Fast deformable image registration uncertainty estimation for contour propagation in daily adaptive proton therapy

Author name(s) withheld EMAIL(S) WITHHELD
Address withheld

Editors: Under Review for MIDL 2022

Abstract

In daily adaptive proton therapy, deformable image registration (DIR) can be used to propagate manually delineated contours from a reference CT to the daily CT for plan reoptimization. However, the ill-posedness of DIR implies uncertainty on the DIR hyperparameters, which results in uncertainty in the displacement field. In this work, a fast deep learning method is developed to predict the uncertainty associated with a DIR result without the need for Monte-Carlo (MC) sampling. It is shown that this results in a significant time reduction compared to MC whilst leading to similar probabilistic contours.

Keywords: Deformable Image Registration, Uncertainty, Proton Therapy, Adaptive Planning, Deep Learning

1. Introduction

In conventional radiotherapy, the treatment is planned based on a reference CT taken before the treatment. However, because of daily set-up and anatomical changes, the delivered dose can deviate from the planned dose, especially in proton therapy (PT), where dose gradients are large because of the peaked depth-dose profile. Target margins are therefore applied, increasing the dose to the healthy tissue. Daily adaptive PT (DAPT) aims to overcome this by reoptimizing the treatment based on a daily CT, reducing the size of the target margins.

In order to reoptimize, the tumor and organs-at-risk need to be delineated on the daily scan. Manual contouring is not feasible as the time between the acquisition of the CT and the treatment should be as low as possible to ensure maximal correspondence between the image and the treated patient. Automatic delineation can be obtained by registering the reference to the daily CT and propagating the reference contours to it. In case of deforming anatomy (e.g. weight loss), deformable image registration (DIR) is used. However, DIR is ill-posed and the optimal hyperparameter set for each scan pair is unknown. This uncertainty can be accounted for by assuming a probability distribution for each hyperparameter and propagating it through the DIR and contour propagation algorithms, yielding probabilistic contours. These can be used in combination with robust optimization techniques, e.g. to adaptively adjust margins in uncertain regions. Although Monte-Carlo (MC) sampling could be used, this requires several iterations of the DIR which increases the time between the scan and the treatment unacceptably. In this work, a deep learning method requiring only one DIR run is developed to speed up the DIR uncertainty quantification.

2. Method

We trained a neural network to predict the uncertainty of a given deformable vector field (DVF) based on example uncertainties which we quantified using MC sampling. First, we explain how the labels are obtained, after which the neural network is discussed.

2.1. Label generation

We used the plastimatch b-spline algorithm for DIR, which includes two important hyperparameters: the grid spacing s, i.e. the spacing between the control points of b-splines and the regularization λ, which trades off the image similarity and the smoothness of the DVF. For s we assumed a uniform distribution (see further). Because $\lambda \geq 0$, we assumed a lognormal distribution.

To estimate the mean and variance of $log(\lambda)$ and the bounds of s, we calculated the target registration error (TRE) on the publicly available DIRLAB 4D lung CT dataset for a range of hyperparameters (Fig. 1). The algorithm performs well for $10^{-4} \leq \lambda \leq 10$ and $5 \leq s \leq 50$ mm, hence the choice of probability densities as in Fig. 1. Note that these bounds are relatively wide to include hyperparameters that might be optimal for other anatomical regions or larger deformations. For 52 CT image pairs, the standard deviation σ_{MC} of each vector with respect to a reference DVF was estimated by running 100 DIR iterations with randomly sampled hyperparameters from their respective distributions.

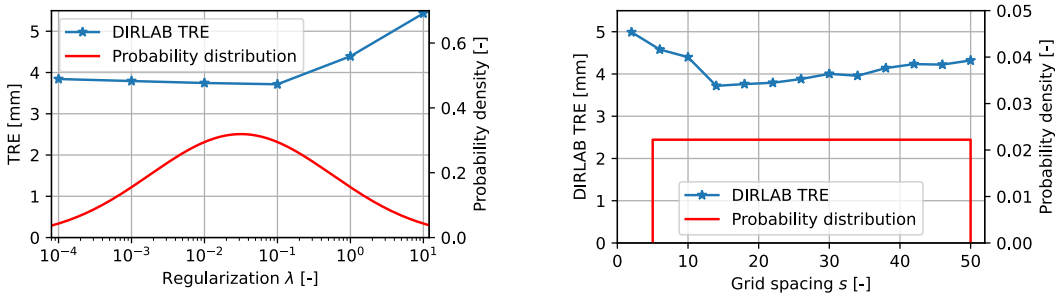

Figure 1: Target registration error (TRE) on the DIRLAB database as a function of (a) the regularization λ and (b) grid spacing s

2.2. Neural network

A 3D UNet was trained to predict the standard deviation σ of a DVF. The network inputs are the reference and daily CT, together with the reference DVF. The network was trained with mean absolute error over 150 epochs, with axis aligned flipping and random cropping.

