



HAL
open science

SegSRGAN: A software solution for super-resolution and segmentation using generative adversarial networks – Application to neonatal brain MRI

Quentin Delannoy, Chi-Hieu Pham, Clément Cazorla, Carlos Tor-Díez,
Guillaume Dollé, Hélène Meunier, Nathalie Bednarek, Ronan Fablet, Nicolas
Passat, François Rousseau

► To cite this version:

Quentin Delannoy, Chi-Hieu Pham, Clément Cazorla, Carlos Tor-Díez, Guillaume Dollé, et al.. SegSRGAN: A software solution for super-resolution and segmentation using generative adversarial networks – Application to neonatal brain MRI. 2019. hal-02189136v1

HAL Id: hal-02189136

<https://hal.science/hal-02189136v1>

Preprint submitted on 19 Jul 2019 (v1), last revised 7 Apr 2020 (v3)

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.

SegSRGAN: A software solution for super-resolution and segmentation using generative adversarial networks — Application to neonatal brain MRI

Quentin Delannoy^{a,b}, Chi-Hieu Pham^{a,c}, Clément Cazorla^b, Carlos Tor-Díez^c, Guillaume Dollé^d, Hélène Meunier^e, Nathalie Bednarek^{b,e}, Ronan Fablet^f, Nicolas Passat^b, François Rousseau^c

^aEqual first co-authors

^bUniversité de Reims Champagne Ardenne, CReSTIC EA 3804, 51097 Reims, France

^cIMT Atlantique, LaTIM U1101 INSERM, UBL, Brest, France

^dUniversité de Reims Champagne Ardenne, LMR FRE 2011, 51097 Reims, France

^eService de médecine néonatale et réanimation pédiatrique, CHU de Reims, France

^fIMT Atlantique, Lab-STICC UMR CNRS 6285, Brest, France

Abstract

One of the main issues in the analysis of clinical neonatal brain MRI is the low anisotropic resolution of the data. In most MRI analysis pipelines, data are first re-sampled using interpolation or single image super-resolution techniques and then segmented using (semi-)automated approaches. In other words, image reconstruction and segmentation are then performed separately. In this article, we propose a methodology and a software solution for carrying out simultaneously high-resolution reconstruction and segmentation of brain MRI data. Our strategy mainly relies on generative adversarial networks. The network architecture is described in details, such as the associated software tool. We illustrate its behaviour for cortex analysis from neonatal MR images, both in a quantitative way on a research MRI dataset, and more qualitatively on real clinical data. Results emphasize the potential of our proposed method / software with respect to practical medical applications.

Keywords: super-resolution, segmentation, 3D generative adversarial networks, neonatal brain MRI, cortex.

1. Introduction

Long-term studies of the outcome of prematurely born infants have clearly documented that the majority of such infants may have significant motor, cognitive, and behavioral deficits [15, 25]. However, there is a limited understanding of the nature of the cerebral abnormality underlying these adverse neurologic outcomes. In this context, Magnetic Resonance Imaging (MRI) provides unique opportunities for in vivo investigation of the early developing human brain. However, the analysis of these clinical neonatal brain MRI data remains challenging. This is mainly due to their low anisotropic resolution. As a consequence, improving morphological data processing such as image resolution enhancement and brain segmentation is the cornerstone for further providing robust morphometry analysis tools to the scientific and clinical communities.

When dealing with anisotropic low resolution images, one of the first key components of the processing pipeline of clinical MRI data is the upsampling image

estimation. Super-resolution (SR) is a post-processing technique that aims at enhancing the resolution of an imaging system [12]. However, SR is a challenging inverse problem; in particular the estimation of texture and details remains difficult. Recently, supervised deep learning-based techniques have shown great improvement over model-based approaches, for that purpose. Indeed, applying 3D convolutional neural networks (CNNs) yields promising results for MRI data [30, 7]. However, it has been shown that the use of ℓ_2 -norm loss leads to smooth, unrealistic high resolution images [17, 20]. In order to tackle these issues, generative adversarial networks (GANs) [11] have been considered for estimating textured and sharper images [20, 6].

Once a high resolution image reconstruction has been performed, the implementation of an automatic segmentation robust approach is crucial for brain structure analysis. Dedicated (in general, automatic) segmentation methods can be classified into four types (see e.g. [23] for a recent survey): unsupervised [29,

13]; parametric [22]; classification [32, 39, 26, 4, 24] / atlas fusion [33, 8, 18, 34]; and deformable models [36, 37]. In addition, GANs were also recently proposed for brain MRI segmentation purpose [27].

Among the structures of interest, thin ones and in particular the cortical gray matter, remain difficult to analyze. However, the cortex is a region of interest, as emphasized by recent works focusing e.g. on brain folding [10, 21, 28], cortical connectivity [2] and cortical development [3, 40]. Since the cerebral cortex is a thin surface object, it remains difficult to segment in neonatal MRI data. In particular, the segmentation step is generally considered separately from image reconstruction.

Our purpose is to propose an end-to-end methodology dedicated to the analysis of anisotropic low resolution MR images. In particular, we aim to deal with complex anatomical structures, and in particular the cortex. To this end, we propose a GAN-based approach, namely SegSRGAN, that generates both the perceptually super-resolved image and a cortical segmentation map from a single low-resolution (LR) MR image.

This article, which is an extended and improved version of the conference paper [31], is organized as follows. In Section 2, we describe our SR/segmentation framework from a methodological point of view. In Section 3, we provide technical details concerning our proposed implementation. In Section 4, we describe experiments carried out on various kinds of data, and we quantitatively and qualitatively assess the effects of various parameters. Section 5 provides concluding remarks on the main contributions and perspective works.

2. Methodological description

In this section, we focus on the methodological aspects of joint reconstruction and cortical segmentation of clinical LR neonatal brain MR images.

2.1. Formulation of single image super-resolution

The goal of single image SR is to estimate a high-resolution (HR) image $\mathbf{X} \in \mathbb{R}^m$ from an observed LR image $\mathbf{Y} \in \mathbb{R}^n$, with $m > n$. Such SR problem can be formulated using the following linear observation model:

$$\mathbf{Y} = H_{\downarrow} \mathbf{B} \mathbf{X} + N = \Theta \mathbf{X} + N \quad (1)$$

where $N \in \mathbb{R}^n$ is the additive noise, $B \in \mathbb{R}^{m \times m}$ is a blur matrix (depending on the point spread function), $H_{\downarrow} \in$

$\mathbb{R}^{n \times m}$ is a downsampling decimation matrix and $\Theta = H_{\downarrow} B \in \mathbb{R}^{n \times m}$.

A popular way of solving this SR problem is to define the matrix Θ^{-1} as the combination of a restoration operator $F \in \mathbb{R}^{m \times m}$ and an upscaling interpolation operator $S^{\uparrow} \in \mathbb{R}^{m \times n}$ that computes the interpolated LR image $\mathbf{Z} \in \mathbb{R}^m$ associated to \mathbf{Y} , i.e. $\mathbf{Z} = S^{\uparrow} \mathbf{Y}$. In the context of supervised learning, given a set of HR images \mathbf{X}_i and their corresponding LR images \mathbf{Y}_i , the restoration operator F can be estimated by minimizing the following loss function:

$$\hat{F} = \arg \min_F \sum_i \|\mathbf{X}_i - F(\mathbf{Z}_i)\|_2^2 \quad (2)$$

Once \hat{F} is estimated, given a LR image \mathbf{Y} , the computation of an HR image \mathbf{X} is straightforwardly expressed as $\mathbf{X} = \hat{F}(S^{\uparrow} \mathbf{Y})$. In [30, 7], it was shown that 3D CNNs could be used to estimate accurately the restoration mapping \hat{F} for brain MRI.

However, it is known that the use of a ℓ_2 -norm may lead to oversmoothing high resolution images. In order to provide realistic HR images, perceptual loss function [17] have been used in a GAN [20]. This constitutes a paradigm shift since it is no longer a question of minimizing only the reconstruction error but of estimating a realistic image, i.e. a high resolution image that corresponds to the observation model with realistic texture details.

A perceptual loss can be formulated as the weighted sum of the content loss (based e.g. on pixel-wise mean squared error loss) and an adversarial loss component. In GAN-based approaches, the purpose is to train a generating network G that estimates for a given LR input image a corresponding HR image. The goal of the discriminator network D is to classify real and simulated HR images.

