



HAL
open science

Exploring Deep Registration Latent Spaces

Théo Estienne, Maria Vakalopoulou, Stergios Christodoulidis, Enzo Battistella, Théophraste Henry, Marvin Lerousseau, Amaury Leroy, Guillaume Chassagnon, Marie-Pierre Revel, Nikos Paragios, et al.

► **To cite this version:**

Théo Estienne, Maria Vakalopoulou, Stergios Christodoulidis, Enzo Battistella, Théophraste Henry, et al.. Exploring Deep Registration Latent Spaces. DART in MICCAI 2021, the 24th International Conference on Medical Image Computing and Computer Assisted Intervention, Sep 2021, Strasbourg, France. pp.112-122, 10.1007/978-3-030-87722-4_11 . hal-03524105

HAL Id: hal-03524105

<https://hal.science/hal-03524105v1>

Submitted on 13 Jan 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Exploring Deep Registration Latent Spaces

Théo Estienne^{1,2}, Maria Vakalopoulou¹, Stergios Christodoulidis¹, Enzo Battistella^{1,2}, Théophraste Henry², Marvin Lerousseau^{1,2}, Amaury Leroy^{2,4}, Guillaume Chassagnon³, Marie-Pierre Revel³, Nikos Paragios⁴, and Eric Deutsch²

¹ Université Paris-Saclay, CentraleSupélec, Mathématiques et Informatique pour la Complexité et les Systèmes, Inria Saclay, 91190, Gif-sur-Yvette, France.

`theo.estienne@centralesupelec.fr`

² Université Paris-Saclay, Institut Gustave Roussy, Inserm, Radiothérapie Moléculaire et Innovation Thérapeutique, 94800, Villejuif, France.

³ Département de Radiology, Hôpital Cochin, AP-HP Centre, Université de Paris, 27 Rue du Faubourg Saint-Jacques, 75014 Paris, France

⁴ TheraPanacea, Pépinière Santé Cochin, Paris, France.

Abstract. Explainability of deep neural networks is one of the most challenging and interesting problems in the field. In this study, we investigate the topic focusing on the interpretability of deep learning-based registration methods. In particular, with the appropriate model architecture and using a simple linear projection, we decompose the encoding space, generating a new basis, and we empirically show that this basis captures various decomposed anatomically aware geometrical transformations. We perform experiments using two different datasets focusing on lungs and hippocampus MRI. We show that such an approach can decompose the highly convoluted latent spaces of registration pipelines in an orthogonal space with several interesting properties. We hope that this work could shed some light on a better understanding of deep learning-based registration methods.

Keywords: Deep Learning-based Medical Image Registration · Deformable Registration · Explainability

1 Introduction

Deep learning methods provide the state of the art performance for various applications currently. This is due to their inherent property to generate highly abstract representations hierarchically. These representations are building on top of each other, making it possible to encode highly non-linear manifolds. Even though such hierarchies can outperform traditional methods, they lack explainability, making their translation difficult to solve real-life problems. This drawback is of great significance in the medical field and especially for the algorithms that are intended to be adapted to clinical practice, addressing problems of precision medicine [12,2]. For these reasons, it is essential to identify ways to understand better the high throughput operations that are applied.

Recently, with the introduction of the differentiable spatial transformer [13], trainable deep learning registration methods are becoming more and more popular, reducing computational times while reporting similar to traditional methods performance [1,25,17]. Meanwhile, the deformation field, which is one of the products of deformable registration methods, has been shown to encode not only the spatial correspondences but also clinical relevant information that could add valuable aspects to a variety of problems related to survival assessment or anomaly detection [18]. Indeed, encoding information between subjects can be very informative for various medical tasks such as medical image segmentation [9]. However, according to our knowledge, there are not many efforts focusing on understanding and analysing this encoding information which could initiate the explainability of deep learning-based registration methods.

In this study, we propose a framework for interpreting the encoded representation of deep learning-based registration methods. In particular, with the appropriate model architecture and by using a simple linear projection, we decompose the encoding space, generating a new basis that captures various geometrical operations. This decomposed encoding space is then driving the generation of the deformation field. The contributions of this work are twofold: *(i)* to the best of our knowledge, this study is one of the first to explore the explainability of deep learning-based registration methods through their encodings using linear projections, *(ii)* we show empirically, using two different datasets, one focusing on lungs and the other on the brain hippocampus that our projections are associated with different types of deformations and in particular rigid transformations. We hope that this work can highlight the very challenging topic of explainability of deep neural networks.

