



HAL
open science

A Deep Residual Learning Implementation of Metamorphosis

Matthis Maillard, Anton François, Joan Glaunès, Isabelle Bloch, Pietro Gori

► **To cite this version:**

Matthis Maillard, Anton François, Joan Glaunès, Isabelle Bloch, Pietro Gori. A Deep Residual Learning Implementation of Metamorphosis. IEEE ISBI 2022, Mar 2022, Kolkata, India. hal-03552709

HAL Id: hal-03552709

<https://telecom-paris.hal.science/hal-03552709v1>

Submitted on 2 Feb 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.

A DEEP RESIDUAL LEARNING IMPLEMENTATION OF METAMORPHOSIS

Matthis Maillard*

Anton François*[‡]

Joan Glaunès[‡]

Isabelle Bloch*[†]

Pietro Gori*

* LTCI, Télécom Paris, Institut Polytechnique de Paris, Paris, France

[†] Sorbonne Université, CNRS, LIP6, Paris, France

[‡] MAP5, Université de Paris, France

ABSTRACT

In medical imaging, most of the image registration methods implicitly assume a one-to-one correspondence between the source and target images (*i.e.*, diffeomorphism). However, this is not necessarily the case when dealing with pathological medical images (*e.g.*, presence of a tumor, lesion, etc.). To cope with this issue, the *Metamorphosis* model has been proposed. It modifies both the shape and the appearance of an image to deal with the geometrical and topological differences. However, the high computational time and load have hampered its applications so far. Here, we propose a deep residual learning implementation of *Metamorphosis* that drastically reduces the computational time at inference. Furthermore, we also show that the proposed framework can easily integrate prior knowledge of the localization of topological changes (*e.g.*, segmentation masks) that can act as spatial regularization to correctly disentangle appearance and shape changes. We test our method on the BraTS 2021 dataset, showing that it outperforms current state-of-the-art methods in the alignment of images with brain tumors.

Index Terms— *Metamorphosis*, Image registration, Deep learning, Brain, Tumors.

1. INTRODUCTION

Image registration has many applications in medical imaging, such as statistical analysis, modality fusion, surgery planning and follow up. Most of the methods assume that the two images have the same topology (*i.e.*, same number of anatomical components) thus looking for a diffeomorphism [1, 2]. However, in some studies one needs to align images characterized by a different number of anatomical components, like a healthy image and an image with a tumor [3].

One of the first strategies to deal with a difference in topology between source and target images was the cost function masking, where the lesion was excluded when computing the image similarity [4]. However, when the lesion is too big, the deformation in the masked zone might be very distorted. To cope with that, geometric metamorphosis [5] adds a specific deformation to the masked area, but it works only when the lesion/tumor is present in both source and target images. Another strategy, when one wants to register two images with and without tumor, is to first model the growth of the tumor in the healthy image by using a biophysical model [6, 7], and then do the registration. This strategy deals with the topological difference, since the images now contain the same number of components, but it requires a user initialization, extensive computations to estimate the model parameters, and they are specific to a particular kind of tumor. In [8], authors use a similar but inverse perspective where, instead of adding a tumor to a healthy image, they remove the pathological region by synthesizing a quasi-normal image via low-rank approaches. This approach can effectively recover tumor regions,

but at the same time distort or blur the healthy regions. Furthermore, it is a statistical technique that needs lesions to be homogeneously (and randomly) distributed across the population [8], which is not the case for all kinds of lesions or tumors (*e.g.*, glioblastoma). With a similar perspective, inpainting techniques have also been proposed [9]. However, they may produce inaccurate or unreal images when dealing with large tumors or lesions. *Metamorphosis* [10] is a more generic approach that offers theoretical guarantees to tackle both shape and appearance differences. It consists in repetitively applying (together) infinitesimal shape and appearance deformations to the source image. To the best of our knowledge, there is no clinical application of this method and a recent implementation [11] indicates that it is very time consuming.

