

UNSUPERVISED MEDICAL IMAGE ALIGNMENT WITH CURRICULUM LEARNING

Mihail Burduja Radu Tudor Ionescu

Faculty of Mathematics and Computer Science, University of Bucharest, Romania

ABSTRACT

We explore different curriculum learning methods for training convolutional neural networks on the task of deformable pairwise 3D medical image registration. To the best of our knowledge, we are the first to attempt to improve performance by training medical image registration models using curriculum learning, starting from an easy training setup in the first training stages, and gradually increasing the complexity of the setup. On the one hand, we consider two existing curriculum learning approaches, namely curriculum dropout and curriculum by smoothing. On the other hand, we propose a novel and simple strategy to achieve curriculum, namely to use purposely blurred images at the beginning, then gradually transit to sharper images in the later training stages. Our experiments with an underlying state-of-the-art deep learning model show that curriculum learning can lead to superior results compared to conventional training. Additionally, we show that curriculum by input blur has the best accuracy versus speed trade-off among the compared curriculum learning approaches.

Index Terms— Unsupervised learning, image registration, medical image alignment, curriculum learning.

1. INTRODUCTION

Image registration is a smooth alignment between two images that maps point coordinates from one image to corresponding point coordinates in the second image. *Medical image registration* maps point coordinates of some anatomical structures from the first image onto point coordinates from the same anatomical structures in the second image. Medical image registration is one of the most studied problems in medical image analysis [1, 2, 3, 4, 5, 6, 7], which helps with the alignment and fusion of scans taken through different means, e.g. computer tomography (CT) or magnetic resonance imaging (MRI), or at different times, leading to better computer assisted diagnosis among other benefits.

The state-of-the-art methods for medical image registration are based on deep neural networks [2, 6, 7]. These models are usually trained using a variant of stochastic gradient descent on mini-batches that are randomly selected from the training set. However, the conventional training based on random data selection is not always optimal. As noted by Bengio et al. [8], curriculum learning represents a training strategy

from easy to hard, which may guide neural models to better local optima. To this end, we investigate a series of curriculum learning strategies for unsupervised medical image registration. Besides considering state-of-the-art curriculum learning approaches such as curriculum dropout [9] and curriculum by smoothing [10] for our experiments, we propose a novel curriculum learning strategy that starts the training process on intentionally blurred images and then makes a gradual transition to sharper images in the later training stages.

We perform our evaluation on the SLIVER [11] data set in order to compare the proposed curriculum learning methods with conventional training. The underlying model in our experiments is the state-of-the-art recursive cascade network proposed by Zhao et al. [7]. The empirical results indicate that curriculum learning can lead to superior performance compared to the conventional learning process based on randomly chosen mini-batches.

2. RELATED WORK

Image registration has been studied extensively in literature [1, 2, 3, 4, 5, 6, 7]. Medical image registration methods are used for the alignment and fusion of different types of scans (CT, PET, MRI, etc.) of the same anatomical structure, being useful in computer aided diagnosis, computer assisted surgery and treatment. Recent approaches use convolutional neural networks (CNNs) trained specifically for registration [2, 6, 7, 12]. One of the first networks that was used for image registration is a spatial transformer network (STN) [12], which was not originally proposed as a method for image registration, requiring a certain degree of adaptation.

Engineered image registration methods. Early methods for image registration were based on solving the differential equations that morph one image into another [13]. DARTTEL [1] is a method that uses a single flow field and is computationally more efficient than methods that use multiple flow fields. Moreover, the resulting deformations are easily invertible. The iterative closest point algorithm [14] represents another method for 3D image registration. The method employs the iterative closest point algorithm to handle the full six degrees of freedom. Other early approaches are intensity-based or feature-based methods, usually relying on handcrafted features. The main drawback of these methods [1, 15, 16] is that the involved differential equations are difficult to work with, leading to suboptimal results in

realistic scenarios that entail non-rigid distortions.