2.2. Formulation of image segmentation

In this work, image segmentation is viewed as a supervised regression problem:

$$\mathbf{S}_{\mathbf{X}} = R(\mathbf{X}) \quad (3)$$

where R denotes a non-linear mapping from the up-scaled image \mathbf{X} to the segmentation map $\mathbf{S}_{\mathbf{X}}$. Similarly to the SR problem, assuming that we have a set of images \mathbf{X}_i and corresponding segmentation maps $\mathbf{S}_{\mathbf{X}_i}$, a general approach for solving this segmentation problem is to find the mapping R by minimizing the following loss function:

$$\hat{R} = \arg \min_R \sum_i \|\mathbf{S}_{\mathbf{X}_i} - R(\mathbf{X}_i)\|_2^2 \quad (4)$$

Contrary to the SR problem, the use of ℓ_2 -norm is less critical as it is expected to estimate smooth segmentation maps.

2.3. Joint mapping by generative adversarial networks

We now propose to use a GAN-based approach for estimating jointly both a HR image and its corresponding segmentation map from a LR image.

The considered loss function can be split into two different terms:

- a term \mathcal{L}_{rec} which directly measures the generator output image distance from the real one (reconstruction loss);
- a term \mathcal{L}_{adv} which expresses the game between generator and discriminator (i.e. adversarial loss).

The convolution-based generator network G takes as input an interpolated LR image \mathbf{Z} and computes a HR image $\widehat{\mathbf{X}}$ and a segmentation map $\widehat{\mathbf{S}}$ by minimizing the reconstruction loss:

$$\mathcal{L}_{rec} = \mathbb{E}_{\mathbf{X} \sim \mathbb{P}_{\mathbf{X}}, \mathbf{S}_{\mathbf{X}} \sim \mathbb{P}_{\mathbf{S}_{\mathbf{X}}}} [\rho((\mathbf{X}, \mathbf{S}_{\mathbf{X}}) - G(\mathbf{Z}))] \quad (5)$$

where the difference is computed voxel-wise, and ρ is the robust Charbonnier loss [5] ($x \in \mathbb{R}^n$):

$$\rho(x_1, \dots, x_n) = \sum_{i=1}^n \sqrt{(x_i^2 + \nu^2)} \quad (6)$$

with ν set here to 10^{-3} .

Traditionally, the adversarial function is modeled as minimax objective¹. However this loss can suffer from vanishing gradient, due to discriminator saturation for instance. Consequently, we propose to alternatively use Wasserstein GAN loss as described in [14]. Indeed, this type of GAN aims to minimize the Earth-Mover distance between two distribution \mathbb{P}_r and \mathbb{P}_g , defined as:

$$W(\mathbb{P}_r, \mathbb{P}_g) = \inf_{\gamma \in \Pi(\mathbb{P}_r, \mathbb{P}_g)} \mathbb{E}_{(x,y) \sim \gamma} [\|x - y\|] \quad (7)$$

$$= \sup_{\|f\|_L \leq 1} \mathbb{E}_{x \sim \mathbb{P}_r} [f(x)] - \mathbb{E}_{x \sim \mathbb{P}_g} [f(x)] \quad (8)$$

where, $\Pi(\mathbb{P}_r, \mathbb{P}_g)$ is the set of all the distributions of which marginals are \mathbb{P}_r and \mathbb{P}_g , respectively, and the

¹That is, by solving

$$\min_G \max_D V(D, G) = \mathbb{E}_{x \in p_{data}(x)} [\log D(x)] + \mathbb{E}_{z \in p_z(z)} [\log(1 - D(G(z)))]$$

The solution to this minimax problem can be interpreted as a Nash equilibrium, a concept from game theory.

supremum is calculated over all the 1-Lipschitz functions $f : A \rightarrow \mathbb{R}$

In this GAN, the discriminator learns the parametrized function f and the generator aims at minimizing this distance. Hence, in practice, the adversarial part of the loss function is:

$$\begin{aligned} \mathcal{L}_{adv} = & \mathbb{E}_{\mathbf{X} \sim \mathbb{P}_{\mathbf{X}}, \mathbf{S}_{\mathbf{X}} \sim \mathbb{P}_{\mathbf{S}_{\mathbf{X}}}} [D((\mathbf{X}, \mathbf{S}_{\mathbf{X}}))] \\ & - \mathbb{E}_{\mathbf{Z} \sim \mathbb{P}_{\mathbf{Z}}} [D(G(\mathbf{Z}))] \\ & - \lambda_{gp} \mathbb{E}_{\widehat{\mathbf{X}\mathbf{S}}} [(\|\nabla_{\widehat{\mathbf{X}\mathbf{S}}} D(\widehat{\mathbf{X}\mathbf{S}})\|_2 - 1)^2] \end{aligned} \quad (9)$$

where

$$\widehat{\mathbf{X}\mathbf{S}} = (1 - \epsilon)(\mathbf{X}, \mathbf{S}_{\mathbf{X}}) + \epsilon G(\mathbf{Z}) \quad (10)$$

is uniformly sampled between $(\mathbf{X}, \mathbf{S}_{\mathbf{X}})$ and $G(\mathbf{Z})$, whereas $\lambda_{gp} > 0$ and ∇ denote the gradient penalty coefficient and gradient operator, respectively.

The images \mathbf{X} , $\mathbf{S}_{\mathbf{X}}$ and \mathbf{Z} are extracted randomly from the data distributions of HR images $\mathbb{P}_{\mathbf{X}}$, HR segmentation maps $\mathbb{P}_{\mathbf{S}_{\mathbf{X}}}$ and LR images $\mathbb{P}_{\mathbf{Z}}$, respectively. The terms $D((\mathbf{X}, \mathbf{S}_{\mathbf{X}}))$, $D(G(\mathbf{Z}))$ and $D(\widehat{\mathbf{X}\mathbf{S}})$ are the responses of the discriminator with respect to the real data, the generated data and the interpolated data, respectively.

Finally, the loss function of the GAN is then:

$$\mathcal{L} = \mathcal{L}_{rec} + \lambda_{adv} \mathcal{L}_{adv} \quad (11)$$

where $\lambda_{adv} > 0$.


Finally the game between the generator and the discriminator is modeled as:

$$\min_G \max_D \mathcal{L} \quad (12)$$


2.4. Neural network architecture

Generator architecture. The generator network is a convolution-based network with residual blocks. It takes as input the interpolated LR image. It is composed of 18 convolutional layers: three for the encoding part, twelve for the residual part and three for the decoding part.

Let $C_j^i - S^k$ be a block consisting of the following layers: a convolution layer of j filters of size i^3 with stride of k , an instance normalization layer (InsNorm) [35] and a rectified linear unit (ReLU). R_k denotes a residual block as Conv-InsNorm-ReLU-Conv-InsNorm that contains 3^3 convolution layers with k filters. U_k denotes layers as Upsampling-Conv-InsNorm-ReLU layer with k filters of 3^3 and stride of 1. The generator architecture is then: $C_{16}^7 - S^1$, $C_{32}^3 - S^2$, $C_{64}^3 - S^2$, R_{64} , R_{64} , R_{64} , R_{64} , R_{64} , R_{64} , R_{64} , R_{64} , U_{32} , U_{16} , $C_2^7 - S^1$ (see Figure 1(a)).



(a) Generator architecture



(b) Discriminator architecture

Figure 1: Neural network architecture. (a) Generator architecture. (b) Discriminator architecture. See Section 2.4.

During the encoding, the number of kernels is multiplied by 2 at each convolution, from 16 to 64. The last convolutional layer produces two 3D images: the first will be turned into a class probability map (using a sigmoid activation); the second will be summed with the original interpolated image. In order to improve the training procedure performance, instance normalization layers are used on the result of each convolution (before application of activation function).

Note that summing the prediction of the generator with the interpolated image implies some constraints on the interpolated image size. Indeed, the decoder uses two upsampling layers which multiply the size of each channel by two. Then, the neural network output size is divisible by 4 in each dimension. Finally, to have the same dimension between input and output, we need an input size which is also divisible by 4.

Discriminator architecture. The discriminator network is fully convolutional. It takes as input an HR image and a segmentation map.

The discriminator contains five convolutional layers with an increasing number of filter kernels, increasing by a factor of 2 from 32 to 512 kernels. Let C_k be a block consisting of the following layers: a convolution layer of k filters of size 4^3 with stride of 2 and a Leaky ReLU with a negative slope of 0.01. The last layer C_1^2 is a 2^3 convolution filter with stride of 1. No activation layer is used after the last layer. The discriminator then consists of $C_{32}, C_{64}, C_{128}, C_{256}, C_{512}, C_1^2$ (see

Figure 1(b)).

For the generator as for the discriminator, the number of output channels for each convolutional layer is multiplied by 2 at each layer. The number of output channels of the first layer is a parameter of the algorithm, as further discussed in Section 4.

2.5. Training and testing

Before each application of the network, the LR images are normalized within $[0, 1]$ and interpolated using cubic splines. These two steps are common to both training and testing.

As we will see in the next section, before applying the neural network, the images are also split into patches of a given size.

For training purpose, the data can also be augmented with respect to (1) contrast modification (LR and HR images) and (2) Gaussian noise addition (LR images only). In particular, this can be of particular interest when the data to be processed are distinct from the data used for training.

Training. Figure 2(a) illustrates the pipeline considered for designing the training, at each epoch. Practically, data augmentation is carried out differently for each epoch and for each image. Indeed, for each image and for each epoch, the coefficient used for contrast modification is randomly drawn, such as the additive Gaussian noise added to the image.

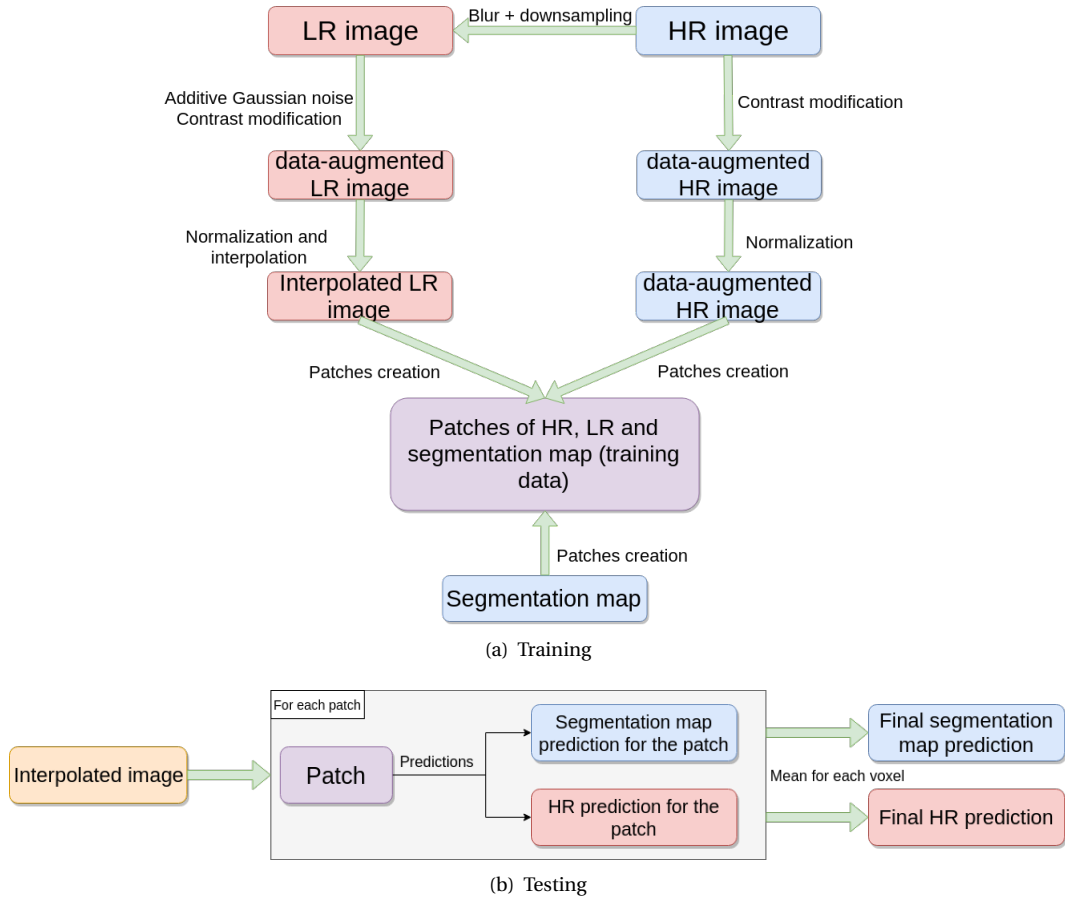


Figure 2: (a) Training data creation for one epoch, and (b) prediction in the testing pipeline. See Section 2.5.

From these training data, the discriminator and the generator are trained the following way. For each generator weights update, the discriminator weights are updated five times. The training data for the discriminator are randomly chosen at each iteration, whereas the generator training data successively pass over each batch.

The parameters used for the training are the following:

- Adam optimizer parameters similar as in the original Adam paper [19];
- batch size of 32 patches;
- $\lambda_{gp} = 100$ and $\lambda_{adv} = 0.001$;
- patch size of 64^3 with a 20 voxels shifting.

For each training carried out, the maximal number of epochs is set to 200 and the final weights are those which maximize the performance on a testing data set.

Testing. Figure 2(b) illustrates the testing pipeline.

3. Technical description

3.1. General statements

The proposed code is written in Python 3.6 (but it also works in Python 2.7). It was packaged on Pipy² and can be easily installed using:

```
1 pip install SegSRGAN
```

It depends on several packages, all of which are automatically installed when installing the SegSRGAN pip package. The dependencies are mainstream, widely used libraries such as Numpy³, Tensorflow⁴, Keras⁵, Pandas⁶, or SimpleITK⁷.

²<https://pypi.org/project/SegSRGAN>

³<https://www.numpy.org>

⁴<https://www.tensorflow.org>

⁵<https://keras.io>

⁶<https://pandas.pydata.org>

⁷<http://www.simpleitk.org>

Table 1: Example of a csv file used for training.

HR_image	Label_image	Base
/home/user/data/HR/1.nii.gz	/home/user/data/Label/1.nii.gz	Train
/home/user/data/HR/2.nii.gz	/home/user/data/Label/2.nii.gz	Test
/home/user/data/HR/3.nii.gz	/home/user/data/Label/3.nii.gz	Test
/home/user/data/HR/4.nii.gz	/home/user/data/Label/4.nii.gz	Train
/home/user/data/HR/5.nii.gz	/home/user/data/Label/5.nii.gz	Test

Since we use tensorflow-gpu dependency, the program can be computed on several GPUs. It is indeed recommended to run the code on GPUs instead of CPUs whenever possible, for the sake of computational efficiency.

Trained weights for the SegSRGAN networks are available in a Github repository⁸. When installing the pip package, these weights are locally downloaded while importing the SegSRGAN package in Python. (The weights folder is contained in the SegSRGAN package directory.) In particular, all the test results presented in Section 4 were obtained with these weights, for the sake of reproducibility.

All the Python files and the weights considered hereafter can be found in the Python package directory. The path of this folder is obtained by executing the following Python command:

```
1 import imp
2 imp.find_module('SegSRGAN')
```

and then entering the subdirectory “SegSRGAN” in the path returned by the result of this command.

3.2. Training

The code runs on multiple GPUs thanks to the Keras library (a parameter allows to use all available GPUs). In our models of generator and discriminator, we used the `multi_gpu_model` function⁹. It is optimized for the merging of the results when the GPUs are connected via NVLinks (wire-based communications protocol serial multi-lane near-range communication link developed by NVIDIA).

3.2.1. An example of basic command line

The package contains a file name “SegSRGAN_training.py” that allows one to train a SegSRGAN model. This file is located in the package folder which can be found as explained in Section 3.1. The basic command for performing the training from an image training database is for instance:

⁸<https://github.com/koopa31/SegSRGAN>

⁹<https://keras.io/utils>

```
1 python SegSRGAN_training.py --new_low_res 0.5,0.5,3
   --csv /home/user/data.csv --snapshot_folder /
   home/user/training_weights --dice_file /home/
   user/dice.csv --mse_file /home/user/
   mse_example_for_article.csv --
   folder_training_data /home/user/
   temporary_file_for_training
```

The only requirement to be fulfilled beforehand is the definition of an appropriate csv file. One can find an example of such file in Table 1. Here, the “HR_image” column must contain all paths to the HR images; the “Label_image” column must contain all paths to segmentation maps; and finally, the “Base” column must contain either the “Train” or “Test” label.

The learned weights are saved into the folder “/home/user/training_weights” that contains one dedicated file per epoch. At the end of each epoch, the dice and mse file (containing the Dice and the MSE scores, respectively) are also saved.

Warning – The argument “folder_training_data” corresponds to a folder, created during the script execution, in which all the training data are split into patches and organized by batch. This folder, which can be large (20 to 60 GB, depending on the chosen arguments) is deleted at the end of each epoch. In case of unexpected interruption of the script, it may be required to delete it manually.

3.2.2. Options

The parameters that control the training from scratch are the following:

- **new_low_res**: resolution of the LR image generated during the training. One value is given per dimension, for fixed resolution (e.g. “--new_low_res 0.5 0.5 3”). Two values are given per dimension if the resolutions have to be drawn between bounds (e.g. “--new_low_res 0.5 0.5 4 --new_low_res 1 1 2” means that for each image at each epoch, x and y resolutions are uniformly drawn between 0.5 and 1, whereas z resolution is uniformly drawn between 2 and 4).

- **csv** (string): file that contains the paths to the files used for training. These files are divided into two categories: train and test. Consequently, it must contain 3 columns, called: HR_image, Label_image and Base (which is equal to either Train or Test), respectively.
- **snapshot_folder** (string): path of the folder in which the weights will be regularly saved after a given number of epochs (this number is given by the **snapshot** (integer) argument).
- **dice_file** (string): CSV file where the Dice scores are stored at each epoch.
- **mse_file** (string): MSE file where the Dice scores are stored at each epoch.
- **folder_training_data** (string): folder where temporary files are written during the training (created at the beginning of each epoch and deleted at the end of it).
- **epoch** (integer): number of epochs up to which the training is performed.
- **batch_size** (integer): number of patches per mini-batch.
- **number_of_discriminator_iteration** (integer): how many times the discriminator is trained before training the generator.

It is also possible to continue a training from weights saved beforehand, by adding the following parameters:

- **init_epoch** (integer): number of the first epoch which will be considered during the continued training (e.g., 21 if the weights given were those obtained at the end of the 20th epoch). This is mainly useful for writing the weights in the same folder as the training which is continued.
Warning – The number of epochs of the remaining training is then **epoch** – **initepoch** + 1.
- **weights** (string): path to the saved weights from which the training will continue.

Important parameters to set the structure of the network (see Section 2.4) are:

- **kernel_gen** (integer): number of output channels of the first convolutional layer of the generator.

- **kernel_dis** (integer): number of output channels of the first convolutional layer of the discriminator.
- **is_conditional** (Boolean): enables to train a conditional network with a condition on the input resolution (discriminator *and* generator are conditionnal)
- **u_net** (Boolean): enables to train U-Net network. This kind of network is not detailed in this article. For more information, please check the Pipy documentation¹⁰.
- **is_residual** (Boolean): determines whether the structure of the network is residual or not.

Finally, the package also allows one to train from an augmented dataset, as described in section 2.5, by using the following parameters :

- **percent_val_max**: multiplicative value that gives the ratio of the maximal value of the image, for defining the standard deviation of the additive Gaussian noise. For instance, a value of 0.03 means that $\sigma = 0.03 \max(X)$ where $\max(X)$ is the maximal value of the image X .
- **contrast_max**: controls the modification of contrast of each image. For instance, a value of 0.4 means that at each epoch, each image will be set to a power uniformly drawn between 0.6 and 1.4.

Complementary information about all the parameters and their default values can be obtained by:

```
python SegSRGAN_train_avec_base_test.py --help
```

Warning – The patch size is set to 64 (argument **patch_size**). It is possible to change it; however this requires that the discriminator architecture be modified too, as it computes successive convolutions until the output is of size one (see Section 2.4).

3.3. Testing

3.3.1. Testing on one or more images (command line)

We now describe the way of performing segmentation on a set of many MR images (instead of segmenting one image many times). This command line works with NIfTI and DICOM images but the results are returned as NIfTI images (located in the package folder which can be found as explained in Section 3.1). An example of corresponding command for such group processing is:

¹⁰<https://pypi.org/project/SegSRGAN>


```

1 python job_model.py --path /home/data.csv --patch "
  64,128" --step "32 64,64 128" --
  result_folder_name "
  weights_without_augmentation" --weights_path "
  weights/Perso_without_data_augmentation"

```

csv path parameter. A csv file, as the one mentioned in the above example, is used to get the paths of all the images to be processed. Only the first column of each entry will be used, and the file it must contain only paths (i.e. no header).

Step and patch parameters. In this example, we run steps 32 and 64 for patch 64 and steps 64 and 128 for patch 128. The list of the paths of the images to be processed must be stored in a CSV file.

Warning – It is mandatory to respect exactly the same shape for the given step and patch.

Weights parameter. The implementation of the algorithm allows one to use two different kinds of weights:

- The weights we have already trained.
- New weights one can obtain through the training.

In order to use the weights we have already trained, the easiest solution is to provide for the `--weights_path` parameters some values as exemplified hereafter:

- **weights/Perso_without_data_augmentation:** corresponding to the weights without data augmentation.
- **weights/Perso_with_constrast_0.5_and_noise_0.03_val_max:** corresponding to the weights with data augmentation as described in Section 4.

Others weights not presented in this article are available (the help of SegSRGAN provides the list of all these available weights).

Organizing the output storage. Each image to be processed has to be stored in its own folder. When processing a given input image (which can be either a NIfTI image or a DICOM folder), a dedicated folder is created for each output. This folder will be located in the folder of the input image which has been processed and will be named with respect to the value of the parameter `--result_folder_name`




Figure 3: Data organization for the output images of a same input image.

(in our example the folder will be named “result_with_Weights_without_augmentation”). In Figure 3, one can see how the data are organized in the folder “result_with_Weights_without_augmentation”. Finally, each folder presented in Figure 3 contains two NIfTI file, namely the SR and the segmentation.

3.3.2. Testing code on one image in Python

The command line presented below allows one to segmented one or more images. Practically, it relies on a Python function that segments one image and write the result on the disk.

This function, called “segmentation”, is located in the “Function_for_application_test_python3.py” file and can be applied as follows:

```

1 from SegSRGAN.SegSRGAN.
  Function_for_application_test_python3 import
  segmentation
2
3 segmentation(input_file_path = "/home/user/data/1/1.
  nii.gz", step = 30, new_resolution =
  (0.5,0.5,0.5), patch=128, path_output_cortex = "/
  home/user/1/cortex.nii.gz", path_output_hr = "/
  home/user/1/SR.nii.gz", weights_path="/home/
  user/anaconda3/lib/python3.7/site-packages/
  SegSRGAN/SegSRGAN/weights/
  Perso_without_data_augmentation")

```

The main two differences between the function and a command line invoking it are the following:

- only one image can be segmented by the function, whereas many can be via a command line;
- the result folders need to be created before applying the Python function.

4. Experiments and results

4.1. Quality metrics

Segmentation. In order to assess the quality of segmentation results, we consider the Dice score [9], which is a standard measure for that purpose. In particular, the adequacy of the computed segmentation S with respect to the ground-truth G is then given as:

$$\text{Dice}(S, G) = \frac{2|S \cap G|}{|S| + |G|} \quad (13)$$

and lies in $[0, 1]$. The closer the Dice score to 1, the better the correlation between S and G .

In addition to the quantitative information carried by the Dice score, we also consider the number of connected components of the segmented results, noted NCC. By assuming that the cortex is a connected object, NCC provide structural information: the higher this value, the lower the topological quality of the object.

Super-resolution. Measuring the performance of SR algorithms is less straightforward. Indeed, for gauging the visual aspect similarity between two images, a distance between the intensity of the SR and HR voxels may not be sufficient. The performance of SR reconstruction is then measured by two different indexes, namely the PSNR and the SSIM [38], defined respectively as:

$$\text{PSNR}(X, Y) = 10 \log_{10} \frac{(\max_i X_i)^2}{\frac{1}{|X|} \sum_i |X_i - Y_i|^2} \quad (14)$$

where X_i and Y_i are the values of X and Y at point i , respectively, and:

$$\text{SSIM}(X, Y) = \frac{(2\mu_X\mu_Y + c_1) + (2\sigma_{XY} + c_2)}{(\mu_X^2 + \mu_Y^2 + c_1)(\sigma_X^2 + \sigma_Y^2 + c_2)} \quad (15)$$

where μ_X (resp. μ_Y) is the mean of X (resp. Y) values, σ_X (resp. σ_Y) is the standard deviation of X (resp. Y) values, σ_{XY} is the covariance between X and Y values, and c_1, c_2 are numerical stabilizers quadratically linked to the dynamics of the image. For both PSNR and SSIM indices, the higher the value, the better the similarity.

4.2. Data

4.2.1. Datasets

We work on two MRI datasets, namely dHCP¹¹ [16], and the French Epirmex¹² dataset.

dHCP. We consider T2-weighted MR images of the Developing Human Connectome Project, provided by the Evelina Neonatal Imaging Centre, London, UK. Forty neonatal data were acquired on a 3T Achieva scanner with repetition (TR) and echo times (TE) of 12000 ms and 156 ms, respectively. All dHCP images

¹¹<http://www.developingconnectome.org>

¹²Epirmex is a part of the French epidemiologic study Epipage 2 [1], <http://epipage2.inserm.fr>.