Recently, motivated by faster computation time at inference, several deep learning approaches have been developed for image registration. Neural networks are trained on a database of image pairs and not only for a fixed pair of images [12]. The training is thus very time consuming but this type of methods is much faster at inference time than previous methods that need to optimize a functional for every pair of images. Similarly, deep learning can also be used to align images with different topology. In [13], authors proposed a *Metamorphic* (variational) Auto-Encoder (MAE) to modify both the geometry and the appearance of an image at the same time. However, this method may result in poor disentanglement, where the intensity changes can actually modify the shape of the input image.

Inspired by [14], we propose a ResNet-based [15] implementation of *Metamorphosis* which overcomes all previous limitations. Our contributions are:

- We use a residual deep network [15] to solve the system of differential equations of *Metamorphosis* (see Eq. 3 next).
- Our method can be used: 1) to optimize an energy between two paired images, or 2) in a learning-based scheme where we use a training set of source images and one fixed target image.
- We introduce a local regularization that consists in limiting the intensity changes to a pre-specified zone (segmentation mask). We show that this produces a better disentanglement between shape and appearance changes.
- We evaluate our method on a synthetic shape dataset and on the BraTS 2021 dataset [16].

2. METHOD

Mathematical formulation. Let $\Omega \subset \mathbb{R}^d$ be a bounded domain, with $d \in \{2, 3\}$. Let V be a Reproducible Kernel Hilbert space (RKHS) with kernel K of vector fields with support Ω and T times continuously differentiable, where $T \in \mathbb{N}^*$. Let I be a gray scale image defined as a square integrable and differentiable function $I : \Omega \rightarrow \mathcal{R}$. The aim of *Metamorphosis* is to modify a source image I

so that it perfectly aligns with a target image J . The model joins diffeomorphic deformations with additive intensity changes. Similarly to [10, 11], the evolution of the image I at time $t \in [0, 1]$ is:

$$\partial_t I_t = v_t \cdot I_t + \mu^2 z_t = -\langle \nabla I_t, v_t \rangle + \mu^2 z_t \quad \text{s.t. } I_0 = I \quad (1)$$

where $v_t \cdot I_t$ implies that I_t is deformed by an infinitesimal vector field $v_t \in V$ and $z_t : \Omega \rightarrow \mathbb{R}$ is the additive part corresponding to the infinitesimal intensity variation (called the residual image or momentum). The hyperparameter $\mu^2 \in \mathbb{R}^+$ balances the intensity and geometric changes. As in [11], we cast the metamorphic registration as an inexact matching problem minimizing the cost function:

$$E = \frac{1}{2} \|I_1 - J\|_2^2 + \lambda \left[\int_0^1 \|v_t\|_V^2 + \|\mu z_t\|_2^2 dt \right] \quad (2)$$

where the first term is the classical L_2 data term (please note that other data terms could be used as well) and the second term, weighted by λ , is the regularization. It is composed of the total geometric and intensity energy of the deformation, respectively. As shown in [11], the geodesic equations for Metamorphosis are:

$$\begin{cases} v_t = -K * (z_t \nabla I_t) & (3.a) \\ \partial_t z_t = -\nabla \cdot (z_t v_t) & (3.b) \\ \partial_t I_t = -\langle \nabla I_t, v_t \rangle + \mu^2 z_t & (3.c) \end{cases}$$

with $\|v_t\|_V^2 = \langle z_t \nabla I_t, K * (z_t \nabla I_t) \rangle$, where K is a Gaussian kernel. From this system of equations, we can notice that v_t is completely defined by z_t and I_t , thus making z_t the only unknown. The momentum z_t has therefore a double role. It represents the additive intensity variation *and* it is also the parameter of the deformation. This eases the computation but at the same time it makes the disentanglement between shape and intensity variations more difficult (see Fig. 3). Inspired by [14, 17], we propose to use a residual neural network (ResNet) to find the solution of this system. We take advantage of the similarity between ResNets and the numerical solutions of PDEs using Euler's method, given an initial value, to solve Eq. 3.b. Indeed, the numerical integration of Eq. 3.b, using discrete time steps t , is:

