

Received February 16, 2020, accepted March 9, 2020, date of current version March 19, 2020.

Digital Object Identifier 10.1109/ACCESS.2020.2980266

# Convolutional Neural Networks With Intermediate Loss for 3D Super-Resolution of CT and MRI Scans

MARIANA-IULIANA GEORGESCU<sup>1</sup>, RADU TUDOR IONESCU<sup>1</sup>, (Member, IEEE),  
AND NICOLAE VERGA<sup>2,3</sup>

<sup>1</sup>Department of Computer Science, University of Bucharest, 010014 Bucharest, Romania

<sup>2</sup>Department of Radiology, Oncology, Hematology, Carol Davila University of Medicine and Pharmacy, 020021 Bucharest, Romania

<sup>3</sup>Colțea Hospital, 030171 Bucharest, Romania

Corresponding author: Radu Tudor Ionescu (raducu.ionescu@gmail.com)

The research leading to these results has received funding from the EEA Grants 2014-2021, under Project contract no. EEA-RO-NO-2018-0496.

**ABSTRACT** Computed Tomography (CT) scanners that are commonly-used in hospitals and medical centers nowadays produce low-resolution images, e.g. one voxel in the image corresponds to at most one-cubic millimeter of tissue. In order to accurately segment tumors and make treatment plans, radiologists and oncologists need CT scans of higher resolution. The same problem appears in Magnetic Resonance Imaging (MRI). In this paper, we propose an approach for the single-image super-resolution of 3D CT or MRI scans. Our method is based on deep convolutional neural networks (CNNs) composed of 10 convolutional layers and an intermediate upscaling layer that is placed after the first 6 convolutional layers. Our first CNN, which increases the resolution on two axes (width and height), is followed by a second CNN, which increases the resolution on the third axis (depth). Different from other methods, we compute the loss with respect to the ground-truth high-resolution image right after the upscaling layer, in addition to computing the loss after the last convolutional layer. The intermediate loss forces our network to produce a better output, closer to the ground-truth. A widely-used approach to obtain sharp results is to add Gaussian blur using a fixed standard deviation. In order to avoid overfitting to a fixed standard deviation, we apply Gaussian smoothing with various standard deviations, unlike other approaches. We evaluate the proposed method in the context of 2D and 3D super-resolution of CT and MRI scans from two databases, comparing it to related works from the literature and baselines based on various interpolation schemes, using  $2\times$  and  $4\times$  scaling factors. The empirical study shows that our approach attains superior results to all other methods. Moreover, our subjective image quality assessment by human observers reveals that both doctors and regular annotators chose our method in favor of Lanczos interpolation in 97.55% cases for an upscaling factor of  $2\times$  and in 96.69% cases for an upscaling factor of  $4\times$ . In order to allow others to reproduce our state-of-the-art results, we provide our code as open source at <https://github.com/lilygeorgescu/3d-super-res-cnn>.

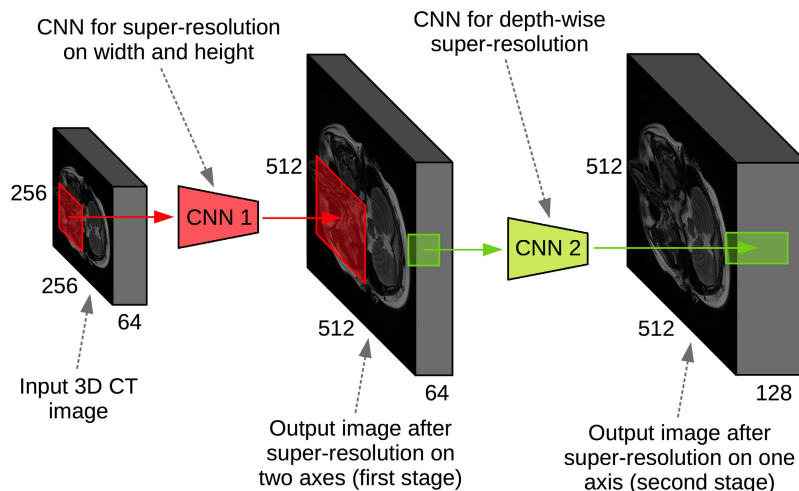
**INDEX TERMS** Convolutional neural networks, single-image super-resolution, CT images, MRI images, medical image super-resolution.

## I. INTRODUCTION

Medical centers and hospitals around the globe are typically equipped with single-energy Computer Tomography (CT) or Magnetic Resonance Imaging (MRI) scanners that produce cross-sectional images (slices) of various body parts. The resulting images are of low-resolution, since one pixel usually corresponds to at most one-millimeter piece of tissue. The thickness of one slice is one millimeter at best,

The associate editor coordinating the review of this manuscript and approving it for publication was Corrado Mencar<sup>1</sup>.

so the 3D CT images are composed of volumetric pixels (voxels) that usually correspond to one cubic millimeter ( $1 \times 1 \times 1 \text{ mm}^3$ ) of tissue. One of the main benefits of this non-invasive scanning technique is that it allows doctors to see if there are malignant tumors inside the body. Nevertheless, doctors, and even machine learning systems [1], are not able to accurately contour (segment) the tumor regions because of the low-resolution of CT or MRI scans. According to a team of radiologists from Colțea Hospital in Bucharest, that provided a set of anonymized CT scans for our experiments, the desired resolution is to have one voxel correspond to one



**FIGURE 1.** Our method for 3D image super-resolution based on two subsequent fully convolutional neural networks. In the first stage, the input volume is resized in two dimensions (width and height). In the second stage, the processed volume is further resized in the third dimension (depth). Using a scale factor of  $2\times$ , an input volume of  $256 \times 256 \times 64$  components is upscaled to  $512 \times 512 \times 128$  components (on all axes). Best viewed in color.

cubic micrometer (a thousandth part of a cubic millimeter) of tissue. In other words, the goal is to increase the resolution of 3D CT and MRI scans by a factor of  $10\times$  in each direction.

The main motivation behind our work is to allow radiologists and oncologists to accurately segment tumors and make better treatment plans. In order to achieve the desired goal, we propose a machine learning method that takes as input a 3D image and increases the resolution of the input image by a factor of  $2\times$  or  $4\times$ , providing as output a high-resolution 3D image. To our knowledge, there are only a few previous works [1]–[18] that study the super-resolution of CT or MRI images. Similar to some of these previous works [2]–[7], [9], [11]–[18], we approach single-image super-resolution (SISR) of CT and MRI scans using deep convolutional neural networks (CNNs). We propose a CNN architecture composed of 10 convolutional layers and an intermediate sub-pixel convolutional (upsampling) layer [19] that is placed after the first 6 convolutional layers. Different from related works [3], [6], [16], [18] that use the sub-pixel convolutional layer of Shi *et al.* [19], we add 4 convolutional layers after the upsampling layer. In order to obtain 3D super-resolution, we employ two CNNs with similar architectures, as illustrated in Figure 1. The first CNN increases the resolution on two axes (width and height), while the second CNN takes the output from the first CNN and further increases the resolution on the third axis (depth). Different from related methods [3], [6], [16], [18], we compute the loss with respect to the ground-truth high-resolution image right after the upsampling layer, in addition to computing the loss after the last convolutional layer. The intermediate loss forces our network to produce a better output, closer to the ground-truth. In order to improve the results and obtain sharper images, a common approach is to apply Gaussian smoothing on top of the input images, using a fixed standard deviation. Different from other

medical image super-resolution methods [3], [14], [17], we use various standard deviations in order to avoid overfitting to a certain standard deviation and improve the generalization capacity of our model.

We note that our model belongs to a class of deep neural networks known as fully convolutional neural networks. The main advantage of using such models, which do not include dense (fully-connected) layers, is that the input samples do not have to be of the same size. This flexibility enables a broad range of applications such as image segmentation [20], object tracking [21], crowd detection [22], time series classification [23] and single-image super-resolution [3], [6], [16], [18]. Different from other fully convolutional neural networks [3], [6], [16], [18], [20]–[23], our network is specifically designed for SISR, having a custom architecture that includes an upsampling layer [19] useful only for SISR, as well as a novel loss function.

We conduct super-resolution experiments on two databases of 3D CT and MRI images, one gathered from the Colțea Hospital (CH) and one that is publicly available online, known as NAMIC.<sup>1</sup> We compare our method with several interpolation baselines (nearest, bilinear, bicubic, Lanczos) and state-of-the-art methods [3], [13], [15], [17], in terms of the peak signal-to-noise ratio (PSNR), the structural similarity index (SSIM) and the information fidelity criterion (IFC). We perform comparative experiments on both 2D and 3D single-image super-resolution under commonly-used upsampling factors, namely  $2\times$  and  $4\times$ . The empirical results indicate that our approach is able to surpass all the other methods included in the experiments. For example, on the NAMIC data set, we obtain a PSNR of 40.57 and an SSIM of 0.9835 for 3D super-resolution by a factor of  $2\times$ , while

<sup>1</sup> Available at <http://hdl.handle.net/1926/1687>.