Trainable image registration methods. The lack of reliance on handcrafted features and the vast ability of deep neural networks to learn patterns of anatomical structures makes them good candidates for the task of medical image registration. The research on image registration based on deep neural networks is rich [17, 18, 19], and it usually involves supervised models that learn from ground-truth labels provided in the form of segmentations or warp fields. However, medical data containing ground-truth labels is scarce, mainly due to the laborious annotation process that is required. Hence, the potential of using supervised learning for medical image registration is limited. To this end, several recent works proposed unsupervised frameworks that use data sets [17, 20] with synthetic annotations and similarity-based loss functions [2, 6, 7]. Similar to these studies, we consider an unsupervised framework. Unlike prior works on medical image registration, we study curriculum learning approaches to improve the results of unsupervised models for medical image registration.

Curriculum learning. Curriculum learning was introduced by Bengio et al. [8] as a technique to optimize machine learning models. It is based on the fact that neural networks are inspired by the human brains and they should be trained the same way humans learn – starting with easy examples and progressively increasing their complexity. Since its introduction, curriculum learning was applied on a wide range of deep learning problems, as noted in the recent survey of Soviany et al. [21]. A problem of the original (data-level) curriculum learning formulation is that the training examples must be ranked according to their complexity [21], preventing the application on tasks and data sets for which the complexity of the examples is unknown or cannot be determined. While this situation is no longer problematic for natural images since the development of image difficulty estimators [22], we are still confronted with the problem when it comes to medical images. To avoid this problem, some (model-level) curriculum learning methods turned to alternative strategies such as gradually increasing the model’s complexity [23], decreasing the dropout rate [9] or smoothing the convolutional filters [10]. These methods provide promising results when the difficulty of the data samples is hard to determine. In this work, we propose a curriculum learning strategy that artificially simplifies the examples in the initial training stages through blurring. Our method can be regarded as a data-level curriculum technique which, unlike other data-level strategies, does not require a difficulty estimator for the data samples.

3. METHOD

The *Volume Tweening Network* (VTN) was introduced by Zhao et al. [6] as an unsupervised end-to-end framework that uses CNNs for 3D medical image registration. The VTN is based on a set of stacked registration sub-networks, achieving state-of-the-art results by predicting a dense flow field using deconvolutional layers. There are two types of sub-

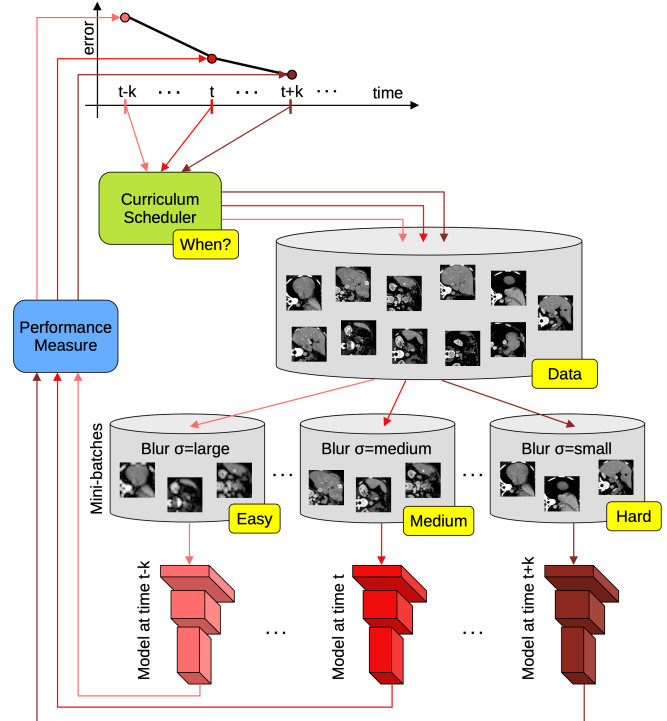


Fig. 1. The proposed curriculum learning method based on input blur. Best viewed in color.

networks used in the VTN, namely affine networks and dense deformable registration networks. The affine network aims to perform an initial alignment, and is only used once, as the first sub-network. The dense deformable networks are based on an encoder-decoder architecture that receives as input the fixed image and the current moving image (obtained from previous sub-networks in the stack). The encoder part is composed of convolutional layers and the decoder is composed of deconvolutional layers, similar to the U-Net architecture [24]. The output of each sub-network is a dense flow field that contains three-axis displacements of the same size as the input. As the underlying model for our study, we employ the 1-cascade VTN network [6, 7] and augment it as needed in order to apply a curriculum learning strategy or the other.