$$z_{t+1} = z_t - \frac{1}{T} \nabla \cdot (z_t v_t) \quad \text{for } t = i/(T-1), i \in 0, \dots, T-1 \quad (4)$$

where T is the number of steps. By replacing the divergence with a neural network, we obtain a ResNet:

$$z_{t+1} = z_t + \frac{1}{T} f_{\theta_t}(z_t, I_t) \quad (5)$$

where f_{θ_t} is modeled as a convolutional layer followed by an activation function (Leaky ReLU) and two other convolutional layers. Compared with Eq. 3.b, v_t is replaced by I_t because v_t is a function of I_t (and z_t). The network is built as a sequence of T convolution blocks f_{θ_t} . At each time step t , z_{t+1} is computed using Eq. 5. Subsequently, v_{t+1} is calculated directly with Eq. 3.a and one determines I_{t+1} by applying the geometric transformation induced by v_t and adding the residuals z_t as in Eq. 3.c. The architecture of the model is detailed in Fig. 1.

Optimization setting. The proposed method can be used in two different contexts: optimization or learning-based. In the former, one wants to compute the transformation between two fixed images. In this case, the back-propagation algorithm is repetitively applied to minimize the following equation until convergence:

$$E_O(\theta, z_0) = \frac{1}{2} \|I_T - J\|_2^2 + \frac{\lambda}{T} \sum_{t=0}^{T-1} [\|v_t\|_V^2 + \mu^2 \|z_t\|_2^2] \quad (6)$$




Fig. 1. Architecture of the proposed method.

The parameters to optimize are the weights of the neural network θ and the initial residuals z_0 . We use Adam optimizer with a learning rate of 10^{-3} . In this context, we have empirically measured that the first layers are similar to the divergence in Eq. 3.b: their L_1 distance is small (between 0 and 1) and the Pearson's correlation coefficient is around 0.8. However, this is not the case for the last layers, where the L_1 distance is high (around 40) and the correlation is close to 0. Future research will dive more into that, trying to understand whether the use of a ResNet may give more generic and/or well adapted results.

In the learning-based setting, the goal is to compute the transformations between every image of a training dataset of N images and a fixed target image. The optimized parameters are the same as in the previous context, the difference is that from one iteration to the other, the source image is not the same. Therefore, one needs to minimize the following energy $E_L(\theta, z_0)$:

$$\sum_{n=1}^N \left[\frac{1}{2} \|I_T^n - J\|_2^2 + \frac{\lambda}{T} \left[\sum_{t=0}^{T-1} \|v_t^n\|_V^2 + \mu^2 (\|z_0\|_2^2 + \sum_{t=1}^{T-1} \|z_t^n\|_2^2) \right] \right] \quad (7)$$

where n is the index of image I^n in the dataset. Please note that the learning setting is possible only because we use a neural network in Eq.3.b. The ResNet learns the optimal network parameters θ and initial residuals z_0 to align *the entire* distribution of source images onto a (fixed) target image. Once the network is trained, it can be used to register any image of the same type as those in the training set to the target image. Differently from the optimization setting, this allows for the registration of a high number of images in a very short time since the inference time of the ResNet is definitely lower compared to the optimization of Eq. 6 (less than one second vs one minute for images of size 200×200). This is particularly important, for instance, when one needs to normalize or align a large set of images to a common template for a statistical analysis.

Local regularization. The main inconvenience with Metamorphosis is that it is hard to control the disentanglement between shape and appearance. For instance, a trivial solution would be to set the overall geometrical deformation function to the identity (no geometrical change) and the overall appearance deformation map to $J - I_0$. In that case, the L_2 distance between the deformed image and J would be 0 but it would not be a satisfactory result since homologous structures should be matched using only geometric deformations whereas appearance and topological changes (*i.e.*, new components) should be taken into account by the intensity modifications. The disentanglement can be controlled by tuning the hyperparameters μ and λ . However, finding the right ones is a difficult task and they are different for each setting. If they are not correctly chosen, the appearance map could, for instance, modify the shape of the image, thus distorting the results and their interpretations.