2 Related Work

Explaining how deep neural networks function is a matter of extensive research the recent years. GradCam [21] is one of the most popular methods that can provide some insights on deep neural networks for many applications, including medical imaging. GradCam highlights the region of the original input that contributes the most to the final prediction, producing coarse heatmaps based on the gradients. Similar to GradCam, there are many additional methods based on the gradient [27,3,24] that are commonly used for the explainability of the models. Moreover, in [10] the authors proposed a general framework of explanations as meta-predictors while they also reinterpret the network’s saliency providing a natural generalisation of the gradient-based saliency techniques. Even though such approaches can provide information on where the models attend, they can be mostly utilised in classification or detection schemes.

Representation disentangling methodologies is a concurrent field of research also investigating explainability topics. Such approaches are mainly focusing on generating interpretable latent representations by enforcing several constraints. This can be achieved either using architecture tricks [22,20] or with appropriate loss functions [4,14]. In medical image computing, several studies focus on

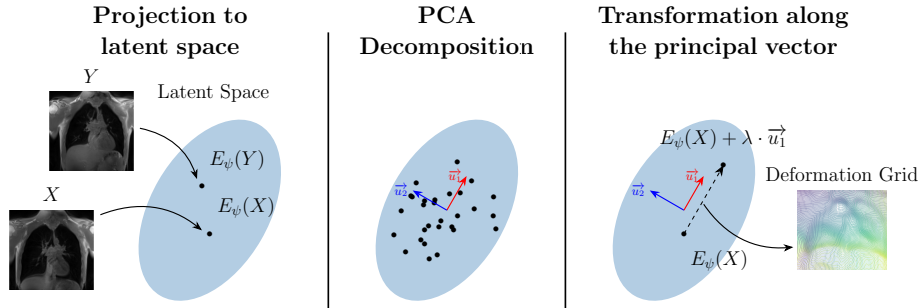


Fig. 1: Overall overview of our proposed framework. The different subjects (X, Y) are projected on the latent representation by the encoder E_ψ and then a linear decomposition of this latent space is calculated to identify a new vector space (\vec{u}_i) .

approaches for generating disentangled or decomposed representations. In [19] for example, the authors proposed a multimodal image registration method by decomposing the volumes into a common latent shape space and separate latent appearance spaces via image-to-image translation approach and generative models. Our method shares many common points with the approaches mentioned above, yet it focuses on exploring the registration latent space decomposition.

3 Methodology

Deep learning-based registration methods have received much attention in the last few years [1,25]. Formally, let us consider two volumes, the moving M and the fixed F . The goal of deep learning-based registration methods is to obtain the best parameters θ^* for the network g_θ that will map most accurately M to F using the predicted deformation grid Φ . The network g_θ usually is composed of an encoding E_ψ and a decoding D_ω part.

There are multiple ways to fuse the input volumes in deep learning-based registration approaches. Most of the methods use an early fusion strategy on which the two volumes are concatenated before they pass through the g_θ . However, some methods investigate late fusion strategies [11,7] where the two volumes pass independently through the encoder, and their merging operation is achieved in the encoding representation using various operations such as concatenation or subtraction. Thanks to this formulation, each volume has a unique encoding representation. In this study, we adopt the second strategy using the subtraction operation to encode each volume independently and calculate its latent space’s linear decomposition. In Figure 1, the overall scheme is presented.

3.1 Deep learning-based registration scheme

To perform our experiments and obtain our embeddings, we defined a network based on a 3D UNet architecture [5]. The encoder and the decoder are composed

of a fixed number of blocks with 3D convolution layers (stride 3, padding 1), instance normalisation layer and leaky ReLU activation function. The down and up-sampling operations are performed with a 3D convolution layer with stride and padding of 2. One of the main differences in our architecture was the absence of skip connections. Indeed, we want to enforce that all information passes through the last encoding layer without any leak due to the skip connections. This modification led us to reduce the downsampling operations from four to three for the lung dataset, to maintain the spatial resolution of the bottleneck.