Curriculum by input blur. One strategy to perform curriculum learning without having a way to estimate the difficulty of data samples is to purposely simplify the inputs. In this paper, we propose to artificially simplify the inputs, reducing the amount of information, by blurring the images using a Gaussian filter. We control the amount of blur by adjusting the parameter σ with respect to the training stage. In the early training stages, we use a higher value for σ , blurring the input images in a more aggressive manner. We conjecture that blurred images are easier to align as the blur hides away fine misalignment errors. As the model begins to learn to align blurry images, we gradually reduce σ until it becomes 0 in the later training stages. Hence, at the end of the training process, the neural model is trained on original images. Our curriculum learning strategy is illustrated in Figure 1.

Curriculum dropout. Dropout [25] was introduced as a regularization technique, to prevent neural networks from overfitting. The idea behind dropout is to randomly deactivate certain neurons in a neural network during training, leading to a distributed and more robust representation. In the original formulation, the dropout rate is fixed. Curriculum dropout was introduced in [9] as an adaptive dropout that decreases the rate of dropout in later stages of training up to some minimum value that is set in advance. Since the baseline VTN network does not use dropout, we set the minimum dropout rate to 0.0.

Curriculum by smoothing. Similar to our curriculum learning strategy that blurs the examples in the initial stages of training, curriculum by smoothing [10] applies a low-pass filter to convolutional filters. Applying low-pass filters in the early stages of training reduces the noise propagated through the network, improving convergence. As the training progresses, more and more high frequency data is let through the network. In [10], a 2D Gaussian kernel is used as the low-pass filter. Since, we are dealing with 3D scans as input, we need to apply a 3D Gaussian kernel:

$$k_{G^3D}(x, y, z, \sigma) = \frac{1}{(\sqrt{2\pi}\sigma)^3} \cdot \exp\left(-\frac{x^2 + y^2 + z^2}{2\sigma^2}\right), \quad (1)$$

where σ controls the amount of blur.

A typical 3D convolutional network for image registration, such as the 1-cascade VTN [6, 7], is constructed of several convolutional blocks of the following form:

$$h_i = \text{activation}(w_i * h_{i-1}), \quad (2)$$

where h_{i-1} and h_i are the 4D input and output tensors of block i , w_i are the learnable parameters of the convolutional block and $*$ denotes the 3D convolution operation. For VTN, the activation function chosen by Zhao et al. [7] is leaky ReLU. The smoothing is applied as follows:

$$h_i = \text{activation}(k_{G^3D} \hat{*}(w_i * h_{i-1})), \quad (3)$$

where k_{G^3D} is a 3D Gaussian kernel and $\hat{*}$ denotes the separable 3D convolution operation. We emphasize that it is not straightforward to apply a 3D Gaussian filter in commonly used libraries for deep neural network training. Indeed, the implementation requires a separable 3D convolution to apply the 3D kernel k_{G^3D} on a 4D input tensor. Another disadvantage of this method is the added model complexity and the additional training time required for applying Gaussian filters on all convolutional layers. In this regard, our novel approach based on blurring the 3D input scans represents a less cumbersome solution to achieve the same result.

4. EXPERIMENTS

Data sets. We train and evaluate the models on the same liver data sets as [7]. More specifically, the training is performed on the MSD [26] and BFH [6] data sets. The validation is conducted on the LiTS [27] data set, while the testing is performed on the SLIVER [11] data set. MSD contains various types of CT scans of liver tumors (70 scans), hepatic vessels

(443 scans) and pancreas tumors (420 scans). BHF contains 92 scans. There are no annotations for MSD and BHF (the training is unsupervised). LiTS contains 131 liver scans with ground-truth segmentations. Similarly, SLIVER contains 20 scans with ground-truth liver segmentations. We used the pre-processed data sets provided by Zhao et al. [7]. The pre-processing includes cropping and resampling into volumes of $128 \times 128 \times 128$ voxels.

Baseline. In our experiments, we consider the 1-cascade recursive VTN [6, 7] as the base model, which we trained for a number of 40,000 iterations on randomly generated mini-batches of four samples each.