To improve the control of the disentanglement, we introduce a

local regularization which consists in limiting the region where the intensities can be modified (e.g. a tumor). To do so, we define such a region as $\alpha \subset \Omega$. We introduce a mask m_0 where $m_0(x) = 1$ if the pixel $x \in \alpha$ and 0 otherwise. During the deformation, the shape of the tumor will be modified, therefore, the mask must undergo the same deformation. Hence, we have $\partial_t m_t = v_t \cdot m_t$. The transformation of the image is then: $\partial_t I_t = v_t \cdot I_t + \mu^2 m_t z_t$. Using this equation and $\|\sqrt{m_t} z_t\|_2^2$ as regularization term for z_t , we obtain the same geodesic Equations 3.a and 3.b.

3. EXPERIMENTS

Datasets. We first evaluate the effectiveness of the method on a synthetic dataset of 2000 images of size 200×200 . It is built from an image of a white “C” on a black background. Each image of the dataset is generated by applying a random elastic deformation to the “C” image. The target image associated with this dataset is a “C” that has been cut in the middle, as shown in Fig. 2.

The second dataset is BraTS 2021 [16] comprising 4 MR modalities and the associated tumor segmentation image for 1251 brains with tumor. For each patient, we select the same slice of the T1 modality and crop it to obtain an image of size 208×208 . We randomly pick 40 images from the dataset to form an evaluation set. The target image associated with this dataset is the linear MNI152 template [18]. To show that our model also works with other modalities, we select a slice from the T1 contrast enhanced (T1ce) modality of two different patients. We extract one slice outside the tumor and the other one exhibiting part of the tumor (see Fig. 3). We use the segmentation of the tumor as mask for the local regularization.

Results All experiments with our model have been computed using $T = 20$ and $\lambda = 10^{-6}$. The effectiveness of the method is first evaluated on the synthetic dataset. For comparison, we select optimization based methods: LDDMM [2] and cost function masking [4], and learning based methods: Voxelmorph [12] and Metamorphic auto-encoders (MAE) [13]. They are compared with our optimization and learning based method respectively. Voxelmorph and our learning model are trained with the “C” dataset as source and the fixed target in Fig. 2. The source image in Fig. 2 is *not* included in the training set. MAE works the other way around, the source image is fixed and the target changes. Therefore, we create a dataset of images made of elastic deformations of the target image in Fig. 2 and train MAE on this dataset. Similarly, the target image in Fig. 2 is *not* included in the training set. Thus, the three learning methods are comparable for the couple of images in Fig. 2. The visual and quantitative results show that, in both contexts, our method outperforms the others. As shown by the arrows, our model can effectively deal with the topological change and correctly aligns the remaining part.

As already discussed, disentanglement between shape and appearance is hard to control. To illustrate this, we show the deformations between two T1ce images for several values of μ without regularization in Fig. 3. For $\mu = 0.035$, there is almost no geometrical deformation of the source, yet the final image is really close to the target. This indicates that μ is too big since all differences have been removed only by the appearance changes. For the lower values of μ : 0.015 and 0.025, the shape deformation is better but still unsatisfactory. As pointed out by the red arrow, the ventricles of the shape deformation image are not aligned with the target ones. However, the appearance deformation is also unsatisfactory because the tumor has not fully disappeared on both deformations. This shows that finding the optimal parameters to properly align two images *without* regularization might be hard and time-consuming. By contrast, when using local regularization with the same μ values




Fig. 2. Registration between Source and Target images for LDDMM [2], Cost Function Masking (CFM) [4], ours in an optimization setting, Voxelmorph [12], Metamorphic Auto-Encoders (MAE) [13] and ours in a learning context. Plots refer to the absolute value of the difference between the corresponding output and the target. SSD refers to the sum of squared differences. Top and bottom lines show the optimization and learning based methods, respectively. Arrows point to the topological difference between Source and Target image.