Different formulations have been proposed to generate the deformation from deep learning schemes, such as displacement field formulation [1], diffeomorphic formulations [6,15] and formulations based on the spatial gradients [25]. In this work, we focused on the last one, with our network regressing the spatial gradients $\nabla_x \Phi_x$, $\nabla_y \Phi_y$ and $\nabla_z \Phi_z$, while the final deformation field is obtained through a cumulative sum operation. We also followed the symmetric formulation proposed in [8], predicting both the forward and backward deformations: $\nabla \Phi_{M \rightarrow F} = D_\omega(E_\psi(M) - E_\psi(F))$ and $\nabla \Phi_{F \rightarrow M} = D_\omega(E_\psi(M) - E_\psi(F))$.

The network was trained with a combination of four losses, one focusing on the intensity similarity using normalised cross-correlation (\mathcal{L}_{sim}), one focusing on anatomical structures using dice loss (\mathcal{L}_{seg}) and two losses for regularisation of the displacements. The first one was the Jacobian loss which is exploited on different works such as [17,16,26] (\mathcal{L}_{jac}) and the second one enforcing smooth gradients similar to [8] (\mathcal{L}_{smooth}). As such our final loss is: $\mathcal{L} = (\mathcal{L}_{sim} + \mathcal{L}_{seg} + \alpha \mathcal{L}_{smooth} + \beta \mathcal{L}_{jac})_{M \rightarrow F} + (\mathcal{L}_{sim} + \mathcal{L}_{seg} + \alpha \mathcal{L}_{smooth} + \beta \mathcal{L}_{jac})_{F \rightarrow M}$ with α and β being the weights of the regularisation losses.

3.2 Decomposition of latent space

Let $A_{train} = \{X_i | i \in [0, n]\}$ be the set of our n training samples. The proposed formulation apply the encoder independently to each volume, and thus we can obtain the set of latent vectors: $E_\psi(A_{train}) = \{E_\psi(X_i) | i \in [0, n]\}$. Then, we decompose this space using principal components analysis (PCA). That way, we obtain a set of principal vectors $\mathcal{U}_K = (\vec{u}_1, \dots, \vec{u}_K)$ with K being a hyperparameter fixing the number of principal components. It worth noting that each vector \vec{u}_i has the same size as the activation map of the encoder’s last layer. This size depends on the number of channels, the size of the input images and the number of downsampling operations. We flatten each encoding representation from its four dimensions representation (channel dimension and the three spatial dimensions) to a one-dimensional array to perform the PCA. Thus, the PCA is not calculated channel-wise, but all the channels are considered together. Each principal vector \vec{u}_i can be converted to a deformation grid ϕ_i using the corresponding decoder D_ω : $\phi_i = D_\omega(\vec{u}_i)$. Therefore, we obtained a set of elementary transformations $\{\phi_i\}_{i=1 \dots K}$. These elementary transformations generate a basis that can be used to approximate and decompose every new deformation. Using such a decomposition, we can obtain a representation in small dimensions of every training volume X_i . These representations are obtained by the projection of $E_\psi(X_i)$ to each principal vector: $a_i^j = E_\psi(X_i) \cdot \vec{u}_j$. For every volume of our

training set we have the approximation: $E_\psi(X_i) \approx \sum_{j=1}^K \alpha_i^j \vec{u}_j$. After calculating the vector of the principal components \mathcal{U}_K with the training set, we projected each image of the validation set to obtain its PCA representation.

3.3 Implementation and Training Details

The Adam optimiser was used for our training, with a constant learning rate set to $1e^{-4}$, a batch size equal to 4 and 8 for lung and hippocampus, respectively. Our models were trained for 600 epochs, and it last approximately 4 and 9 hours for the lung and hippocampus dataset. Concerning data augmentation, we applied random flip, rotation, translation and zoom. Moreover, the weights of the different loss components were set to 1 except the loss for smoothness set to $\alpha = 0.1$ for both datasets and the weight for the jacobian loss β that was discarded for the hippocampus dataset. During the training process, we registered random pairs of different patients. Our training has been performed using the framework PyTorch and one GPU card Nvidia Tesla V100 with 32G memory. The PCA decomposition was calculated using the library scikit-learn, and the number of principal components K was set to 32. Using 32 components, our decomposition covered 95% and 93% of the variance ratio for the lung and hippocampus dataset, respectively, while 42% and 62% are covered by the first four components for each dataset, respectively.