Parameter tuning. We set the training hyperparameters according to [7], regardless of the fact that the VTN is trained with conventional or curriculum learning. However, the curriculum learning strategies have additional hyperparameters that need to be set, as described below. All curriculum strategies are applied over the first 20,000 training iterations, using a linear curriculum scheduler. Then, in the last 20,000 iterations, the conventional training regime is resumed. For curriculum by input blur, we apply a Gaussian kernel starting with a σ of 1.0 that linearly decreases to 0.0 during the first 20,000 training iterations. Analogously, for curriculum by smoothing, the parameter σ of the Gaussian filter is decreased linearly from 1.0 to 0.0 in the first 20,000 iterations. For curriculum dropout, we linearly decrease the dropout rate from 0.5 to 0.0 in the first 20,000 training iterations.

Evaluation metrics. We employ the evaluation framework provided in [7] for a direct comparison with the baseline VTN. In the evaluation framework, the considered metrics are the Dice coefficient and the Jaccard index. The Dice coefficient measures the mean overlap as the intersection between the warped and target volumes divided by their mean volume. The intersection-over-union (IoU), also known as the Jaccard index, is an alternative way to measure the overlap as the intersection between the warped and target volumes divided by their union. The two metrics are computed as follows:

$$\text{Dice}(A, B) = 2 \cdot \frac{|A \cap B|}{|A| + |B|}; \text{Jaccard}(A, B) = \frac{|A \cap B|}{|A \cup B|}. \quad (4)$$

For both metrics, higher values indicate better performance.

Results. We report the results on the SLIVER data set in Table 1. First, we observe that two of the studied curriculum learning methods achieve superior results compared to the conventional training regime. These are curriculum by input blur and curriculum by smoothing. Their improvements are above 0.72% in terms of the Dice coefficient and above 1.19% in terms of the Jaccard index. We notice that curriculum dropout is not useful in our case. In a set of extra experiments, we observed that dropout alone also degrades performance. This might explain why curriculum dropout is rendered ineffective for the 1-cascade VTN. In Figure 2, we show a series of examples in which the output of the baseline 1-cascade VTN is compared with the outputs of the 1-cascade

Table 1. Dice and Jaccard scores on the SLIVER data set. Three curriculum learning strategies are compared with the conventional training regime. The best results are highlighted in bold. The times required for executing a step with each training strategy are measured on a Google Colab machine with an Intel(R) Xeon(R) CPU at 2.20GHz with 12GB of RAM and an Nvidia Tesla P100 GPU with 16GB of VRAM.

Method	#Steps	Time per step (seconds)	Dice	Jaccard
1-cascade VTN [7]	40000	0.187	0.90268	0.82365
1-cascade VTN based on curriculum by input blur	40000	0.193	0.91037	0.83636
1-cascade VTN based on curriculum dropout	40000	0.195	0.90097	0.82088
1-cascade VTN based on curriculum by smoothing	40000	0.276	0.90990	0.83557