(here we only show $\mu = 0.025$), the model is able to properly align the two images. The shape deformation image, where we only deform the image without adding the intensity changes, shows that the ventricles have been correctly deformed. In the total deformation image, where we consider both geometric deformations and intensity changes, the tumor has disappeared. Hence, with the proposed local regularization, shape and appearance are properly disentangled and the method is more robust to the choice of μ , since the results do not vary much when changing it.

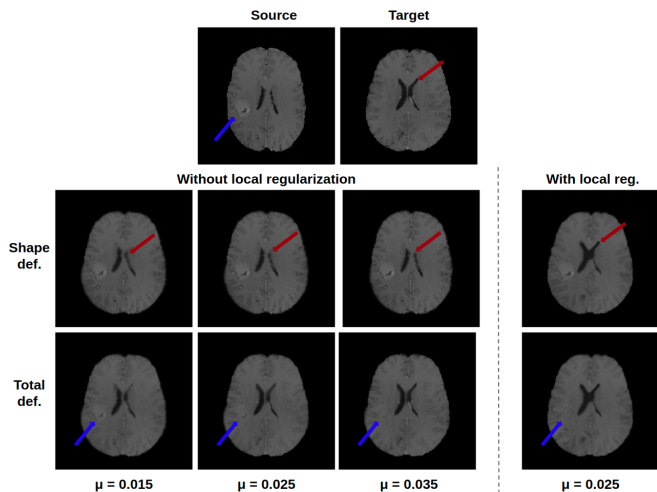


Fig. 3. Comparison of the deformation between two T1ce images with and without local regularization. The Shape def. line only shows the geometrical deformation of the source, while Total def. also includes the estimated appearance changes. The tumor segmentation is used as mask m_0 for the local regularization.

Eventually, to show a potential application of the method in a clinical setting, we align a set of T1 images onto the MNI template.




Fig. 4. Deformation of a T1 image onto the MNI atlas.

An example is shown in Fig. 4. To evaluate the quality of the alignment, we segment the ventricles of the deformed image and compute the Dice score with the reference segmentation of the ventricles in the MNI atlas (see Table 1). As segmentation algorithm, we trained a U-Net model on 2D slices from the OASIS dataset [19]. Reference segmentations were obtained using FreeSurfer. The trained U-Net showed an average Dice score of 0.91 on the validation set. We compare our method, with and without local regularization ($\mu = 0.015$), against cost function masking [4] and Voxelmorph [12]. The presence of the tumor mainly modifies the shape of the ventricles (mass effect), but not their intensity. All methods should therefore correctly match the ventricles using only geometric transformations. For this reason, when using our methods, we compute the Dice score on the shape deformation images and not on the total deformation images. In Table 1, we also show the SSD and inference time. Results clearly show that our method with local regularization better aligns the ventricles than the other methods and has the best SSD. This indicates that taking into account the topological changes during the deformation, rather than just masking them out or ignoring them, improves the alignment of homologous components between the two images.

Method	SSD	Dice	Inference Time (s)
CFM	167 ± 68	69.1 ± 22.4	55 ± 1.2
Voxelmorph	187 ± 73	63.4 ± 22.6	0.01 ± 8e-4
Ours w/o reg (learn.)	220 ± 20	65.4 ± 20.7	0.15 ± 0.003
Ours w/ reg (learn.)	161 ± 16	69.8 ± 15.2	0.16 ± 0.009
Ours w/ reg (opt.)	135 ± 45	71.9 ± 19.5	59 ± 1.2

Table 1. Quantitative evaluation for Cost Function Masking (CFM), Voxelmorph and our method with and without regularization. learn. (respectively opt.) indicates a learning context (respectively optimization context). Results were computed on a test set of 40 patients. Best results in bold.

4. CONCLUSION

This paper proposes a deep residual learning implementation of Metamorphosis and adds a local regularization to improve the shape and appearance disentanglement. Qualitative and quantitative results on BRATS dataset show the effectiveness of the method. Moreover, the model can be used in an optimization or learning context, the latter offering very fast computation time during inference.