4 Experiments and Results

We performed our experiments on two different datasets, one public and one private. Starting with the public dataset, we conduct experiments with the hippocampus⁵ [23]. This dataset comprises 394 MRI with the segmentations of two small structures, the head and the body hippocampus. The images have been cropped around the hippocampus into small patches of $64 \times 64 \times 64$ voxels. The second dataset is composed of 41 lung MRI patients (12 healthy and 29 diseased with pulmonary fibrosis) together with their lung segmentations. Each patient had been acquired in two states, the inspiration and the expiration. Each volume has been resampled to 1.39mm on the x and z-axis and 1.69 on the y-axis and cropped to $128 \times 64 \times 128$ volumes. The same normalisation strategy has been applied for the two datasets: $\mathcal{N}(0, 1)$ standardisation, clip to $[-5, 5]$ to remove outliers values and min-max normalisation to $(0, 1)$. Both datasets were split into training and validation, resulting in 200 and 60 patients for hippocampus and 28 and 13 patients for the lung dataset.

As the first step of our evaluation, we benchmarked the performance of the registration network g_θ , on which our decomposition is based on. More specifically, we obtained a Dice coefficient of 0.90 ± 0.04 for the lungs and 0.76 ± 0.05 for the hippocampus, while the initial unregistered cases reported a Dice of 0.74 ± 0.14 and 0.59 ± 0.15 respectively. Moreover, we calculated the registration

⁵ <http://medicaldecathlon.com/>




Fig. 2: Visualisation of the displacements following the first four principal components. For each component, we depicted coronal and sagittal views. In red, the contours of the lungs of the M , and in gold, the $\mathcal{W}(D_\omega(\lambda\vec{u}_i), M)$ lung's contours. The deformation field is represented with arrows. The arrows' norm is represented with a colour map, red being the smallest and white the largest. Other patients and components are displayed on supplementary materials.

for g_θ with the skip connections to measure their impact on the registration. The Dice is then equal to 0.92 ± 0.02 and 0.85 ± 0.03 respectively. Thus, by removing the skip-connections, we decrease the performance of the registration, slightly on the lungs, more importantly, on the hippocampus. However, both strategies register the pair of volumes properly.

4.1 Qualitative Evaluation

To understand and evaluate the calculated components of \mathcal{U}_K per dataset, we perform a qualitative analysis. In particular, for each principal vector \vec{u}_i , we calculated the corresponding deformation $\{\phi_i\}_{i=1\dots K}$ and we applied to the moving image M together with its corresponding segmentation map M_{seg} . More formally, the deformed contour correspond to $\mathcal{W}(D_\omega(\lambda\vec{u}_i), M_{seg})$ with \mathcal{W} being the warping operation and λ the parameter to control the strength of the displacements for better visualisation.

In Figure 2, we show the principal components obtained for one validation subject for the lung dataset. Interestingly, one can observe that each ϕ_i corresponds to a different elementary transformation. More precisely, the 1st component is associated with translation, the 2nd with a deformation focusing on the bottom of the lungs, the 3rd with a deformation on the right lung focusing also on the heart region and lastly the 4th with a deformation focusing on the top region of the lung and shoulders.




Fig. 3: Representation of the deformed MR together with its lung contours following the first and fourth principal components u_1, u_4 . The red contour represents the position of the lung segmentation of the input image while the gold contour the position of the deformed lung. The values of lambda range from: $-200, -100, 0, 100$ and 200 (left to right). Negative values of lambda correspond to an upward translation, while positive values to a downward translation.