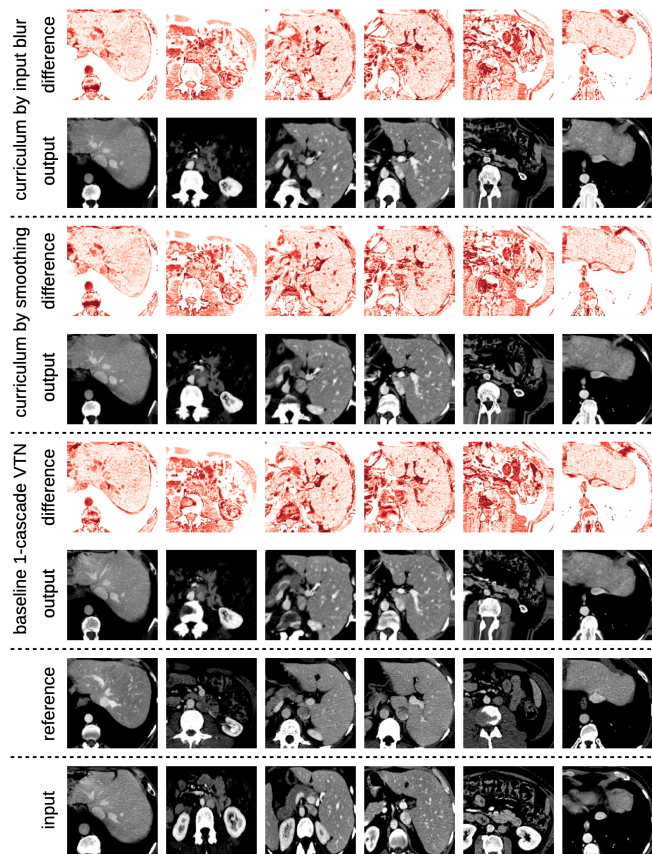


Fig. 2. Comparative results of the baseline 1-cascade VTN versus two of our best curriculum learning strategies. The mean absolute differences between the reference and the warped output are represented on a white-to-red scale, where the red intensity is proportional to the magnitude of the difference. Shown samples are 2D excerpts from full 3D scans of different patients. Best viewed in color.

VTNs based on curriculum by smoothing and curriculum by input blur, respectively. Consistent with the results presented in Table 1, we observe that curriculum learning helps the network to minimize the differences between the warped outputs and the reference images depicted in Figure 2.

While curriculum by input blur is the top scoring approach, we emphasize that the training time compared to

curriculum by smoothing is significantly lower, being very close to the training time of the baseline VTN. This happens because blurring the input images is very fast. Curriculum by smoothing applies a 3D Gaussian blur on all convolutional filters, which is more time consuming. Curriculum dropout is as fast as our approach based on curriculum by input blur, but, as noted earlier, curriculum dropout does not bring any performance improvements over the baseline VTN. In conclusion, curriculum by input blur provides the best trade-off between accuracy and speed.

5. CONCLUSION

In this work, we studied the possibility of employing a series of curriculum learning methods for medical image registration. Additionally, we proposed a novel and effective curriculum learning regime based on input blur. Our empirical results showed that curriculum learning can bring performance gains to a state-of-the-art unsupervised medical image registration network [7]. In future work, we aim to apply curriculum learning strategies on additional medical image registration frameworks and evaluate the resulting models on multiple benchmarks.

Acknowledgments. The research leading to these results has received funding from the NO Grants 2014-2021, under project ELO-Hyp contract no. 24/2020.

6. REFERENCES

- [1] J. Ashburner, “A fast diffeomorphic image registration algorithm,” *NeuroImage*, vol. 38, no. 1, pp. 95–113, 2007.
- [2] G. Balakrishnan, A. Zhao, M. R. Sabuncu, J. Guttag, and A. V. Dalca, “An unsupervised learning model for deformable medical image registration,” in *Proceedings of CVPR*, 2018, pp. 9252–9260.
- [3] L. G. Brown, “A survey of image registration techniques,” *ACM Computing Surveys*, vol. 24, no. 4, pp. 325–376, 1992.
- [4] D. L. G. Hill, P. G. Batchelor, M. Holden, and D. J. Hawkes, “Medical image registration,” *Physics in Medicine & Biology*, vol. 46, no. 3, pp. R1, 2001.