Compliance with ethical standards This research was conducted retrospectively using human subject data made available in open access (OASIS and BRATS 2021 databases).

Acknowledgments Authors report no conflict of interest. M. Maillard was supported by a grant of IMT, Fondation Mines-

Télécom and Institut Carnot TSN.

5. REFERENCES

- [1] J. Ashburner, “A fast diffeomorphic image registration algorithm,” *NeuroImage*, vol. 38, no. 1, pp. 95–113, Oct. 2007.
- [2] M.F. Beg et al, “Computing Large Deformation Metric Mappings via Geodesic Flows of Diffeomorphisms,” *IJCV*, vol. 61, no. 2, pp. 139–157, 2005.
- [3] A. Roux et al, “MRI atlas of IDH wild-type supratentorial glioblastoma: Probabilistic maps of phenotype, management, and outcomes,” *Radiology*, vol. 293, no. 3, pp. 633–643, 2019.
- [4] M. Brett et al, “Spatial Normalization of Brain Images with Focal Lesions Using Cost Function Masking,” *NeuroImage*, vol. 14, no. 2, pp. 486–500, Aug. 2001.
- [5] M. Niethammer et al, “Geometric Metamorphosis,” in *MICCAI*, 2011, vol. LNCS 6892, pp. 639–646.
- [6] A. Gooya et al, “GLISTR: Glioma Image Segmentation and Registration,” *IEEE TMI*, vol. 31, no. 10, pp. 1941–1954, 2012.
- [7] K. Scheufele et al, “Coupling brain-tumor biophysical models and diffeomorphic image registration,” *Comput Methods Appl Mech Eng*, vol. 347, pp. 533–567, 2019.
- [8] X. Liu et al, “Low-Rank Atlas Image Analyses in the Presence of Pathologies,” *IEEE TMI*, vol. 34, pp. 2583–2591, 2015.
- [9] M. Sdika and D. Pelletier, “Nonrigid registration of multiple sclerosis brain images using lesion inpainting for morphometry or lesion mapping,” *Hum. Brain Mapp.*, vol. 30, no. 4, 2009.
- [10] A. Trouvé et al, “Metamorphoses Through Lie Group Action,” *Foundations of Computational Mathematics*, vol. 5, no. 2, pp. 173–198, 2005.
- [11] A. François et al, “Metamorphic image registration using a semi-Lagrangian scheme,” in *GSI*, 2021.
- [12] G. Balakrishnan et al, “VoxelMorph: A Learning Framework for Deformable Medical Image Registration,” *IEEE TMI*, vol. 38, no. 8, pp. 1788–1800, Aug. 2019.
- [13] A. Bône et al, “Learning Joint Shape and Appearance Representations with Metamorphic Auto-Encoders,” in *MICCAI*, 2020, pp. 202–211.
- [14] B. Ben Amor et al, “ResNet-LDDMM: Advancing the LDDMM Framework Using Deep Residual Networks,” *CoRR*, 2021.
- [15] K. He et al, “Deep Residual Learning for Image Recognition,” in *CVPR*, 2016, pp. 770–778, IEEE.
- [16] B.H. Menze et al, “The multimodal brain tumor image segmentation benchmark (BRATS),” *IEEE TMI*, vol. 34, no. 10, pp. 1993–2024, 2015.
- [17] F. Rousseau, L. Drumetz, and R. Fablet, “Residual Networks as Flows of Diffeomorphisms,” *JMIV*, vol. 62, no. 3, pp. 365–375, 2020.
- [18] J. C. Mazziotta et al, “A probabilistic atlas of the human brain: Theory and rationale for its development: ICBM,” *NeuroImage*, vol. 2, no. 2, Part A, pp. 89–101, 1995.
- [19] DS. Marcus et al, “OASIS: Cross-sectional MRI data in young, middle aged, nondemented, and demented older adults.,” *Journal of Cognitive Neuroscience*, vol. 19, pp. 1498–1507.