3. Results

The MSE on the validation and test sets is significantly lower than using linear regression (Tab. 1), which was previously used to estimate DIR uncertainty in Amstutz et al. (2021). This is likely due to the inclusion of spatial and contrast information in our method, whereas the linear regression only uses the magnitude of the deformation. Indeed, the trained network is able to spatially locate regions with high and low uncertainty (Fig. 2).

In a probabilistic contour, the value for each voxel depicts its probability of being part of the organ. It can be generated by sampling a normally distributed DVF with the reference DVF as mean and the predicted σ_p as σ and propagating the contours with each sample. We compare this to the MC approach for a representative image, both using 100 samples. First, we find that our method is a factor 60 faster (Tab. 1). Further, the probabilistic contours of both methods are similar, with a mean absolute deviation of 15% for the voxels with probabilities between 1% and 99% (Fig. 3). This difference is likely due to the assumption of a Gaussian vector field and the use of a limited number of samples.

Table 1: Left: Mean squared error (MSE) on the validation and test sets using linear regression and our method. Right: Comparison of time required for a Monte Carlo simulation with 100 samples versus our method for a representative image pair.

MSE [mm^2]	**Linear Regression**	**Our Method**
Validation	0.181	0.086
Test	0.158	0.089

Time [s]	**Monte Carlo**	**Our Method**
DIR samples	4563	46
Inference	-	5
Contour prop	21	21
Total	4584	72

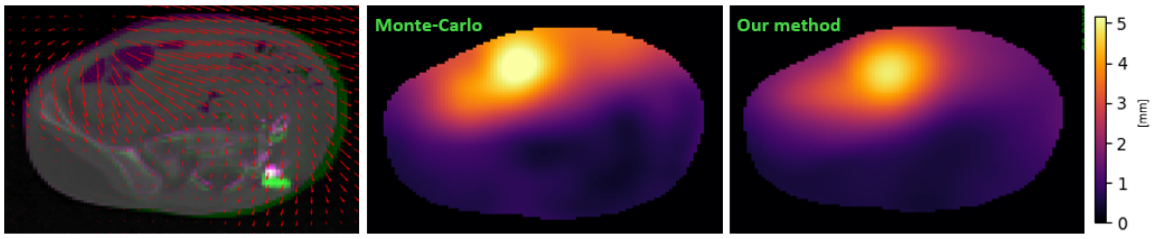

Figure 2: Left: Example image in the validation set. Middle: Uncertainty label σ_{MC}. Right: Predicted uncertainty σ_p.

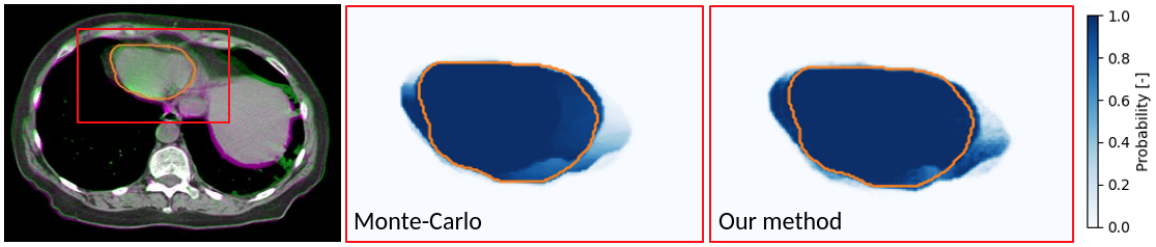

Figure 3: Comparison of the propagated contour of a heart with MC and our method. The orange contour depicts the manually annotated contour. Left: overlayed reference and daily CTs. Middle: MC propagated contour. Right: propagated contour with our method.

4. Conclusion and outlook

This work presents a deep learning method to estimate the uncertainty of a DIR result. It is shown that the method outperforms linear regression for this task. Furthermore, we demonstrate that the resulting probabilistic DVFs allow to create probabilistic contours similar to ones created with Monte-Carlo simulations with a time reduction factor around 60. Future work will look into relaxing the assumption of a normally distributed DVF to improve the accuracy of the probabilistic contours.

References

Florian Amstutz, Lena Nenoff, Francesca Albertini, Cássia O Ribeiro, Antje C Knopf, Jan Unkelbach, Damien C Weber, Antony J Lomax, and Ye Zhang. An approach for estimating dosimetric uncertainties in deformable dose accumulation in pencil beam scanning proton therapy for lung cancer. *Physics in Medicine & Biology*, 66(10):105007, 2021.