- [5] F. P. M. Oliveira and J. M. R. S. Tavares, "Medical image registration: a review," *Computer Methods in Biomechanics and Biomedical Engineering*, vol. 17, no. 2, pp. 73–93, 2014.
- [6] S. Zhao, T. Lau, J. Luo, E. I. Chang, and Y. Xu, "Unsupervised 3D end-to-end medical image registration with volume tweening network," *IEEE Journal of Biomedical and Health Informatics*, vol. 24, no. 5, pp. 1394–1404, 2019.
- [7] S. Zhao, Y. Dong, E. I. Chang, and Y. Xu, "Recursive cascaded networks for unsupervised medical image registration," in *Proceedings of ICCV*, 2019, pp. 10600–10610.
- [8] Y. Bengio, J. Louradour, R. Collobert, and J. Weston, "Curriculum learning," in *Proceedings of ICML*, 2009, pp. 41–48.
- [9] P. Morerio, J. Cavazza, R. Volpi, R. Vidal, and V. Murino, "Curriculum dropout," in *Proceedings of ICCV*, 2017, pp. 3544–3552.
- [10] S. Sinha, A. Garg, and H. Larochelle, "Curriculum by smoothing," in *Proceedings of NeurIPS*, 2020, pp. 21653–21664.
- [11] T. Heimann, B. Van Ginneken, M. A. Styner, Y. Arzhaeva, V. Aurich, C. Bauer, A. Beck, C. Becker, R. Beichel, G. Bekes, et al., "Comparison and evaluation of methods for liver segmentation from CT datasets," *IEEE Transactions on Medical Imaging*, vol. 28, no. 8, pp. 1251–1265, 2009.
- [12] M. Jaderberg, K. Simonyan, A. Zisserman, and K. Kavukcuoglu, "Spatial transformer networks," in *Proceedings of NIPS*, 2015, pp. 2017–2025.
- [13] G. E. Christensen, R. D. Rabbitt, and M. I. Miller, "3D brain mapping using a deformable neuroanatomy," *Physics in Medicine & Biology*, vol. 39, no. 3, pp. 609, 1994.
- [14] P. J. Besl and N. D. McKay, "Method for registration of 3-d shapes," in *Sensor Fusion IV: Control Paradigms and Data Structures*, 1992, vol. 1611, pp. 586–606.
- [15] M. F. Beg, M. I. Miller, A. Trouvé, and L. Younes, "Computing large deformation metric mappings via geodesic flows of diffeomorphisms," *International Journal of Computer Vision*, vol. 61, no. 2, pp. 139–157, 2005.
- [16] M. Vaillant, M. I. Miller, L. Younes, and A. Trouvé, "Statistics on diffeomorphisms via tangent space representations," *NeuroImage*, vol. 23, pp. S161–S169, 2004.
- [17] A. Dosovitskiy, P. Fischer, E. Ilg, P. Hausser, C. Hazirbas, V. Golkov, P. Van Der Smagt, D. Cremers, and T. Brox, "FlowNet: Learning optical flow with convolutional networks," in *Proceedings of ICCV*, 2015, pp. 2758–2766.
- [18] J. Krebs, T. Mansi, H. Delingette, L. Zhang, F. C. Ghesu, S. Miao, A. K. Maier, N. Ayache, R. Liao, and A. Kamen, "Robust non-rigid registration through agent-based action learning," in *Proceedings of MICCAI*, 2017, pp. 344–352.
- [19] M. M. Rohé, M. Datar, T. Heimann, M. Sermesant, and X. Pennec, "SVF-Net: Learning deformable image registration using shape matching," in *Proceedings of MICCAI*, 2017, pp. 266–274.
- [20] N. Mayer, E. Ilg, P. Hausser, P. Fischer, D. Cremers, A. Dosovitskiy, and T. Brox, "A large dataset to train convolutional networks for disparity, optical flow, and scene flow estimation," in *Proceedings of CVPR*, 2016, pp. 4040–4048.
- [21] P. Soviany, R. T. Ionescu, P. Rota, and N. Sebe, "Curriculum learning: A survey," *arXiv preprint arXiv:2101.10382*, 2021.
- [22] R. T. Ionescu, B. Alexe, M. Leordeanu, M. Popescu, D. P. Papadopoulos, and V. Ferrari, "How hard can it be? Estimating the difficulty of visual search in an image," in *Proceedings of CVPR*, 2016, pp. 2157–2166.
- [23] T. Karras, T. Aila, S. Laine, and J. Lehtinen, "Progressive growing of GANs for improved quality, stability, and variation," in *Proceedings of ICLR*, 2018.
- [24] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional networks for biomedical image segmentation," in *Proceedings of MICCAI*, 2015, pp. 234–241.
- [25] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, "Dropout: a simple way to prevent neural networks from overfitting," *Journal of Machine Learning Research*, vol. 15, no. 1, pp. 1929–1958, 2014.
- [26] A. L. Simpson, M. Antonelli, S. Bakas, M. Bilello, K. Farahani, B. Van Ginneken, A. Kopp-Schneider, B. A. Landman, G. Litjens, B. Menze, et al., "A large annotated medical image dataset for the development and evaluation of segmentation algorithms," *arXiv preprint arXiv:1902.09063*, 2019.
- [27] P. Bilic, P. F. Christ, E. Vorontsov, G. Chlebus, H. Chen, Q. Dou, C. W. Fu, X. Han, P. A. Heng, J. Hesser, et al., "The Liver Tumor Segmentation Benchmark (LiTS)," *arXiv preprint arXiv:1901.04056*, 2019.