Pham *et al.* [13] reported a PSNR of 38.28 and an SSIM of 0.9781 in the same setting. Furthermore, we conduct a subjective image quality assessment by human observers, asking 6 doctors and 12 regular annotators to choose between the CT images produced by our method and those produced by Lanczos interpolation (the best interpolation method). The annotators opted for our method in favor of Lanczos interpolation in 97.55% cases for an upscaling factor of  $2\times$  and in 96.69% cases for an upscaling factor of  $4\times$ . These results indicate that our method is significantly better than Lanczos interpolation. In order to allow further developments and results replication, we provide our code as open source in a public repository.<sup>2</sup>

To summarize, our contribution is threefold:

- We propose a novel CNN model for 3D super-resolution of CT and MRI scans, which is based on an intermediate loss added to the standard output loss and on smoothing the input using random standard deviations for the Gaussian blur.
- We conduct a subjective image quality assessment by human observers to determine the quality and the utility of our super-resolution results, as in [15].
- We provide our code online for download, allowing our results to be easily replicated.

We organize the rest of this paper as follows. We present related work in Section II. We describe our method in detail in Section III. We present experiments and results in Section IV. Finally, we draw our conclusions in Section V.

## II. RELATED WORK

The purpose of SISR is to reconstruct a high-resolution (HR) image from its low-resolution (LR) counterpart. Before the deep learning age, researchers have used exemplar-based or sparse coding methods for SISR. Exemplar-based methods learn mapping functions from external LR and HR exemplar pairs [24]–[26]. Sparse coding methods [27] are representative for external exemplar-based SR methods. For example, the method of Yang *et al.* [27] builds a dictionary with LR patches and the corresponding HR patches.

To our knowledge, the first work to present a deep learning approach for SISR is [28]. Dong *et al.* [28] proposed a CNN composed of 8 convolutional layers. The network was trained in an end-to-end fashion, minimizing the reconstruction error between the HR image and the output of the network. They used bicubic interpolation to resize the image, before giving it as input to the network. Hence, the CNN takes a blurred HR image as input and learns how to make it sharper. Since the input is an HR image, this type of CNN is time consuming. Therefore, Shi *et al.* [19] introduced a new method for upsampling the image using the CNN activation maps produced by the last layer. Their network is more efficient, because it builds the HR image only at the very end. Other works, such as [29], proposed deeper architectures, focusing strictly on accuracy. Indeed, Zhang *et al.* [29] presented one of the

deepest CNNs used for SR, composed of 400 layers. They used a channel attention mechanism and residual blocks to handle the depth of the network.

For medical SISR, some researchers have focused on sparse representations [8], [10], while others on training convolutional neural networks [1]–[7], [9], [11], [12], [14]–[16], [18].

The authors of [8] proposed a weakly-supervised joint convolutional sparse coding method to simultaneously solve the problems of super-resolution and cross-modal image synthesis. In [10], the authors adopted a method based on compressed sensing and self-similarity constraint, obtaining better results than [17] in terms of SSIM and PSNR.

Some works [1], [3], [5], [6], [9]–[11], [14]–[16], [18] focused on 2D upsampling, i.e. on increasing the width and height of CT/MRI slices, while other works [2], [4], [8], [12] focused on 3D upsampling, i.e. on increasing the resolution of full 3D CT/MRI scans on all three axes (width, height and depth).

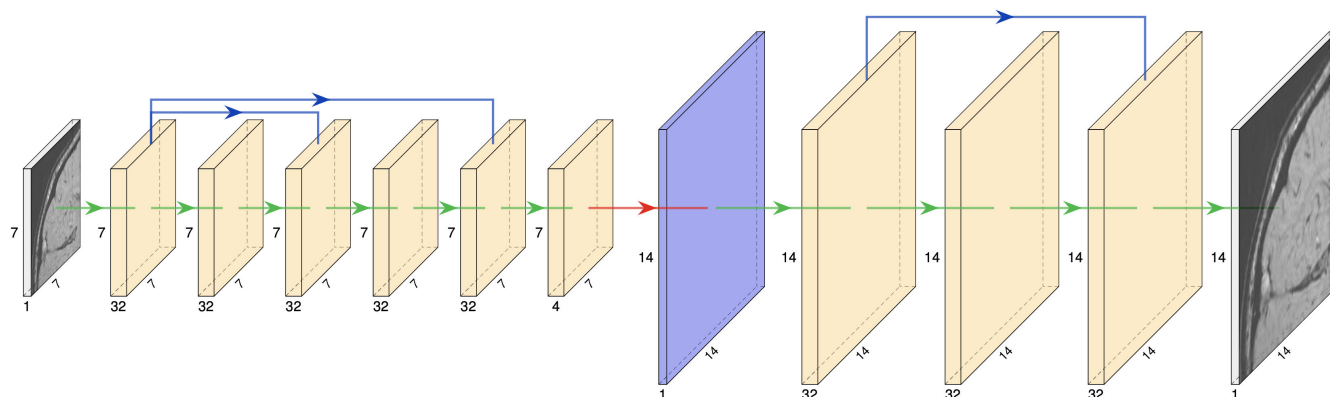
For 2D upsampling, some works [1], [5], [9], [14] used interpolated low resolution (ILR) images, while other works [3], [6], [16], [18] used the efficient sub-pixel convolutional neural network (ESPCN) introduced in [19]. Similar to the latter approaches [3], [6], [16], [18], we employ the sub-pixel convolutional layer of Shi *et al.* [19]. Different from these related works [3], [6], [16], [18], we add a convolutional block after the sub-pixel convolutional layer, in order to enhance the HR output image. Furthermore, we propose a novel loss function for our CNN model. Instead of computing the loss between the output image and the ground-truth high-resolution image, we also compute the loss between the intermediate image given by the sub-pixel convolutional layer and the high-resolution image. This forces our neural network to learn a better intermediate representation, increasing its performance.

There are some works [2], [11], [15] that employed generative adversarial networks (GANs) [30] to upsample medical images. Although our approach based on fully convolutional neural networks is less related to GAN-based SISR methods, we decided to include the approach of You *et al.* [15] in our experiments, as a recent and relevant baseline.

For 3D upsampling, Chen *et al.* [2] trained a CNN with 3D convolutions and used a GAN-based loss function to produce sharper and more realistic images. In order to upsample a 3D image, Du *et al.* [4] employed a deconvolutional layer composed of 3D filters to upsample the LR image, in an attempt to reduce the computational complexity. As [2], [4], [8], [12], we tackle the problem of 3D CT/MRI image super-resolution. However, instead of using inefficient 3D filters to upsample the LR images in a single step, we propose a more efficient two-stage approach that uses 2D filters. Our approach employs a CNN to increase the resolution in width and height, and another CNN to further increase the resolution depth-wise.

Most SISR works [3], [14], [17], apply Gaussian smoothing using a fixed standard deviation on the training images,

<sup>2</sup>Available at <https://github.com/lilygeorgescu/3d-super-res-cnn>.



**FIGURE 2.** Our convolutional neural network for super-resolution on two axes, height and width. The network is composed of 10 convolutional layers and an upsampling (sub-pixel convolutional) layer. It takes as input low-resolution patches of  $7 \times 7$  pixels and, for the  $r = 2$  scale factor, it outputs high-resolution patches of  $14 \times 14$  pixels. The convolutional layers are represented by green arrows. The sub-pixel convolutional layer is represented by the red arrow. The long-skip and short-skip connections are represented by blue arrows. Best viewed in color.

thus training the models in more difficult conditions. However, we believe that using a fixed standard deviation can harm the performance, as deep models tend to overfit to the training data. Different from the standard methodology, each time we apply smoothing on a training image, we chose a different standard deviation, randomly. This simple change improves the generalization capacity of our model, yielding better performance at test time.

While many works focus only on the super-resolution task, the work of Sert *et al.* [1] is focused on the gain brought by the upsampled images in solving a different task. Indeed, the authors [1] obtained an improvement of 7.5% in the classification of segmented brain tumors when the upsampled images were used.

We note that there is also some effort in designing and obtaining CT scan results of higher resolution directly from CT scanners. For example, X-ray microtomography (micro-CT) [31], which is based on pixel sizes of the cross-sections in the micrometer range, has applications in medical imaging [32], [33]. Another alternative to standard (single-energy) CT is dual-energy or multi-energy CT [34], [35]. Different from the expensive alternatives such as dual-energy CT and micro-CT, our approach to increasing the resolution of single-energy CT images using a machine learning algorithm represents an economical and accessible mode.