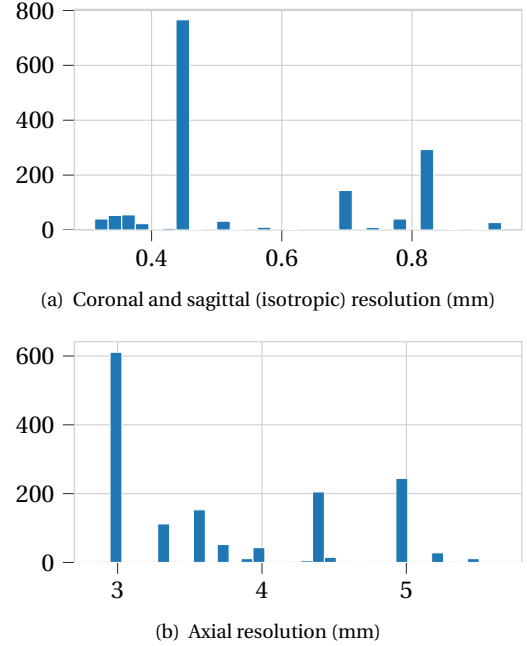


Figure 4: Histograms of the image resolution for the MR images from Epirmex: (a) coronal and sagittal; (b) axial.




Figure 5: Two examples of MR images (axial slices) from the Epirmex dataset. One can observe the differences in terms of contrast.

have a size of $290 \times 290 \times 203$ voxels, and a resolution of $0.5 \times 0.5 \times 0.5 \text{ mm}^3$.

32 images are used for the training and the 8 remaining ones are used to evaluate the model performance (test). Since dHCP MR images are HR data endowed with ground-truth, this dataset was mainly considered in our experiments, and in particular for quantitative evaluations.

Epirmex. The Epirmex dataset is composed of 1500 MR images acquired in 11 different French hospitals. These MR images were acquired with different scanners.

Most of Epirmex images present a number of rows




Figure 6: Dice and PSNR evolution during the training on the dHCP dataset.

(coronal) and columns (sagittal) comparable to those of dHCP images. The number of slices (axial) is also comparable, once the images are modeled at the same resolution.

Similarly to dHCP, all MR images in Epirmex are isotropic in their axial slices. One can find histograms for the resolutions in each dimension in Figure 4.

For half of these images, the resolution of the axial slices (between 0.4 and 0.6 mm) is comparable to that of dHCP images. The mode of this distribution is approximately 0.45 mm, and is much more frequent than the other values. The mode of the resolution along the axial dimension is near to 3 mm, and 40% of the dataset have exactly 3 mm of resolution. (This is the value which is used to generate LR image in dHCP database). The particularity of the Epirmex dataset is indeed its high anisotropy between coronal-sagittal and axial dimensions.

The TR and TE parameters can exhibit very different values from one image to another. This heterogeneity also constitutes a strong difference with dHCP. In practice, this implies that the images have different contrasts, as one can see e.g. in Figure 5.

Since Epirmex MR images are clinical, LR data non-endowed with ground-truth, this dataset was considered for more illustrative purpose, and to qualitatively emphasize the ability of the SR and segmentation methods to deal with true clinical applications.

4.2.2. Preprocessing: LR image generation (dHCP)

The dHCP dataset is composed of HR images equipped with binary segmentation (ground-truth), in particular for the cerebral cortex at the same high

resolution. As a consequence, in order to train a super-resolution model, we need to determine corresponding LR images for these couples HR images / segmentation maps.

From a given HR image X , the associated LR image X_{LR} is generated using the following model, as proposed in [12]:

$$X_{LR} = H_1 B X \quad (16)$$

where B is a blur matrix and H_1 is a downsampling decimation. In particular, we consider a Gaussian filter B with a standard deviation:

$$\sigma = \frac{\text{res}}{2\sqrt{2\log 2}} \quad (17)$$

where res is the resolution of the HR image.

The generated LR images considered in our experimental study are of resolution $0.5 \times 0.5 \times 3 \text{ mm}^3$, which is compliant with true clinical data, which are usually strongly anisotropic, as emphasized by the above Epirmex data description.

4.3. Convergence – Training (dHCP)

First, we observe the evolution the Dice and PSNR scores, along the training. The values presented here are calculated as follows. The test images are split into patches of size 64^3 voxels, using a 20 voxel shifting (the same used for the training image set). The PSNR and the Dice scores are then calculated in each patch. The final PSNR and the Dice scores are obtained by averaging the values computed patch-wise.

The results are depicted in Figure 6, that provides the evolution of the Dice and the PSNR scores at the

Table 2: Dice, PSNR and SSIM scores (mean values over 8 images) for SR reconstruction and segmentation, compared with cubic spline interpolation on dHCP (see text).

	SegSRGAN	LR interpolation	Segmentation map interpolation
Dice	0.856 ± 0.014	–	0.735 ± 0.012
PSNR	26.96	24.22	–
SSIM	0.73	0.63	–

end of each epoch. The initial Dice has a very low value, close to 0.5, and then increases up to 0.8 rather smoothly. It seems to converge from the 100th epoch. The PSNR also converges, but in a more noisy way. However, the size of the peaks progressively decreases whereas the score tends to stabilize, following the same behaviour as the Dice score.

4.4. Results – Testing (dHCP)

The results presented here were computed from the 8 images from dHCP used as testing dataset. These results were obtained with patches of size 128^3 voxels, with a 30-voxel shifting.

Table 2 summarizes the various quality scores for SR and segmentation. On the one hand, we can observe that the two quality scores for the SR image reconstruction exhibit much better results with SegSRGAN than with cubic spline interpolation (which constitutes a standard baseline and is the input of the network). On the other hand, applying downsampling and interpolation to the segmentation map allows one to understand the difficulty to retrieve the information actually lost during the downsampling of a high resolution image, in particular due to blurring effects.

4.4.1. Impact of patch overlapping

Impact on the computation time. The time cost of the algorithm is directly influenced by the choice of the step (that controls patch overlapping). Indeed, the value of step directly implies the number of patches which have to pass through the neural network. Actually, this number of patches depends directly from the value

$$\prod_{\star \in \{x,y,z\}} \frac{n_{\star} - patch_{\star}}{step} \quad (18)$$

where (n_x, n_y, n_z) is the size of the smallest image containing the interpolated image and which allows to create an integer number of full patches of size $(patch_x, patch_y, patch_z)$ with $step$ between the successive patches.

From Equation (18), it is plain that the number of patches is calculated from a quantity that evolves like




Figure 7: Computation time (in seconds) against the step value.


$step^{-3}$. Furthermore, even though (n_x, n_y, n_z) also depends on the step, one can approximate the computation time by an affine function of $step^{-3}$, as confirmed by Figure 7.

The computation times presented in this figure were obtained in the following configuration:


- patch size of 128^3 ;
- on dHCP, which implies that all the HR images which led to the interpolated images are of size $(290, 290, 198)$;
- computation on Tesla P100 GPU and 2 Intel® Xeon™ Gold “Skylake” 6132 CPU.

Here, the computation time varies from 600 seconds (10 min) for a step value of 10, to 4 seconds for a step value of 100, with a $step^{-3}$ decrease between these two values. It is worth mentioning that the device used here processes the patches quickly (approximately 5 patches per second) which tends to reduce the impact of the variation of (n_x, n_y, n_z) ; this could not be the same on other kinds of devices.


Impact on algorithm performance. We now investigate the impact of the step value on the algorithm performance.



(a) Dice



(b) PSNR



(c) SSIM

Figure 8: (a) Dice scores of segmentation results on dHCP, depending on the step between successive patches of size 128^3 voxels. PSNR (b) and SSIM (c) scores of SR reconstruction results on dHCP, depending on the step between successive patches of size 128^3 voxels.

For segmentation purpose, the Dice scores obtained on dHCP are provided in Figure 8(a). Here, the




Figure 9: Number of connected components (NCC) of the segmented cortex, depending on the step between successive patches of size 128^3 voxels. Calculated on dHCP database

patches are still of size 128^3 voxels. We observe that in any cases, overlapping between patches (i.e. with steps lower than 128) provides better results than without (i.e. with a step of 128). In overlapping cases, the Dice score grows roughly linearly with respect to the size of the step, for reaching values around 0.86 with maximal overlaps, i.e. for steps close to 1.

This behaviour for segmentation is confirmed by the behaviour for SR reconstruction, as illustrated by Figure 8(b–c), which provides the PSNR and SSIM scores with respect to the step.

The cortex is a continuous ribbon-like structure. In particular, the connectedness of the segmentation results provided by SegSRGAN is a relevant property, in complement to the three considered quantitative scores. In Figure 9, we show how the value of step between patches influences the number of connected components of the segmented cortex. In theory, only one connected component should be obtained. In other words, obtaining n connected components means that $n - 1$ parts of the segmented cortex are erroneously disjoint from the topologically correct structure. Once again, one can observe a linear correlation between the error and the value of the step, with close to 1 values for the higher overlapping / lower steps.

4.4.2. Impact of noise

MR images are generally affected by noise. We now investigate in which extent the performance of the segmentation and SR reconstruction is impacted by adding Gaussian noise to the image. In practice, a Gaussian noise with $\sigma = 2$ is added to each voxel of the 8 test MRI images.