In Figure 3, we show the effect of the values of λ . In the figure, we present the lung contours of the scaled component (in red) and the corresponding component of the warped of the first ad third components. As we have indicated, the 1st component is associated with translation, which we can also be observed in this visualisation. In particular, for this experiment we sample λ from the values $\{-200, -100, 0, 100, 200\}$. One can observe that we retrieve a near identity deformation for a value of 0, while for negative and positive values, the lung moves up and down, respectively. On the other hand, the fourth component is responsible for deforming the shoulders and the top of the lungs. Indeed, one can observe that through the different λ values, the top lungs region is the one that reports the most changes. In Figure 7, similarly, the 4 deformations produced by the first 4 principal components of the hippocampus dataset are presented. In this case, the 1st component seems to capture rotation on the sagittal plane,




Fig. 4: Visualisation of the displacements following the first four principal components. We depicted a sagittal view of one patient of the validation set of the hippocampus dataset. We represented the ground truth hippocampus contours (red) and the deformed one (gold), following the principal components.

the 2nd translation and shrinking towards the bottom right, while the 3rd seems to be the same operation towards the top left corner. Finally, the 4th seems to be related to scaling, inflating both the hippocampus’s head and tail. We observed that the decomposition of the two datasets created different elementary transformations ϕ_i , with transformations closer to affine for the hippocampus and more complex for the lung.

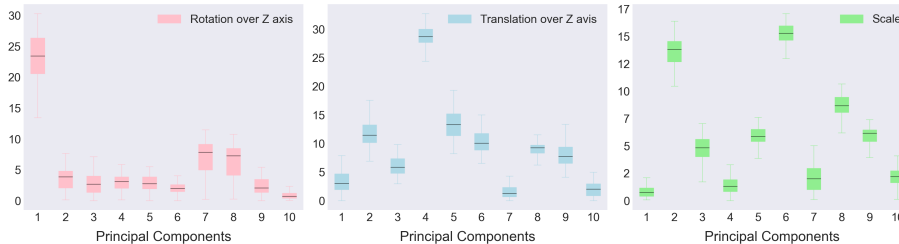


Fig. 5: Visualisation of the differences between the components of a reference image and the same image to which we applied a predefined transformation. The first ten components have been displayed. From right to left: rotation, translation along Z-axis and scaling.

Finally, to verify the obtained decomposition, we performed a case study for all the validation subjects of the hippocampus dataset. More specifically, we applied some predefined translation using 10 pixels on the z axis, rotation using 20 degrees on the z axis and scaling using a factor of 0.2, transforming each subject X to X' . Then we calculated the difference between the projection of $E_\psi(X)$ and $E_\psi(X')$ on the PCA decomposition. In Figure 5, a box plot for all the validation subjects of the absolute difference is presented. Specifically, the amount $\|a_{E_\psi(X)}^j - a_{E_\psi(X')}^j\|$ is shown for each principal component j , with $a_{E_\psi(X)}^j$ being the projection of $E_\psi(X)$ on the principal vectors \mathcal{U}_K , for the three different applied deformations. One can observe that for rotation and translation, only one component is significantly different from the rest. In the case of scaling, however, two components seem to be more activated. Moreover, these findings are in accordance with Figure 7 for the rotation and translation. In supplementary materials, we upgraded the Figure 5 by comparing the network with and without skip-connections. Contrary to our proposed formulation, many components are activated with the skip-connections, demonstrating the necessity of removing them to have a good decomposition.

5 Discussion & Conclusion

In this work, we proposed an approach to decompose and explain the representations of deep learning-based registration methods. The proposed method

utilises a linear decomposition on the latent space projecting it to principal components closely associated with anatomically aware deformations. Our method’s dynamics are demonstrated in two different MRI datasets, focusing on lung and hippocampus anatomies. We hope that these results will take some steps towards a better understanding of latent representations learned by the deep learning registration architectures. We also explored a direct application of the PCA on the deformation’s grid instead of the latent representation. However, we did not observe any qualitative correlations with types of deformations, which is the case for our proposed formulation. One of the main limitations of our approach is the difficulty of quantitative evaluation. Our future steps include the more extensive evaluation of our method, including new anatomies such as abdominal volumes and its clinical significance. More specifically, we want to apply our approach to multi-temporal follow-up of patients, monitoring diseases’ progression.

Funding

This work has been partially funding by the ARC: Grant SIGNIT201801286, the Fondation pour la Recherche Médicale: Grant DIC20161236437, SIRIC-SOCRATE 2.0, ITMO Cancer, Institut National du Cancer (INCa) and Amazon Web Services (AWS).

References

1. Balakrishnan, G., Zhao, A., Sabuncu, M.R., Guttag, J., Dalca, A.V.: Voxelmorph: a learning framework for deformable medical image registration. *IEEE transactions on medical imaging* **38**(8), 1788–1800 (2019)
2. Castro, D.C., Walker, I., Glocker, B.: Causality matters in medical imaging. *Nature Communications* **11**(1), 1–10 (2020)
3. Chattopadhyay, A., Sarkar, A., Howlader, P., Balasubramanian, V.N.: Gradcam++: Generalized gradient-based visual explanations for deep convolutional networks. In: 2018 IEEE Winter Conference on Applications of Computer Vision (WACV). pp. 839–847. IEEE (2018)
4. Chen, X., Duan, Y., Houthoofd, R., Schulman, J., Sutskever, I., Abbeel, P.: InfoGAN: Interpretable representation learning by information maximizing generative adversarial nets. *arXiv preprint arXiv:1606.03657* (2016)
5. Çiçek, Ö., Abdulkadir, A., Lienkamp, S.S., Brox, T., Ronneberger, O.: 3d u-net: learning dense volumetric segmentation from sparse annotation. In: International conference on medical image computing and computer-assisted intervention. pp. 424–432. Springer (2016)
6. Dalca, A.V., Balakrishnan, G., Guttag, J., Sabuncu, M.R.: Unsupervised learning for fast probabilistic diffeomorphic registration. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 729–738. Springer (2018)
7. Estienne, T., Lerousseau, M., Vakalopoulou, M., Alvarez Andres, E., Battistella, E., Carré, A., Chandra, S., Christodoulidis, S., Sahasrabudhe, M., Sun, R., et al.: Deep learning-based concurrent brain registration and tumor segmentation. *Frontiers in computational neuroscience* **14**, 17 (2020)
8. Estienne, T., Vakalopoulou, M., Battistella, E., Carré, A., Henry, T., Lerousseau, M., Robert, C., Paragios, N., Deutsch, E.: Deep Learning Based Registration Using Spatial Gradients and Noisy Segmentation Labels. In: Segmentation, Classification, and Registration of Multi-Modality Medical Imaging Data. pp. 87–93. Lecture Notes in Computer Science, Cham (2021)
9. Estienne, T., Vakalopoulou, M., Christodoulidis, S., Battistella, E., Lerousseau, M., Carre, A., Klausner, G., Sun, R., Robert, C., Mougiakakou, S., et al.: U-resnet: Ultimate coupling of registration and segmentation with deep nets. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 310–319. Springer (2019)
10. Fong, R.C., Vedaldi, A.: Interpretable explanations of black boxes by meaningful perturbation. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 3429–3437 (2017)
11. Heinrich, M.P.: Closing the gap between deep and conventional image registration using probabilistic dense displacement networks. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 50–58. Springer (2019)
12. Holzinger, A., Langs, G., Denk, H., Zatloukal, K., Müller, H.: Causability and explainability of artificial intelligence in medicine. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery* **9**(4), e1312 (2019)
13. Jaderberg, M., Simonyan, K., Zisserman, A., Kavukcuoglu, K.: Spatial Transformer Networks. *arXiv:1506.02025 [cs]* (Feb 2016)
14. Kingma, D.P., Welling, M.: Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114* (2013)