### III. METHOD

Our method for solving the 3D image super-resolution problem relies on deep neural networks. The universal approximation theorem [36], [37] states that neural networks with at least one hidden layer are universal function approximators. Hence, the hypothesis class represented by neural networks is large enough to accommodate any hypothesis explaining the data. This high model capacity seems to help deep neural networks in attaining state-of-the-art results in many domains [38], including image super-resolution [28]. This is the main reason behind our decision to use neural networks. Interestingly, Dong *et al.* [28] show that deep

convolutional neural networks for super-resolution are equivalent to previously-used sparse-coding methods [27]. However, the recent literature indicates that deep neural networks attain better results than handcrafted methods in practice [28], [39], mainly because the parameters (weights) are learned from data in an end-to-end fashion. Further specific decisions, such as the neural architecture or the loss function, are taken based on empirical observations.

Our approach is divided into two stages, as illustrated in Figure 1. In the first stage, we upsample the image on height and width using a deep fully convolutional neural network. Then, in the second stage, we further upsample the resulting image on the depth axis using another fully convolutional neural network. Therefore, our complete method is designed for resizing the 3D input volume on all three axes. While the CNN used in the first stage resizes the image on two axes, the CNN used in the second stage resizes the image on a single axis. Both CNNs share the same architecture, the only difference being in the upsampling layer (the second CNN upsamples in only one direction). At training time, our CNNs operate on patches. However, since the architecture is fully convolutional, the models can operate on entire slices at inference time, for efficiency reasons.

We hereby note that 3D super-resolution is not equivalent to 2D super-resolution in a slice-by-slice order. Our first CNN performs 2D super-resolution in a slice-by-slice order, increasing an input of size  $h \times w \times d$  to the size  $r \cdot h \times r \cdot w \times d$ , where  $r$  is the scale factor. Since we end up with the same number of slices ( $r$ ), this is not enough. This is why we need the second CNN to further increase the image from  $r \cdot h \times r \cdot w \times d$  voxels to  $r \cdot h \times r \cdot w \times r \cdot d$  voxels. The final output of  $r \cdot h \times r \cdot w \times r \cdot d$  voxels could also be obtained by employing a single CNN with 3D convolutions. In most of our convolutional layers, each 2D convolutional filter is formed of  $3 \cdot 3 \cdot 32 + 1 = 289$  weights to be learned. As we employ two networks for 3D super-resolution, we learn  $2 \cdot 289 = 578$  weights. For an equivalent model based on 3D convolutions, each 3D convolutional filter would be formed



of  $3 \cdot 3 \cdot 3 \cdot 32 + 1 = 865$  weights. This analysis proves that our two CNNs put together have less weights than a single 3D CNN. We thus conclude that our approach is more efficient. We note that our approach is essentially based on decomposing the 3D convolutional filters in a product of two 2D convolutional filters. We note that the same principle is applied in literature [40], [41] to build more efficient CNN models by decomposing 2D convolutional layers in two subsequent 1D convolutional layers that operate on independent dimensions.

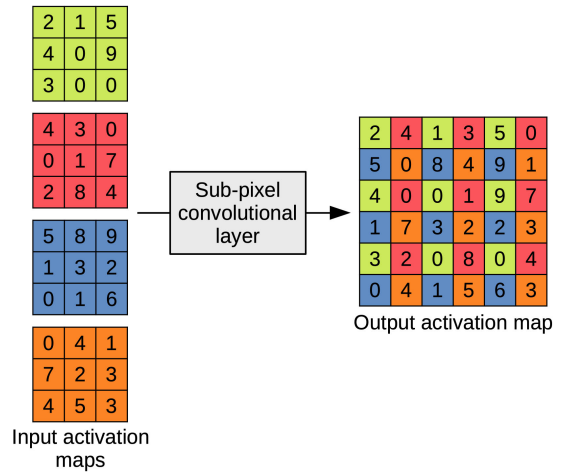
We further describe in detail the proposed CNN architecture, loss function and data augmentation procedure.

### A. ARCHITECTURE

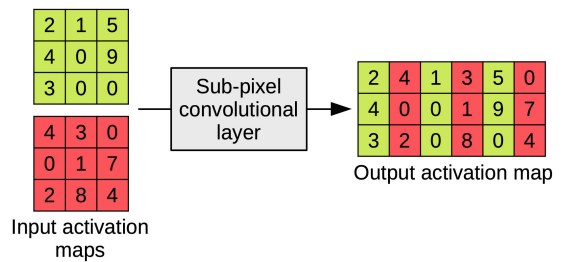
Our architecture, depicted in Figure 2 and used for both CNNs, is composed of 10 convolutional (conv) layers, each followed by Rectified Linear Units (ReLU) [42] activations. We decided to use ReLU activations, as this represents the most popular choice of transfer function in current research based on deep learning. All convolutional layers contain filters with a spatial support of  $3 \times 3$ . While older deep models were based on larger filters, e.g. the AlexNet [43] architecture contains filters of  $11 \times 11$ , the recent trend is towards using smaller filters, e.g. the ResNet [44] architecture does not contain filters larger than  $3 \times 3$ .

Our 10 conv layers are divided into two blocks. The first block, formed of the first 6 conv layers, starts with the input of the neural network and ends just before the upscaling layer. Each of the first 5 convolutional layers are formed of 32 filters. For the CNN used in the first stage, the number of filters in the sixth convolutional layer is equal to the square of the scale factor, e.g. for a scale factor of  $4 \times$  the number of filters is 16. For the CNN used in the second stage, the number of filters in the sixth convolutional layer is equal to the scale factor, e.g. for a scale factor of  $4 \times$  the number of filters is 4. The difference is caused by the fact that the first CNN upscales on two axes, while the second CNN upscales on one axis. The first convolutional block contains a short-skip connection, from the first conv layer to the third conv layer, and a long-skip connection, from the first conv layer to the fifth conv layer.

The first convolutional block is followed by a sub-pixel convolutional (upsampling) layer, which was introduced in [19]. In the upscaling layer, the activation maps produced by the sixth conv layer are assembled into a single activation map. Throughout the first convolutional block, the spatial size of the low-resolution input is preserved, i.e. the activation maps of the sixth conv layer have  $h_I \times w_I$  components, where  $h_I$  and  $w_I$  are the height and the width of the input image  $I$ . In order to increase the input  $r$  times on both axes, the output of the sixth conv layer must be a tensor of  $h_I \times w_I \times r^2$  components. The activation map resulting from the sub-pixel convolutional layer is a matrix of  $(h_I \cdot r) \times (w_I \cdot r)$  components. For super-resolution on two axes, the pixels are rearranged as shown in Figure 3. In a similar fashion, we can increase the input  $r$  times on one axis. In this case, the output of the



**FIGURE 3.** An example of low-resolution input activation maps and the corresponding high-resolution output activation map given by the sub-pixel convolutional layer for upscaling on two axes. For a scaling factor of  $r = 2$  in both directions, the sub-pixel convolutional layer requires  $r^2 = 4$  activation maps as input. Best viewed in color.



**FIGURE 4.** An example of low-resolution input activation maps and the corresponding high-resolution output activation map given by the sub-pixel convolutional layer for upscaling on one axis. For a scaling factor of  $r = 2$  in one direction, the sub-pixel convolutional layer requires  $r = 2$  activation maps as input. Best viewed in color.

sixth conv layer must be a tensor of  $h_I \times w_I \times r$  components. This time, the activation map resulting from the sub-pixel convolutional layer can be either a matrix of  $(h_I \cdot r) \times w_I$  components or a matrix of  $h_I \times (w_I \cdot r)$  components, depending on the direction we aim to increase the resolution. For super-resolution on one axis, the pixels are rearranged as shown in Figure 4. To our knowledge, we are the first to propose a sub-pixel convolutional (upsampling) layer for super-resolution in one direction.

When Shi *et al.* [19] introduced the sub-pixel convolutional layer, they used it as the last layer of their CNN. Hence, the output depth of the upscaling layer is equal to the number of channels in the output image. Different from Shi *et al.* [19], we employ further convolutions after the upscaling layer. Nevertheless, since we are working with CT/MRI (grayscale) images, our final output has a single channel. In our architecture, the upscaling layer is followed by our second convolutional block, which starts with the seventh convolutional layer and ends with the tenth convolutional layer. The first three conv layers in our second block are formed of 32 filters. The tenth conv layer contains a single convolutional filter,

since our output is a grayscale image that has a single channel. The second convolutional block contains a short skip connection, from the seventh conv layer to the ninth conv layer. The spatial size of  $h_O \times w_O$  components of the activation maps is preserved throughout the second convolutional block, where  $h_O$  and  $w_O$  are the height and the width of the output image  $O$ .

## B. LOSSES AND OPTIMIZATION

In order to obtain a CNN model for single-image super-resolution, the aim is to minimize the differences between the ground-truth high-resolution image and the output image provided by the CNN. Researchers typically employ the mean absolute difference as the objective function. Given a low-resolution input image  $I$  and the corresponding ground-truth high-resolution image  $O$ , the loss based on the mean absolute value is formally defined as follows:

$$\mathcal{L}(\theta, I, O) = \sum_{i=1}^{w_O} \sum_{j=1}^{h_O} |f(\theta, I) - O|, \quad (1)$$

where  $\theta$  are the CNN parameters (weights),  $f$  is the transformation function learned by the CNN, and  $w_O$  and  $h_O$  represent the width and the height of the ground-truth image  $O$ , respectively.

When we train our CNN model, we do not employ the standard approach of minimizing the difference between the output provided by the CNN and the ground-truth HR image. Instead, we propose a novel approach based on an intermediate loss function. Since the conv layers after the upscaling layer are meant to refine the high-resolution image without taking any additional information from the low-resolution input image, we note that the high-resolution image resulting immediately after the upscaling layer should be as similar as possible to the ground-truth high-resolution image. Therefore, we propose a loss function that aims to minimize the difference between the intermediately-obtained HR image and the ground-truth HR image, in addition to minimizing the difference between the final HR output image and the ground-truth HR image. Let  $f_1$  denote the transformation function that corresponds to the first convolutional block and the upscaling layer, and let  $f_2$  denote the transformation function that corresponds to the second convolutional block. With these notations, the transformation function  $f$  of our full CNN architecture can be written as follows:

$$f(\theta, I) = f_2(\theta_2, f_1(\theta_1, I)), \quad (2)$$

where  $\theta$  are the parameters of the full CNN,  $\theta_1$  are the parameters of the first convolutional block and  $\theta_2$  are the parameters of the second convolutional block, i.e.  $\theta$  is a concatenation of  $\theta_1$  and  $\theta_2$ . Having defined  $f_1$  and  $f_2$  as above, we can formally write our loss function as follows:

$$\mathcal{L}_{full} = \mathcal{L}_{standard} + \lambda \cdot \mathcal{L}_{intermediate}, \quad (3)$$

where  $\lambda$  is a parameter that controls the importance of the intermediate loss with respect to the standard loss,  $\mathcal{L}_{standard}$

is the standard loss defined in Equation (1) and  $\mathcal{L}_{intermediate}$  is defined as follows:

$$\mathcal{L}_{intermediate}(\theta_1, I, O) = \sum_{i=1}^{w_O} \sum_{j=1}^{h_O} |f_1(\theta_1, I) - O|. \quad (4)$$

In the experiments, we set  $\lambda = 1$ , since we did not find any strong reason to assign a lower or higher importance to the intermediate loss. By replacing  $\lambda$  with 1 and the loss values from Equation (1) and Equation (4), Equation (3) becomes:

$$\mathcal{L}_{full}(\theta, I, O) = \sum_{i=1}^{w_O} \sum_{j=1}^{h_O} |f(\theta, I) - O| + \sum_{i=1}^{w_O} \sum_{j=1}^{h_O} |f_1(\theta_1, I) - O|. \quad (5)$$

In order to optimize towards the objective defined in Equation (5), we employ the Adam optimizer [45], which is known to converge faster than Stochastic Gradient Descent.

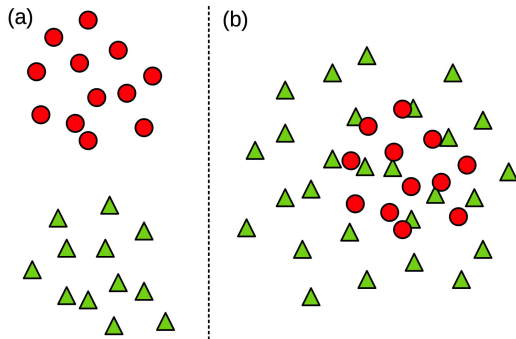
## C. DATA AUGMENTATION

A common approach to force the CNN to produce sharper output images is to apply Gaussian smoothing using a fixed standard deviation at training time [3], [14], [17]. By training the CNN on blurred low-resolution images, the super-resolution task becomes harder. During inference, when the input images are no longer blurred, the task will be much easier. However, smoothing only the training images with a fixed standard deviation will inherently generate a distribution gap between training and test data. If the CNN fits the training distribution well, it might not produce the desired results at inference time. This is because machine learning models are based on the assumption that training data and test data are sampled from the same distribution. We propose to solve this problem by using a randomly-chosen standard deviation for each training image. Although the training data distribution will still be different from the testing data distribution, it will include the distribution of test samples, as illustrated in Figure 5. In order to augment the training data, we apply Gaussian blur with a probability of 0.5 (only half of the images are smoothed) using a kernel of  $3 \times 3$  components and a randomly-chosen standard deviation between 0 and 0.5. In this way, we increase the variance of the training data without introducing any bias. In this case, if the CNN fits the training distribution well, it will produce good super-resolution outputs during inference, since there is no distribution gap between training and test.

## IV. EXPERIMENTS

### A. DATA SETS

The first data set used in the experiments consists of 10 anonymized 3D images of brain CT provided by the Medical School at Colțea Hospital. We further refer to this data set as the Colțea Hospital (CH) data set. In order to fairly train and test our CNN models and baselines, we randomly selected 6 images for training and used the remaining 4 images for



**FIGURE 5.** Distribution of training samples (represented by green triangles) and test samples (represented by red circles), when the training samples are smoothed using a fixed standard deviation (left-hand side) versus using a randomly-chosen standard deviation (right-hand side). In example (a), overfitting on the training data leads to poor results on test data. In example (b), the danger of overfitting is diminished because the test distribution is included in the training distribution. Best viewed in color.

testing. The training set has 359 slices (2D images) and the testing set has 238 slices. The height and the width of the slices vary between 192 and 512 pixels, while the depth of the 3D images varies between 3 and 176 slices. The resolution of a voxel is  $1 \times 1 \times 1 \text{ mm}^3$ .

The second data set used in our experiments is the National Alliance for Medical Image Computing (NAMIC) Brain Multimodality data set. The NAMIC data set consists of 20 3D MRI images, each composed of 176 slices of  $256 \times 256$  pixels. As for the CH data set, the resolution of a voxel is  $1 \times 1 \times 1 \text{ mm}^3$ . For our experiments, we used T1-weighted (T1w) and T2-weighted (T2w) images independently. Following [13], we split the NAMIC data set into a training set containing 10 3D images and a test set containing the other 10 images. We kept the same split for T1w and T2w.

## B. EXPERIMENTAL SETUP

### 1) EVALUATION METRICS

As most previous works [3], [5], [10], [12]–[18], we employed the two most common evaluation metrics, namely the peak signal-to-noise ratio (PSNR) and the structural similarity index (SSIM). The PSNR is the ratio between the maximum possible power of a signal and the power of corrupting noise that affects the fidelity of its representation. Although the PSNR is one of the most used metrics for image reconstruction, some researchers [46], [47] argued that it is not highly indicative of the perceived similarity. The SSIM [46] aims to address this shortcoming by taking contrast, luminance and texture into account. The result of the SSIM is a number between -1 and 1, where a value of 1 means the ground-truth image and the reconstructed image are identical. Similarly, a higher PSNR value indicates a better reconstruction, although the PSNR does not have an upper bound.

Since PSNR and SSIM values cannot guarantee a visually favorable result, we employ an additional metric for the final results, namely the information fidelity criterion (IFC) [48]. Although IFC is scarcely used in literature [15],

**TABLE 1.** Preliminary 2D super-resolution results on the CH data set for an upscaling factor of  $2\times$ . The PSNR and the SSIM values are reported for various patch sizes and different numbers of filters. For models with  $7 \times 7$  patches, we report the inference time (in seconds) per CT slice measured on an Nvidia GeForce 940MX GPU with 2GB of RAM.

Number of filters	Input size	SSIM	PSNR	Time (in seconds)
32	$4 \times 4$	0.9165	35.36	-
32	$7 \times 7$	0.9270	36.22	0.05
32	$14 \times 14$	0.8987	33.83	-
32	$17 \times 17$	0.8934	33.42	-
64	$7 \times 7$	0.9279	36.28	0.08
128	$7 \times 7$	0.9276	36.28	0.16

Yang *et al.* [49] pointed out that IFC is correlated well with the human perception of SR images. As for PSNR and SSIM, higher IFC values indicate better results.

### 2) IMAGE QUALITY ASSESSMENT BY HUMAN OBSERVERS

Because the above metrics rely only on the pixel values and not on the perceived visual quality, we decided to evaluate our method with the help of human annotators. Although a deep learning method can provide better PSNR, SSIM or IFC values, it might produce artifacts that could be misleading for right diagnostics and treatment. We thus have to make sure that our approach does not produce any unwanted artifacts visible to humans. We conducted the image quality assessment on the CH data set, testing our CNN-based method against Lanczos interpolation. We used CT slices extracted from high-resolution 3D images resulting after applying super-resolution on all three axes. For each upscaling factor,  $2\times$  and  $4\times$ , we extracted 100 CT slices at random from the test set. Hence, each human evaluator had to annotate 200 image pairs (100 for each upscaling factor). For each evaluation sample, an annotator would have seen the original image in the middle and the two reconstructed images on its sides, one on the left side and the other on the right side. The annotators had a magnifying glass tool that allowed them to look at details and discover artifacts. The locations (left or right) of the images reconstructed by our CNN and by Lanczos interpolation were randomly picked every time. To prevent any form of cheating, the randomly picked locations were unknown to the annotators. For each test sample, we asked each annotator to select the image that best reconstructed the original image. Our experiment was completed by 18 human annotators, 6 of them being doctors specialized in radiotherapy and oncology. In total, we collected 3600 annotations (18 annotators  $\times$  200 samples).

### 3) BASELINES

We compared our method with standard resizing methods based on various interpolation schemes, namely nearest neighbors, bilinear, bicubic and Lanczos. In addition to these baselines, we compared with three methods [3], [15], [17] that focused on 2D SISR and one method [13] that focused on 3D SISR. We note that You *et al.* [15] did not report results on NAMIC. Nonetheless, You *et al.* [15] were kind to provide access to their source code. We thus applied their method

**TABLE 2.** Ablation 2D super-resolution results on the CH data set for an upscaling factor of  $2\times$ . The PSNR and the SSIM values are reported various ablated versions of our CNN model. The best results are highlighted in bold. Results of ablated models marked with † are significantly worse than our complete model, according to paired McNemar’s testing [50] for the significance level 0.001.

Second conv block	Intermediate loss	Short-skip connections	Long-skip connection	Variable standard deviation	SSIM	PSNR
X	X	X	X	X	0.9224 <sup>†</sup>	35.58 <sup>†</sup>
X	X	✓	✓	✓	0.9236 <sup>†</sup>	35.94 <sup>†</sup>
✓	X	✓	✓	✓	0.9256 <sup>†</sup>	36.15 <sup>†</sup>
✓	✓	X	✓	✓	0.9260 <sup>†</sup>	36.17 <sup>†</sup>
✓	✓	✓	X	✓	0.9234 <sup>†</sup>	36.11 <sup>†</sup>
✓	✓	✓	✓	X	0.9236 <sup>†</sup>	35.69 <sup>†</sup>
✓	✓	✓	✓	✓	<b>0.9270</b>	<b>36.22</b>

on both CH and NAMIC data sets, keeping the same settings and hyperparameters as recommended by the authors. As You *et al.* [15], we employed their method only on 2D super-resolution for the  $2\times$  upscaling factor. For the other three baselines, we included the NAMIC scores reported in the respective articles [3], [13], [17].

### C. PARAMETER TUNING AND PRELIMINARY RESULTS

We conducted a series of preliminary experiments to determine the optimal patch size as well as the optimal width (number of convolutional filters) for our CNN. In order to find the optimal patch size, we tried out patches of  $4\times 4$ ,  $7\times 7$ ,  $14\times 14$  and  $17\times 17$  pixels. In term of the number of filters, we tried out values in the set {32, 64, 128} for all conv layers in our network. These parameters were tuned in the context of 2D super-resolution on the CH data set. The corresponding results are presented in Table 1. First of all, we note that our method produces better SSIM and PSNR values, i.e. 0.9270 and 36.22, for patches of  $7\times 7$  pixels. Second of all, we observe that adding more filters on the conv layers slightly increases the SSIM and PSNR values. However, the gains in terms of SSIM and PSNR come with a great cost in terms of time. For example, using 128 filters on each conv layer triples the processing time in comparison with using 32 filters on each conv layer. For the subsequent experiments, we thus opted for patches of  $7\times 7$  pixels and conv layers with 32 filters.

We believe that it is important to note that, although the number of training CT slices is typically in the range of a few hundreds, the number of training patches is typically in the range of hundreds of thousands. For instance, the number of  $7\times 7$  training patches extracted from the CH data set for the  $2\times$  upscaling factor is 326,000. We thus stress out that the number of training samples is high enough to train highly-accurate deep learning models.

During training, we used mini-batches of 128 images throughout all the experiments. In a set of preliminary experiments, we did not observe any significant differences when using mini-batches of 64 or 256 images. In each experiment, we trained the CNN for 40 epochs, starting with a learning rate (step size) of  $10^{-3}$  and decreasing the learning rate to  $10^{-4}$  after the first 20 epochs.

### D. ABLATION STUDY RESULTS

We performed an ablation study to emphasize the effect of various components over the overall performance.

The ablation results obtained on the CH data set for super-resolution on height and width by a factor of  $2\times$  are presented in Table 2.

In our first ablation experiment, we have eliminated all the enhancements in order to show the performance level of a baseline CNN on the CH data set. Since there are several SISR works [3], [6], [16], [18] based on the standard ESPCN model [19], we have eliminated the second convolutional block in the second ablation experiment, transforming our architecture into a standard ESPCN architecture. The performance drops from 0.9270 to 0.9236 in terms of SSIM and from 36.22 to 35.94 in terms of PSNR. In the subsequent ablation experiments, we have removed, in turns, the intermediate loss, the short-skip connections and the long-skip connection. The results presented in Table 2 indicate that all these components are relevant to our model, bringing significant performance benefits in terms of both SSIM and PSNR. In our last ablation experiment, we used a fixed standard deviation instead of a variable one for the Gaussian blur added on training patches. We notice that our data augmentation approach based on a variable standard deviation brings the highest gains in terms of SSIM (from 0.9236 to 0.9270) and PSNR (from 35.69 to 36.22), with respect to the other ablated components.

Since the differences in terms of PSNR or SSIM for the ablated models are hard to quantify as small or large with respect to the complete CNN, we conducted paired McNemar’s significance testing [50] to determine if the differences are statistically significant or not. We considered a p-value threshold of 0.001 for our statistical testing. Every ablated model that is significantly different from the complete model is marked with † in Table 2. We note that the complete CNN is significantly better than each ablated version, although the actual differences in terms of PSNR or SSIM might seem small. We thus conclude that all the proposed enhancements provide significant performance gains.

### E. RESULTS ON CH DATA SET

We first compared our CNN-based model with a series of interpolation baselines and a state-of-the-art method [15] on the CH data set. We present the results for super-resolution on two axes (width and height) in Table 3. Among the considered baselines, it seems that the Lanczos interpolation method provides better results than the bicubic, the bilinear or the nearest neighbor methods. Our CNN model is able to surpass



**TABLE 3. 2D super-resolution results of our CNN model versus a state-of-the-art method [15] and several interpolation baselines on the CH data set. The PSNR, the SSIM and the IFC values are reported for two upscaling factors, 2× and 4×. The best result on each column is highlighted in bold.**

Method	2×			4×		
	SSIM	PSNR	IFC	SSIM	PSNR	IFC
Nearest neighbor	0.8893	32.71	4.40	0.7659	29.06	1.32
Bilinear	0.8835	33.34	3.73	0.7725	29.73	1.49
Bicubic	0.9077	34.71	4.59	0.7965	30.41	1.72
Lanczos	0.9111	35.08	4.93	0.8012	30.57	1.84
You et al. [15]	0.8874	32.73	4.40	-	-	-
Our CNN model	<b>0.9291</b>	<b>36.39</b>	<b>5.36</b>	<b>0.8308</b>	<b>31.59</b>	<b>1.92</b>

**TABLE 4. 3D super-resolution results of our CNN model versus several interpolation baselines on the CH data set. The PSNR, the SSIM and the IFC values are reported for two upscaling factors, 2× and 4×. The best result on each column is highlighted in bold.**

Method	2×			4×		
	SSIM	PSNR	IFC	SSIM	PSNR	IFC
Nearest neighbor	0.8430	30.36	2.19	0.7152	27.32	0.72
Bilinear	0.8329	30.72	2.18	0.7206	27.93	0.95
Bicubic	0.8335	26.51	2.47	0.7200	24.05	1.04
Lanczos	0.8423	27.85	2.58	0.7263	25.06	1.09
Our CNN model	<b>0.8926</b>	<b>33.04</b>	<b>2.83</b>	<b>0.7819</b>	<b>29.36</b>	<b>1.13</b>

all baselines for both upscaling factors, 2× and 4×. Compared to the best interpolation method (Lanczos), our method is 0.0180 better in terms of SSIM, 1.31 better in terms of PSNR and 0.43 better in terms of IFC. Furthermore, our CNN provides superior results to the GAN-based method of You *et al.* [15].

We note that, in Table 2, we reported an SSIM of 0.9270 and a PSNR of 36.22 for our method, while in Table 3, we reported an SSIM of 0.9291 and a PSNR of 36.39. In order to boost the performance of our method in accordance with the observed differences between Tables 2 and 3, we employed the self-ensemble strategy used by Lim *et al.* [51]. For each input image, the self-ensemble strategy consists in generating additional images using geometric transformations, e.g. rotations and flips. Following Lim *et al.* [51], we generated 7 augmented images from the LR input image, upsampling all 8 images (the original image and the 7 additional ones) using our CNN. We then applied the inverse transformations to the resulting 8 HR images in order to obtain 8 output images that are aligned with the ground-truth HR images. The final output image is obtained by taking the median of the HR images. In the following experiments on CH and NAMIC data sets, the reported results always include the described self-ensemble strategy.

We provide the results for super-resolution on all three axes in Table 4. First of all, we notice that the SSIM, the PSNR and the IFC values are lower for all methods when dealing with 3D super-resolution (Table 4) instead of 2D super-resolution (Table 3). This shows that the task of 3D super-resolution is much harder than 2D super-resolution. This is an expected result, considering that the dimensionality of the reconstruction space increases significantly for 3D super-resolution, i.e. there are many more HR outputs corresponding to a single LR input, while the training data is the same. Nevertheless,

our method exhibits smaller performance drops when going from 2D super-resolution to 3D super-resolution. As for the 2D super-resolution experiments on CH data set, our CNN model for 3D super-resolution is superior to all baselines for both upscaling factors. We thus conclude that our CNN model is better than all interpolation baselines on the CH data set, for both 2D and 3D super-resolution and for all upscaling factors.

## F. RESULTS ON NAMIC DATA SET

On the NAMIC data set, we compared our method with the best-performing interpolation method on the CH data set, namely Lanczos interpolation, as well as some state-of-the-art 2D [3], [15], [17] and 3D [13] super-resolution methods. We note that most previous works, including [3], [13], [17], used bicubic interpolation as a relevant baseline. Unlike these works, we opted for Lanczos interpolation, which provided better results than bicubic interpolation and other interpolation methods on the CH data set.

We first present the 2D super-resolution results in Table 5. The 2D SR results indicate that the GAN-based method of You *et al.* [15] is superior to the CNN baselines [3], [17]. However, none of the state-of-the-art methods [3], [15], [17] is able to attain better performance than Lanczos interpolation. This proves that Lanczos interpolation is a much stronger baseline. Among the deep learning methods, our CNN is the only one to surpass Lanczos interpolation for 2D SR on NAMIC. We believe that this result is noteworthy.

We also present the 3D super-resolution results in Table 6. The 3D SR results show that the approach of Pham *et al.* [13] is better than Lanczos interpolation, which is remarkable. Our CNN is even better, surpassing both the Lanczos interpolation and the approach of Pham *et al.* [13].

As for the CH data set, we observe that the PSNR, the SSIM and the IFC scores for 2D super-resolution are higher than the corresponding scores for 3D super-resolution. The same explanation applies to the NAMIC data set, i.e. the CNNs have to produce likely reconstruction patterns in a much larger space.

While some of the considered state-of-the-art methods [3], [13], [17] presented results only for some cases on NAMIC, either 2D super-resolution on T1w images or 3D super-resolution on T2w images, we provide our results for all possible cases. We note that our CNN model surpasses Lanczos interpolation in each and every case. Furthermore, our model provides superior results than all the state-of-the-art methods [3], [13], [15], [17] considered in our evaluation on the NAMIC data set.

In addition to the quantitative results shown in Tables 5 and 6, we present qualitative results in Figure 6. We selected 5 examples of 2D super-resolution results generated by Lanczos interpolation, by the GAN-based method of You *et al.* [15] and by our CNN model. A close inspection reveals that our results are generally sharper than those of Lanczos interpolation and those of You *et al.* [15]. As also confirmed by the SSIM, the PSNR and the IFC values presented in Tables 5 and 6, the images generated by our CNN

**TABLE 5.** 2D super-resolution results of our CNN model versus several state-of-the-art methods [3], [15], [17] and the Lanczos interpolation baseline on the NAMIC data set. For Zeng et al. [17], we included results for both single-channel super-resolution (SCSR) and multi-channel super-resolution (MCSR). The PSNR, the SSIM and the IFC values are reported for both T1w and T2w images and for two upscaling factors, 2× and 4×. The best results on each column are highlighted in bold.

Method	T1-weighted						T2-weighted					
	2×			4×			2×			4×		
	SSIM	PSNR	IFC	SSIM	PSNR	IFC	SSIM	PSNR	IFC	SSIM	PSNR	IFC
Lanczos	0.9635	37.03	3.06	0.8955	32.71	1.28	0.9788	39.64	3.22	0.9143	33.80	1.37
Zeng et al. [17] (SCSR)	0.9220	36.86	—	0.7120	28.33	—	—	—	—	—	—	—
Zeng et al. [17] (MCSR)	0.9450	38.32	—	0.8110	30.84	—	—	—	—	—	—	—
Du et al. [3]	0.9390	37.21	—	0.7370	29.05	—	—	—	—	—	—	—
You et al. [15]	0.9448	35.09	2.80	—	—	—	0.9594	36.19	2.89	—	—	—
Our CNN model	<b>0.9775</b>	<b>39.29</b>	<b>3.56</b>	<b>0.9193</b>	<b>33.74</b>	<b>1.29</b>	<b>0.9882</b>	<b>42.20</b>	<b>3.79</b>	<b>0.9382</b>	<b>34.86</b>	<b>1.39</b>

**TABLE 6.** 3D super-resolution results of our CNN model versus a state-of-the-art method [13] and the Lanczos interpolation baseline on the NAMIC data set. The PSNR, the SSIM and the IFC values are reported for both T1w and T2w images and for two upscaling factors, 2× and 4×. The best results on each column are highlighted in bold.

Method	T1-weighted						T2-weighted					
	2×			4×			2×			4×		
	SSIM	PSNR	IFC	SSIM	PSNR	IFC	SSIM	PSNR	IFC	SSIM	PSNR	IFC
Lanczos	0.9423	35.72	2.00	0.8690	31.81	0.95	0.9615	37.80	2.29	0.8829	32.08	1.03
Pham et al. [13]	—	—	—	—	—	—	0.9781	38.28	—	—	—	—
Our CNN model	<b>0.9687</b>	<b>37.85</b>	<b>2.38</b>	<b>0.9050</b>	<b>32.88</b>	<b>0.99</b>	<b>0.9835</b>	<b>40.57</b>	<b>2.67</b>	<b>0.9251</b>	<b>33.54</b>	<b>1.10</b>

**TABLE 7.** Image quality assessment results collected from 6 doctors and 12 regular annotators, for the comparison between our CNN-based method versus Lanczos interpolation. For each upscaling factor, each annotator had to select an option for a number of 100 image pairs. To prevent cheating, the randomly picked locations (left or right) for the generated HR images were unknown to the annotators.

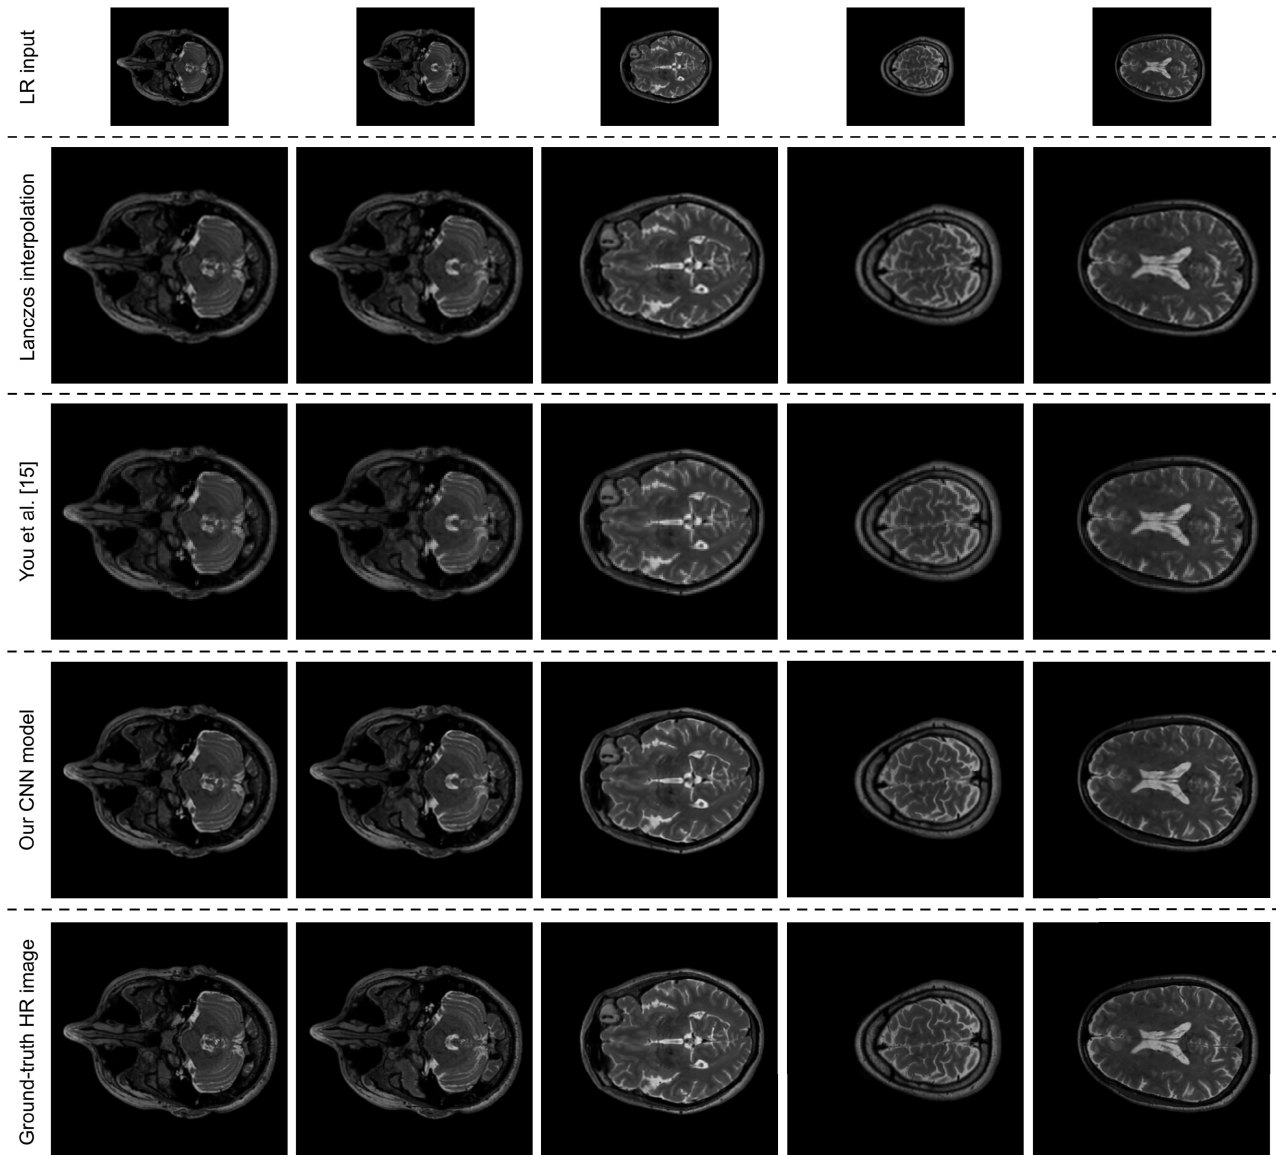
Annotator ID	2×		4×	
	Our CNN	Lanczos	Our CNN	Lanczos
Doctor #1	100	0	100	0
Doctor #2	85	15	100	0
Doctor #3	100	0	100	0
Doctor #4	86	14	71	29
Doctor #5	100	0	68	32
Doctor #6	95	5	95	5
Person #1	100	0	99	1
Person #2	97	3	100	0
Person #3	98	2	100	0
Person #4	98	2	100	0
Person #5	100	0	98	2
Person #6	100	0	76	24
Person #7	100	0	93	7
Person #8	99	1	98	2
Person #9	100	0	98	2
Person #10	100	0	98	2
Person #11	98	2	99	1
Person #12	100	0	98	2
Overall in %	97.55%	2.45%	96.69%	3.31%

are closer to the ground-truth images. At the scale factor of 2× considered in Figure 6, our CNN does not produce any patterns or artifacts that deviate from the ground-truth.

### G. IMAGE QUALITY ASSESSMENT RESULTS

We provide the outcome of the subjective image quality assessment by human observers in Table 7. The study reveals that both doctors and regular annotators opted for our approach in favor of Lanczos interpolation at an overwhelming rate (97.55% at the 2× scale factor and 96.69% at the

4× scale factor). For the 2× scale factor, 10 out of 18 annotators preferred the output of our CNN in all the 100 presented cases. We note that doctors #2 and #4 opted for Lanczos interpolation in 15 and 14 cases (for the 2× scale factor), respectively, which was not typical to the other annotators. Similarly, for the 4× scale factor, there are 3 annotators (doctor #4, doctor #5 and person #6) that seem to prefer Lanczos interpolation at a higher rate than the other annotators. After discussing with the doctors about their choices, we discovered that, in most cases, they prefer the sharper output of our CNN. However, the CNN seems to introduce some reconstruction patterns (learned from training data) that do not correspond exactly to the ground-truth. This phenomenon seems to be more prevalent at the 4× scale factor, although the phenomenon is still rarely observed. This explains why doctors #4 and #5 preferred Lanczos interpolation in more cases than the other doctors, although the majority of their votes are still in favor of our CNN. When they opted for Lanczos interpolation, they considered that it is safer to consider its blurred and less informative output. In trying to find an explanation for these reconstruction patterns, we analyzed the output of the CNN without data augmentation. We observed such reconstruction patterns even when training data augmentation was removed, ruling out this hypothesis. Given a low-resolution input patch, the CNN finds the most likely high-resolution patch corresponding to the input. This likelihood is learned by the CNN when it is trying to minimize the loss over the entire training set. Although producing the most probable output works well in most cases, using a machine learning model, e.g. a CNN, is not the perfect solution. The explanation becomes clear if we consider that multiple HR patches can correspond to the same LR input patch and that choosing the most likely HR patch is not always the right answer. We thus conclude that our CNN suffers from the same



**FIGURE 6.** Image super-resolution examples selected from the NAMIC data set. In order to obtain the input images of  $128 \times 128$  pixels, the original NAMIC images were downsampled by a scale factor of  $2\times$ . HR images of  $256 \times 256$  pixels generated by Lanczos interpolation, by the GAN-based method of You *et al.* [15] and by our CNN model are compared with the original (ground-truth) HR images.

problem as any other machine learning model. Furthermore, we stress out that the reconstruction patterns in question are plausible from a biological point of view, i.e. the doctors were able to spot them only by comparing the HR output with the ground-truth HR image. We note that these patterns should not be mistaken with artifacts that could be caused by underfitting or a poor architectural choice. Our CNN does not introduce such artifacts.

Based on our subjective image quality assessment, we concluded with the doctors that going beyond the  $4\times$  scale factor, solely with a method based on algorithmic super-resolution, is neither safe (a CNN might introduce too many patterns far from the ground-truth) nor helpful (a standard interpolation method is not informative). However, the doctors agree that either super-resolution method is desirable in favor of the

input low resolution images. Therefore, in order to reach the scale factor of  $10\times$  desired by the doctors, we have to look in other directions in future work. A promising direction is to combine multiple inputs, e.g. by using CT and MRI scans of the same person or by using CT/MRI scans taken at different moments in time (before and after the contrast agent is introduced).

## V. CONCLUSION

In this paper, we have presented an approach based on fully convolutional neural networks for the super-resolution of CT/MRI scans. Our method is able to reliably upscale 3D CT/MRI image up to a scale factor of  $4\times$ . We have compared our approach with several baseline interpolation and state-of-the-art methods [3], [13], [15], [17]. The empirical

results indicated that our approach provides superior results on both CH and NAMIC data sets. We have also conducted a subjective image quality assessment by human observers, showing that our method is significantly better than Lanczos interpolation. The subjective image quality assessment also revealed the limitations of a pure algorithmic approach. The doctors invited to take part in our study concluded that going to a scale factor higher than  $4\times$  requires alternative solutions. In future work, we aim to continue our research by extending the proposed CNN method to multi-channel input. This will likely help us in achieving higher upscaling factors, e.g.  $10\times$ , required for the accurate diagnostics and treatment of cancer, an actively studied and extremely important research topic [52], [53].

## ACKNOWLEDGMENT

The authors would like to thank the reviewers for their helpful comments. In memory of Tatiana-Mihaela Ionescu, who offered her CT scans to facilitate our study, but died of cancer.

## REFERENCES

- [1] E. Sert, F. Özyurt, and A. Doğantekin, "A new approach for brain tumor diagnosis system: Single image super resolution based maximum fuzzy entropy segmentation and convolutional neural network," *Med. Hypotheses*, vol. 133, Dec. 2019, Art. no. 109413.
- [2] 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 *Proc. MICCAI*, 2018, pp. 91–99.
- [3] X. Du and Y. He, "Gradient-guided convolutional neural network for MRI image super-resolution," *Appl. Sci.*, vol. 9, no. 22, p. 4874, 2019.
- [4] J. Du, L. Wang, A. Gholipour, Z. He, and Y. Jia, "Accelerated super-resolution MR image reconstruction via a 3D densely connected deep convolutional neural network," in *Proc. IEEE Int. Conf. Bioinf. Biomed. (BIBM)*, Dec. 2018, pp. 349–355.
- [5] J. Du, Z. He, L. Wang, A. Gholipour, Z. Zhou, D. Chen, and Y. Jia, "Super-resolution reconstruction of single anisotropic 3D MR images using residual convolutional neural network," *Neurocomputing*, to be published.
- [6] J. Hatvani, A. Horvath, J. Michetti, A. Basarab, D. Kouame, and M. Gyongy, "Deep learning-based super-resolution applied to dental computed tomography," *IEEE Trans. Radiat. Plasma Med. Sci.*, vol. 3, no. 2, pp. 120–128, Mar. 2019.
- [7] J. Hatvani, A. Basarab, J.-Y. Tournet, M. Gyongy, and D. Kouame, "A tensor factorization method for 3-D super resolution with application to dental CT," *IEEE Trans. Med. Imag.*, vol. 38, no. 6, pp. 1524–1531, Jun. 2019.
- [8] Y. Huang, L. Shao, and A. F. Frangi, "Simultaneous super-resolution and cross-modality synthesis of 3D medical images using weakly-supervised joint convolutional sparse coding," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 5787–5796.
- [9] J. Jurek, M. Kociński, A. Materka, M. Elgalal, and A. Majos, "CNN-based superresolution reconstruction of 3D MR images using thick-slice scans," *Biocybernetics Biomed. Eng.*, vol. 40, no. 1, pp. 111–125, Jan. 2020.
- [10] Y. Li, B. Song, J. Guo, X. Du, and M. Guizani, "Super-resolution of brain MRI images using overcomplete dictionaries and nonlocal similarity," *IEEE Access*, vol. 7, pp. 25897–25907, 2019.
- [11] D. Mahapatra, B. Bozorgtabar, and R. Garnavi, "Image super-resolution using progressive generative adversarial networks for medical image analysis," *Computerized Med. Imag. Graph.*, vol. 71, pp. 30–39, Jan. 2019.
- [12] O. Oktay, W. Bai, M. Lee, R. Guerrero, K. Kamnitsas, J. Caballero, A. de Marvao, S. Cook, D. O'Regan, and D. Rueckert, "Multi-input cardiac image super-resolution using convolutional neural networks," in *Proc. MICCAI*, 2016, pp. 246–254.
- [13] C.-H. Pham, C. Tor-Díez, H. Meunier, N. Bednarek, R. Fablet, N. Passat, and F. Rousseau, "Multiscale brain MRI super-resolution using deep 3D convolutional networks," *Computerized Med. Imag. Graph.*, vol. 77, Oct. 2019, Art. no. 101647.
- [14] J. Shi, Z. Li, S. Ying, C. Wang, Q. Liu, Q. Zhang, and P. Yan, "MR image super-resolution via wide residual networks with fixed skip connection," *IEEE J. Biomed. Health Informat.*, vol. 23, no. 3, pp. 1129–1140, May 2019.
- [15] C. You, W. Cong, M. W. Vannier, P. K. Saha, E. A. Hoffman, G. Wang, G. Li, Y. Zhang, X. Zhang, H. Shan, M. Li, S. Ju, Z. Zhao, and Z. Zhang, "CT super-resolution GAN constrained by the identical, residual, and cycle learning ensemble (GAN-CIRCLE)," *IEEE Trans. Med. Imag.*, vol. 39, no. 1, pp. 188–203, Jan. 2020.
- [16] H. Yu, D. Liu, H. Shi, H. Yu, Z. Wang, X. Wang, B. Cross, M. Bramler, and T. S. Huang, "Computed tomography super-resolution using convolutional neural networks," in *Proc. IEEE Int. Conf. Image Process. (ICIP)*, Sep. 2017, pp. 3944–3948.
- [17] K. Zeng, H. Zheng, C. Cai, Y. Yang, K. Zhang, and Z. Chen, "Simultaneous single- and multi-contrast super-resolution for brain MRI images based on a convolutional neural network," *Comput. Biol. Med.*, vol. 99, pp. 133–141, Aug. 2018.
- [18] X. Zhao, Y. Zhang, T. Zhang, and X. Zou, "Channel splitting network for single MR image super-resolution," *IEEE Trans. Image Process.*, vol. 28, no. 11, pp. 5649–5662, Nov. 2019.
- [19] W. Shi, J. Caballero, F. Huszar, J. Totz, A. P. Aitken, R. Bishop, D. Rueckert, and Z. Wang, "Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 1874–1883.
- [20] J. Long, E. Shelhamer, and T. Darrell, "Fully convolutional networks for semantic segmentation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2015, pp. 3431–3440.
- [21] L. Wang, W. Ouyang, X. Wang, and H. Lu, "Visual tracking with fully convolutional networks," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Dec. 2015, pp. 3119–3127.
- [22] G. Castellano, C. Castiello, C. Mencar, and G. Vessio, "Crowd detection for drone safe landing through fully-convolutional neural networks," in *Proc. SOFSEM*, 2020, pp. 301–312.
- [23] F. Karim, S. Majumdar, H. Darabi, and S. Chen, "LSTM fully convolutional networks for time series classification," *IEEE access*, vol. 6, pp. 1662–1669, 2017.
- [24] M. Bevilacqua, A. Roumy, C. Guillemot, and M.-L.-A. Morel, "Low-complexity single-image super-resolution based on nonnegative neighbor embedding," in *Proc. Brit. Mach. Vis. Conf.*, 2012, p. 135.
- [25] H. Chang, D.-Y. Yeung, and Y. Xiong, "Super-resolution through neighbor embedding," in *Proc. CVPR*, vol. 1, 2004, p. 1.
- [26] D. Dai, R. Timofte, and L. Van Gool, "Jointly optimized regressors for image super-resolution," *Comput. Graph. Forum*, vol. 34, no. 2, pp. 95–104, May 2015.
- [27] J. Yang, J. Wright, T. S. Huang, and Y. Ma, "Image super-resolution via sparse representation," *IEEE Trans. Image Process.*, vol. 19, no. 11, pp. 2861–2873, Nov. 2010.
- [28] C. Dong, C. C. Loy, K. He, and X. Tang, "Image super-resolution using deep convolutional networks," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 38, no. 2, pp. 295–307, Feb. 2016.
- [29] Y. Zhang, K. Li, K. Li, L. Wang, B. Zhong, and Y. Fu, "Image super-resolution using very deep residual channel attention networks," in *Proc. ECCV*, 2018, pp. 294–310.
- [30] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *Proc. Adv. Neural Inf. Process. Syst.*, 2014, pp. 2672–2680.
- [31] J. C. Elliott and S. D. Dover, "X-ray microtomography," *J. Microsc.*, vol. 126, no. 2, pp. 211–213, 1982.
- [32] C. Enders, E.-M. Braig, K. Scherer, J. U. Werner, G. K. Lang, G. E. Lang, F. Pfeiffer, P. Noël, E. Rummeny, and J. Herzen, "Advanced non-destructive ocular visualization methods by improved X-ray imaging techniques," *PLoS ONE*, vol. 12, no. 1, Art. no. e0170633.
- [33] G. Tromba et al., "The SYRMEP Beamline of Elettra: Clinical mammography and bio-medical applications," in *Proc. AIP Conf. Proc.*, 2010, vol. 1266, no. 1, pp. 18–23.
- [34] C. H. McCollough, S. Leng, L. Yu, and J. G. Fletcher, "Dual- and multi-energy CT: Principles, technical approaches, and clinical applications," *Radiology*, vol. 276, no. 3, pp. 637–653, Sep. 2015.
- [35] J. Doerner, M. Hauger, T. