, S., Khamene, A., Callstrom, M.R., Navab, N., 2008. Automatic CT-ultrasound registration for diagnostic imaging and image-guided intervention. *Med. Image Anal.* 12 (5), 577–585.
- Xiao, Y., Rivaz, H., Chabanas, M., Fortin, M., Machado, I., Ou, Y., Heinrich, M.P., Schnabel, J.A., Zhong, X., Maier, A., et al., 2019. Evaluation of MRI to ultrasound registration methods for brain shift correction: the CuRIOUS2018 challenge. *IEEE Trans. Med. Imaging* 39 (3), 777–786.
- Yoon, J., Arik, S., Pfister, T., 2020. Data valuation using reinforcement learning. In: *International Conference on Machine Learning*. PMLR, pp. 10842–10851.
- Zeng, Y.Z., Zhao, Y.Q., Liao, M., Zou, B.J., Wang, X.F., Wang, W., 2016. Liver vessel segmentation based on extreme learning machine. *Phys. Medica* 32 (5), 709–716.