15. Krebs, J., Delingette, H., Mailhé, B., Ayache, N., Mansi, T.: Learning a probabilistic model for diffeomorphic registration. *IEEE transactions on medical imaging* **38**(9), 2165–2176 (2019)
16. Kuang, D., Schmah, T.: FAIM – A ConvNet Method for Unsupervised 3D Medical Image Registration. arXiv:1811.09243 [cs] (Jun 2019)
17. Mok, T.C., Chung, A.: Fast symmetric diffeomorphic image registration with convolutional neural networks. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 4644–4653 (2020)
18. Ou, Y., Weinstein, S.P., Conant, E.F., Englander, S., Da, X., Gaonkar, B., Hsieh, M.K., Rosen, M., DeMichele, A., Davatzikos, C., et al.: Deformable registration for quantifying longitudinal tumor changes during neoadjuvant chemotherapy. *Magnetic resonance in medicine* **73**(6), 2343–2356 (2015)
19. Qin, C., Shi, B., Liao, R., Mansi, T., Rueckert, D., Kamen, A.: Unsupervised Deformable Registration for Multi-modal Images via Disentangled Representations. In: Chung, A.C.S., Gee, J.C., Yushkevich, P.A., Bao, S. (eds.) *Information Processing in Medical Imaging*. pp. 249–261 (2019)
20. Sahasrabudhe, M., Shu, Z., Bartrum, E., Alp Guler, R., Samaras, D., Kokkinos, I.: Lifting autoencoders: Unsupervised learning of a fully-disentangled 3d morphable model using deep non-rigid structure from motion. In: Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops. pp. 0–0 (2019)
21. Selvaraju, R.R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., Batra, D.: Grad-cam: Visual explanations from deep networks via gradient-based localization. In: Proceedings of the IEEE international conference on computer vision. pp. 618–626 (2017)
22. Shu, Z., Sahasrabudhe, M., Guler, R.A., Samaras, D., Paragios, N., Kokkinos, I.: Deforming autoencoders: Unsupervised disentangling of shape and appearance. In: Proceedings of the European conference on computer vision (ECCV). pp. 650–665 (2018)
23. Simpson, A.L., Antonelli, M., Bakas, S., Bilello, M., Farahani, K., Van Ginneken, B., Kopp-Schneider, A., Landman, B.A., Litjens, G., Menze, B., et al.: A large annotated medical image dataset for the development and evaluation of segmentation algorithms. arXiv preprint arXiv:1902.09063 (2019)
24. Springenberg, J.T., Dosovitskiy, A., Brox, T., Riedmiller, M.: Striving for simplicity: The all convolutional net. arXiv preprint arXiv:1412.6806 (2014)
25. Stergios, C., Mihir, S., Maria, V., Guillaume, C., Marie-Pierre, R., Stavroula, M., Nikos, P.: Linear and Deformable Image Registration with 3D Convolutional Neural Networks. In: *Image Analysis for Moving Organ, Breast, and Thoracic Images*. pp. 13–22. Lecture Notes in Computer Science, Springer International Publishing, Cham (2018)
26. Zhang, S., Liu, P.X., Zheng, M., Shi, W.: A diffeomorphic unsupervised method for deformable soft tissue image registration. *Computers in Biology and Medicine* **120**, 103708 (May 2020). <https://doi.org/10.1016/j.compbimed.2020.103708>
27. Zhou, B., Khosla, A., Lapedriza, A., Oliva, A., Torralba, A.: Learning deep features for discriminative localization. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2921–2929 (2016)

Supplementary Material




Fig. 6: Extension of the figure 2 with two others patients and the component 1 to 4.




Fig. 7: Extension of the Figure 4 with two other patients and the components 1 to 8 (first row 1-4, second row 5-8).

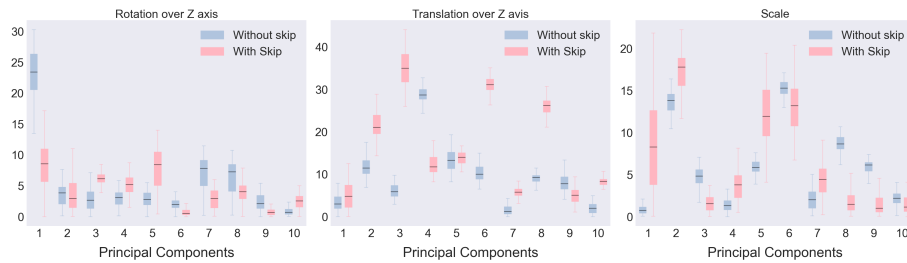


Fig. 8: Comparison between the networks with and without skip connections. We displayed the difference between the components of an image and the same image to which we applied a predefined transformation. Only one or two components are activated without the skip-connections, while many of them are with the skip. The first ten components have been displayed. The results are in blue and pink for respectively without and with the skip-connections. From right to left: rotation, translation along Z-axis and scaling.