Figure 10: Difference between the Dice scores with non-noisy and noisy images from dHCP image. Positive values mean that the Dice scores are higher without than with Gaussian noise.

The difference between the Dice scores of the segmentation results with and without noise, respectively, is depicted in Figure 10. As expected, the relative Dice scores are slightly better for non-noisy images than for noisy ones, with a gap within $[0.03, 0.07]$. (Note that from a topological point of view, we also observed that the number of connected components is approximately two times greater in the noisy images compared to non-noisy ones.) We observe that the difference between noisy and non-noisy data increases (i.e. the impact of adding noise increases) when the step grows. One can then conclude that, the larger the step (i.e. the lower the patch overlapping), the lower the robustness of the method versus noise. In terms of SR reconstruction, we also compared the evolution of SSIM and PSNR with respect to the step value. The results are given in Tables 3–4. They emphasize a strong degradation of SSIM, and a much lower degradation of PSNR. As already observed for segmentation results, when increasing patch overlapping, both scores are significantly improved. For SSIM, the gain between steps sizes of 128 and 30 is of approximately 10% in both noisy and non-noisy cases. For PSNR, it is slightly lower (approximately 8%) with than without noise (approximately 10%).

4.4.3. Impact of data augmentation

Motivation. The experiments described until now have been carried out on dHCP, i.e. an image dataset designed for research purpose, with good properties in terms of noise and signal homogeneity. In real, clinical cases, the images forming a dataset are generally of lower quality both in terms of signal-to-noise ratio

Table 3: Mean SSIM scores for SR reconstruction on noisy and non-noisy images (dHCP).

Step	Non-noisy	Noisy
30	0.73	0.49
80	0.69	0.46
128	0.65	0.44

Table 4: Mean PSNR scores for SR reconstruction on noisy and non-noisy images (dHCP).

Step	Non-noisy	Noisy
30	26.96	25.40
80	25.85	24.45
128	24.24	23.24

and signal homogeneity. For instance, all the images in dHCP were acquired with the same TR and TE (implying inter-image signal homogeneity); by contrast the TR and TE in Epirmex can strongly vary. Moreover, by contrast with research-oriented datasets, clinical data are generally not equipped with ground-truths; indeed, defining such ground-truths is a complex and time-consuming task that requires heavy work by medical experts. (The same remark also holds for the non-availability of HR images associated to the native LR ones.) Although our final purpose is to be able to process real, clinical data, it is generally not possible to perform the training on comparable images. Such training then need to rely on dHCP-like data. As a consequence, it is indeed relevant to consider data augmentation strategies for increasing the ability of the trained networks to take into account real noise and inhomogeneity properties of usual MR images.

Data augmentation. The considered data augmentation on dHCP images consists of applying contrast modification uniformly drawn in $[0.5, 1.5]$, and adding Gaussian noise with a standard deviation set to 3% of the highest value.

Two questions then arise: (1) Does this training improve the result on noisy / signal-biased data? (2) Does this training make the performance decrease on the original images? In order to answer these questions, three supplementary families of images, created from the original database, are considered during the testing, namely images:

- with additive Gaussian noise ($\sigma = 2$);

Table 5: Mean values of Dice scores for segmentation of dHCP images (noisy and non-noisy) based on standard or data-augmented training.

Step	Original			Noisy		
	30	80	128	30	80	128
Without data augmentation	0.855	0.831	0.805	0.819	0.785	0.740
With data augmentation	0.849	0.827	0.806	0.836	0.812	0.790

- with values increased by a square function;
- with values decreased by a square root function.

Hereafter, these three types of images are denoted as “augmented test dataset”.

Segmentation. The results obtained for segmentation purpose are depicted in Figure 11. First, one can observe that in case of large steps, data augmentation improves more significantly the segmentation results (the best example is when the algorithm is applied without overlapping). However, it is important to keep in mind that the Dice score decreases with respect to the step. In other words, data augmentation allows, in a certain extent, to compensate the defects caused by the choice of a large step. In the case of overlapping, the value of step seems to have a negligible quantitative impact on the improvements. In most cases, data augmentation allows to slightly increase the robustness of segmentation results, with Dice scores increase between [0.00,0.03], when considering augmented test datasets. This is not the case for the native dHCP dataset, where the Dice scores are slightly lowered by data augmentation training. Indeed, data augmentation allows the network to learn from a wider range of data; making it more robust with respect to data variability. As a counterpart, a such versatile network becomes less robust for handling a homogeneous population of data such as dHCP.

In complement to the relative evaluation provided by Figure 11, we also propose, in Table 5, an absolute comparison of the Dice scores between segmentation performed with and without data augmentation, both on the dHCP images and their noisy versions. One can observe that the segmentation results for original dHCP images are slightly degraded (Dice decrease lower than 0.01) with data augmentation, with reasons similar to those evoked in the analysis of Figure 11. In the meantime, the segmentation results for noisy images are improved (Dice increase from 0.015 to 0.050). In particular, in the case of data augmentation, the difference between Dice scores on standard and noisy images becomes low (approximately 0.015), compared to the case without data augmentation (ap-

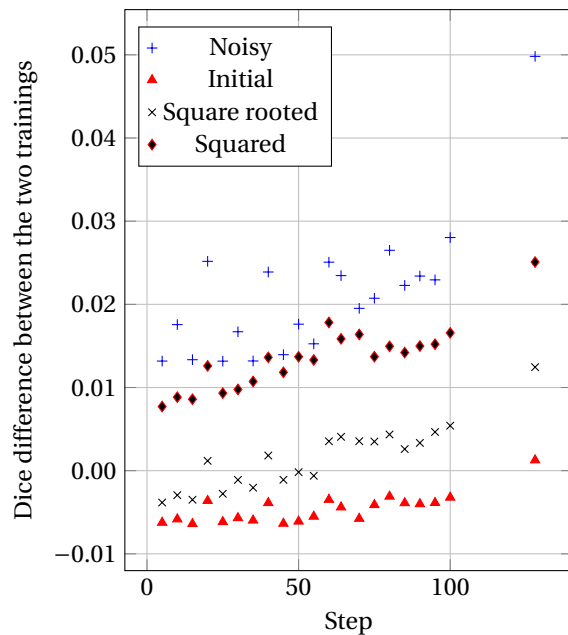


Figure 11: Difference between the mean Dice scores of segmentation results when training with and without data augmentation, depending on the step value. Positive (resp. negative) values mean that the Dice scores are better with (resp. without) data augmentation. Calculated on dHCP database

proximately 0.030 to 0.060). In other words, data augmentation makes the segmentation more robust to the tested images, at the price of a loss of specialization with respect to the trained images. However, such trade-off is relevant with respect to real clinical applications.

In order to complete this quantitative segmentation analysis, we also propose, in Figure 12 and Table 6, a topological analysis of the results, by providing the mean number of connected components of the segmented cortex. These topological results strengthen the conclusions obtained for Dice analysis on augmented data. In particular, on original data, an improvement is also observed from this topological point of view. Indeed, the number of connected components is reduced by data augmentation for all types of data studied. In particular, the improvement is important for noisy images. In addition, (see Table 6) as

Table 6: Mean number of connected components for segmentation of dHCP images (noisy and non-noisy) based on standard or data-augmented training.

Step	Original			Noisy		
	30	80	128	30	80	128
Without data augmentation	8.6	45.3	104.3	26.6	98.5	193.1
With data augmentation	5.7	33.6	61.7	5.8	33.5	61.8

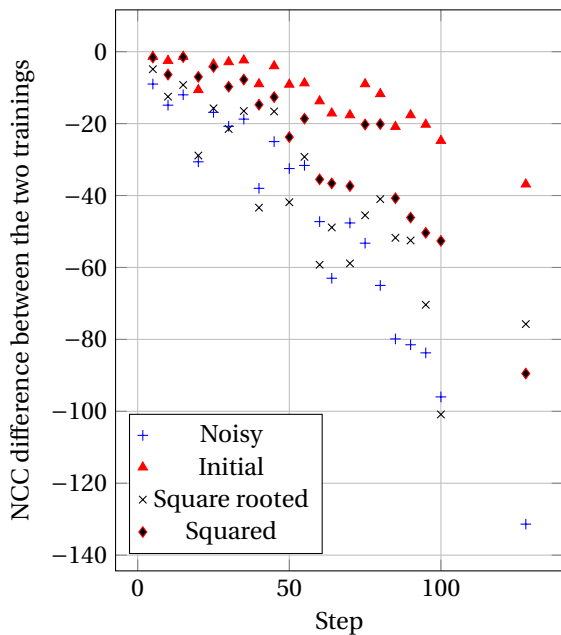


Figure 12: Difference between the number of connected components of segmentation results when training with and without data augmentation, depending on the step value. Negative values mean that the number of connected components is lower with data augmentation. Calculated on dHCP database

for Dice analysis, it appears that segmentation results present very similar topological properties independently of the noise level of the input images.

SR reconstruction. Concerning SR reconstruction, when observing Figures 13 and 14, it appears that for both SSIM and PSNR, the training with data augmentation (noise and contrast variation) slightly decreases the quality of reconstruction on initial images. This behaviour, already observed in the case of segmentation can be explained by the same factors. However, we also observe that the reconstruction is slightly improved for noisy data, whereas it is slightly degraded for contrast-modified data. These results on noisy images corroborate those of segmentation, whereas they are antagonistic in the case of contrast variation. This last point would require to more accurately investi-




Figure 13: SSIM difference of SR reconstruction between the training with and without data augmentation. Positive (resp. negative) values mean that the SSIM is higher with (resp. without) data augmentation. Calculated on dHCP database

gate the specific effects of data augmentation with respect to either noise and contrast variation.

4.5. Results – Testing (Epirmex)

In this last part of the experimental section, we apply our SR reconstruction and segmentation method to the MR images from the Epirmex dataset.

Epirmex contains more than 1500 images, with very different properties, e.g. the TE and TR values, as discussed above (see Figure 5). These MR images are also of low resolution and a strong anisotropy (see Figure 4).

In addition, Epirmex is endowed neither with ground-truth nor with HR images associated to the LR ones, with two main side effects. The first is that the training cannot be carried out on Epirmex images, and it is then mandatory to learn from another dataset, with a data augmentation paradigm; this is what we did

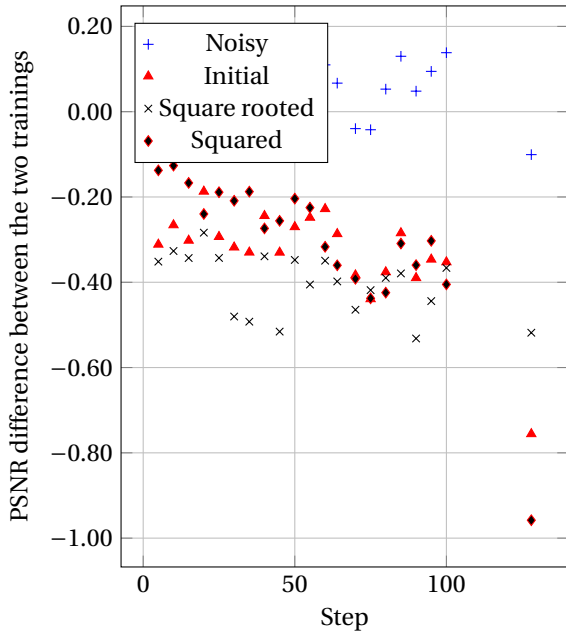


Figure 14: PSNR difference of SR reconstruction between the training with and without data augmentation. Positive (resp. negative) values mean that the PSNR is higher with (resp. without) data augmentation. Calculated on dHCP database

with dHCP. Second, there is no objective way of assessing the quality of the results, both in terms of segmentation and SR reconstruction. As a consequence, the results provided hereafter have mainly a qualitative value.

In particular we propose, in Figure 15, some SR reconstruction and cortex segmentation results for an MR image representative of the data within Epirmex. Those results were obtained from the parameters learned during the training on dHCP, with data augmentation. A visual analysis of these results leads to satisfactory conclusions. In particular, the segmented cortex is geometrically consistent, with a good visual correlation to the input MR image. Regarding the SR reconstruction, the contrast between the cortex and the surrounding tissues seems higher in the SR image than in the LR one, even in the axial slices, where the resolution is however not modified between LR and SR. Globally, these qualitative experiments on Epirmex tend to corroborate the quantitative ones carried out on dHCP, and to suggest that the proposed method is indeed relevant both for reconstruction and segmentation in the context of real, clinical data analysis.

Due to the lack of ground-truth both for segmentation and SR reconstruction, we can rely neither on Dice

Table 7: Mean number of connected components (NCC) computed from 7 Epirmex images.

Step	30	50	80	128
NCC with data augmentation	30	35	96	387
NCC without data augmentation	36	53	102	720


nor on SSIM and PSNR for assessing the quality of the results on Epirmex. From a topological point of view, it is however possible to assess the structural quality of the segmented cortex, that is assumed to be fully connected.

In particular, the results stated in Table 7 confirm those previously obtained on dHCP (see Table 6). Indeed, one can observe that, in case of patch overlapping, the number of connected components is better when training with data augmentation. However, the highest difference between this two trainings occurs when the algorithm is used without patch overlapping (i.e. with a step of 128).


Similarly to the experimental results obtained on dHCP, the evolution of the number of connected components with respect to the step value argues in favour of reducing as much as possible the step for improving the topological coherence of the segmented cortex. In particular, on Epirmex, the sensitivity of NCC to the number of steps is higher than with dHCP. This can be easily explained by the lower quality of the data, that tends to induce erroneous disconnections, partially avoided by increasing patch overlapping.

5. Conclusion


In this article, we have proposed a new methodological and software solution for performing SR reconstruction and segmentation from complex 3D MR images. Our framework, based on generative adversarial networks has been described in details, both from theoretical and technical points of view. In particular, a free, documented software version is available, for dissemination to the scientific and clinical communities and for the sake of reproducibility of the results. Although our purpose was not to carry out a comparative work with other methods (we do not claim the superiority of our method, but its usefulness in certain contexts), we have proposed a consequent experimental analysis in the case of cortex investigation in the neonate. This experimental analysis, carried out on both research and clinical datasets, and from qualitative and quantitative points of view, tend to prove



(a) LR (initial) MR image



(b) SR (reconstructed) MR image



(c) Segmented image (cortex)

Figure 15: (a) LR image from the Epirmex dataset. (b) SR reconstructed image obtained from (a). (c) Segmentation of the cortex obtained from (a).

the relevance of our approach, with satisfactory results both in terms of SR reconstruction and segmentation. Of course, the proposed approach is not specific to the cortex, and it could also be used for other kinds of cerebral structures in MRI. Our first perspec-

tive indeed will consist of processing large image cohorts for automatically extracting regions of interest, in a clinical context.

From a methodological point of view, we will more deeply investigate the trade-off between computa-

tional cost and result quality, in particular with respect to data variability within MR image cohorts. In this context, we will also experimentally assess the various consequences on data augmentation with respect to usual features likely to degrade MR image quality, namely noise, signal bias, contrast heterogeneity, or movement artifacts.

Acknowledgements

The research leading to these results has been supported by the ANR MAIA project, grant ANR-15-CE23-0009 of the French National Research Agency (<http://recherche.imt-atlantique.fr/maia>); INSERM and Institut Mines Télécom Atlantique (Chaire “Imagerie médicale en thérapie interventionnelle”); the *Fondation pour la Recherche Médicale* (grant DIC20161236453); and the American Memorial Hospital Foundation. We also gratefully acknowledge the support of NVIDIA Corporation with the donation of the Titan Xp GPU used for this research.

References

- [1] P.-Y. Ancel, F. Goffinet, and EPIPAGE 2 Writing Group. EPIPAGE 2: A preterm birth cohort in France in 2011. *BMC Pediatrics*, 14:97, 2014.
- [2] G. Ball, J. P. Boardman, P. Aljabar, A. Pandit, T. Arichi, N. Merchant, D. Rueckert, A. D. Edwards, and S. J. Counsell. The influence of preterm birth on the developing thalamocortical connectome. *Cortex*, 49(6):1711–1721, 2013.
- [3] G. Ball, L. Srinivasan, P. Aljabar, S. J. Counsell, G. Durighel, J. V. Hajnal, M. A. Rutherford, and A. D. Edwards. Development of cortical microstructure in the preterm human brain. *Proceedings of the National Academy of Sciences of the United States of America*, 110(23):9541–9546, 2013.
- [4] M. J. Cardoso, A. Melbourne, G. S. Kendall, M. Modat, N. J. Robertson, N. Marlow, and S. Ourselin. AdaPT: An adaptive preterm segmentation algorithm for neonatal brain MRI. *NeuroImage*, 65:97–108, 2013.
- [5] P. Charbonnier, L. Blanc-Féraud, G. Aubert, and M. Barlaud. Deterministic edge-preserving regularization in computed imaging. *IEEE Transactions on Image Processing*, 6(2):298–311, 1997.
- [6] Y. Chen, F. Shi, A. G. Christodoulou, Y. Xie, Z. Zhou, and D. Li. Efficient and accurate MRI super-resolution using a generative adversarial network and 3D multi-level densely connected network. In *MICCAI, Proceedings*, pages 91–99, 2018.
- [7] Y. Chen, Y. Xie, Z. Zhou, F. Shi, A. G. Christodoulou, and D. Li. Brain MRI super resolution using 3D deep densely connected neural networks. In *ISBI, Proceedings*, pages 739–742, 2018.
- [8] P. Coupé, J. V. Manjón, V. Fonov, J. Pruessner, M. Robles, and D. L. Collins. Patch-based segmentation using expert priors: Application to hippocampus and ventricle segmentation. *NeuroImage*, 54:940–954, 2011.
- [9] L.R. Dice. Measures of the amount of ecologic association between species. *Ecology*, 26:297–302, 1945.
- [10] J. Dubois, M. Benders, A. Cachia, F. Lazeyras, R. Ha-Vinh Leuchter, S. V. Sizonenko, C. Borradori-Tolsa, J. F. Mangin, and P. S. Hüppi. Mapping the early cortical folding process in the preterm newborn brain. *Cerebral Cortex*, 18(6):1444–1454, 2008.
- [11] I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. C. Courville, and Y. Bengio. Generative adversarial nets. In *NIPS, Proceedings*, pages 2672–2680, 2014.
- [12] H. Greenspan. Super-resolution in medical imaging. *The Computer Journal*, 52(1):43–63, 2008.
- [13] L. Gui, R. Lisowski, T. Faundez, P. Hüppi, F. Lazeyras, and M. Kocher. Morphology-based segmentation of newborn brain MR images. In *MICCAI NeoBrainS12, Proceedings*, pages 1–8, 2012.
- [14] I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, and A. C. Courville. Improved training of Wasserstein GANs. In *NIPS, Proceedings*, pages 5769–5779, 2017.
- [15] M. Hack and A. A. Fanaroff. Outcomes of children of extremely low birthweight and gestational age in the 1990s. *Seminars in Neonatology*, 5:89–106, 2000.
- [16] E. Hughes, L. Cordero-Grande, M. Murgasova, J. Hutter, A. Price, A. Santos Gomes, J. Allsop, J. Steinweg, N. Tumor, J. Wurie, J. Bueno-Conde, J.-D. Tournier, M. Abaei, S. Counsell, M. Rutherford, M. Pietsch, D. Edwards, J. Hajnal, S. Fitzgibbon, E. Duff, M. Bastiani, J. Andersson, S. Jbabdi, S. Sotiropoulos, M. Jenkinson, S. Smith, S. Harrison, L. Griffanti, R. Wright, J. Bozek, C. Beckmann, A. Makropoulos, E. Robinson, A. Schuh, J. Passerat Palmbach, G. Lenz, F. Mortari, T. Tenev, and D. Rueckert. The Developing Human Connectome: Announcing the first release of open access neonatal brain imaging. In *OHBM, Proceedings*, 2017.
- [17] J. Johnson, A. Alahi, and L. Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *ECCV, Proceedings*, pages 694–711, 2016.
- [18] H. Kim, C. Lepage, A. C. Evans, J. Barkovich, and D. Xu. NEO-CIVET: Extraction of cortical surface and analysis of neonatal gyrification using a modified CIVET pipeline. In *MICCAI, Proceedings*, pages 571–579, 2015.
- [19] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980, 2014.
- [20] C. Ledig, L. Theis, F. Huszár, J. Caballero, A. Cunningham, A. Acosta, A. Aitken, A. Tejani, J. Totz, Z. Wang, and W. Shi. Photo-realistic single image super-resolution using a generative adversarial network. In *CVPR, Proceedings*, pages 105–114, 2017.
- [21] J. Lefèvre, D. Germainaud, J. Dubois, F. Rousseau, I. de Macedo Santos, H. Angleys, J.-F. Mangin, P. S. Hüppi, N. Girard, and F. De Guio. Are developmental trajectories of cortical folding comparable between cross-sectional datasets of fetuses and preterm newborns? *Cerebral Cortex*, 26(7):3023–3035, 2016.
- [22] D. Mahapatra. Skull stripping of neonatal brain MRI: Using prior shape information with graph cuts. *Journal of Digital Imaging*, 25:802–814, 2012.
- [23] A. Makropoulos, S. J. Counsell, and D. Rueckert. A review on automatic fetal and neonatal brain MRI segmentation. *NeuroImage*, 170:231–248, 2017.
- [24] A. Makropoulos, C. Ledig, P. Aljabar, A. Serag, J. V. Hajnal, A. D. Edwards, S. J. Counsell, and D. Rueckert. Automatic tissue and structural segmentation of neonatal brain MRI using Expectation-Maximization. In *MICCAI NeoBrainS12, Proceedings*, pages 9–15, 2012.
- [25] N. Marlow, D. Wolke, M. A. Bracewell, M. Samara, and EPI-Cure Study Group. Neurologic and developmental disability at six years of age after extremely preterm birth. *The New Eng-*

- land *Journal of Medicine*, 352:9–19, 2005.
- [26] A. Melbourne, M. J. Cardoso, G. S. Kendall, N. J. Robertson, N. Marlow, and S. Ourselin. NeoBrainS12 Challenge: Adaptive neonatal MRI brain segmentation with myelinated white matter class and automated extraction of ventricles I-IV. In *MICCAI NeoBrainS12, Proceedings*, pages 16–21, 2012.
 - [27] P. Moeskops, M. Veta, M. W. Lafarge, K. A. J. Eppenhof, and J. P. W. Pluim. Adversarial training and dilated convolutions for brain MRI segmentation. In *DLMA and ML-CDS, Proceedings*, pages 56–64, 2017.
 - [28] E. Orasanu, A. Melbourne, M. J. Cardoso, H. Lomabert, G. S. Kendall, N. J. Robertson, N. Marlow, and S. Ourselin. Cortical folding of the preterm brain: A longitudinal analysis of extremely preterm born neonates using spectral matching. *Brain and Behavior*, 6(8):e00488, 2016.
 - [29] M. Péporté, D. E. I. Ghita, E. Twomey, and P. F. Whelan. A hybrid approach to brain extraction from premature infant MRI. In *SCLA, Proceedings*, pages 719–730, 2011.
 - [30] C.-H. Pham, A. Ducournau, R. Fablet, and F. Rousseau. Brain MRI super-resolution using deep 3D convolutional networks. In *ISBI, Proceedings*, pages 197–200, 2017.
 - [31] C.-H. Pham, C. Tor-Diez, H. Meunier, N. Bednarek, R. Fablet, N. Passat, and F. Rousseau. Simultaneous super-resolution and segmentation using a generative adversarial network: Application to neonatal brain MRI. In *ISBI, Proceedings*, pages 991–994, 2019.
 - [32] M. Prastawa, J. H. Gilmore, W. Lin, and G. Gerig. Automatic segmentation of MR images of the developing newborn brain. *Medical Image Analysis*, 9:457–466, 2005.
 - [33] A. Serag, M. Blesa, E. J. Moore, R. Pataky, S. A. Sparrow, A. G. Wilkinson, G. Macnaught, S. I. Semple, and J. P. Boardman. Accurate learning with few atlases (ALFA): An algorithm for MRI neonatal brain extraction and comparison with 11 publicly available methods. *Scientific Reports*, 6:23470, 2016.
 - [34] C. Tor-Diez, N. Passat, I. Bloch, S. Faisan, N. Bednarek, and F. Rousseau. An iterative multi-atlas patch-based approach for cortex segmentation from neonatal MRI. *Computerized Medical Imaging and Graphics*, 70:73–82, 2018.
 - [35] D. Ulyanov, A. Vedaldi, and V. S. Lempitsky. Improved texture networks: Maximizing quality and diversity in feed-forward stylization and texture synthesis. In *CVPR, Proceedings*, pages 4105–4113, 2017.
 - [36] L. Wang, F. Shi, W. Lin, J. H. Gilmore, and D. Shen. Automatic segmentation of neonatal images using convex optimization and coupled level sets. *NeuroImage*, 58:805–817, 2011.
 - [37] L. Wang, F. Shi, P.-T. Yap, J. H. Gilmore, W. Lin, and D. Shen. 4D multi-modality tissue segmentation of serial infant images. *PLoS One*, 7:e44596, 2012.
 - [38] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli. Image quality assessment: From error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4):600–612, 2004.
 - [39] H. Xue, L. Srinivasan, S. Jiang, M. Rutherford, A. D. Edwards, D. Rueckert, and J. V. Hajnal. Automatic segmentation and reconstruction of the cortex from neonatal MRI. *NeuroImage*, 38:461–477, 2007.
 - [40] Q. Yu, A. Ouyang, L. Chalak, T. Jeon, J. Chia, V. Mishra, M. Sivarajan, G. Jackson, N. Rollins, S. Liu, and H. Huang. Structural development of Human fetal and preterm brain cortical plate based on population-averaged templates. *Cerebral Cortex*, 26(11):4381–4391, 2016.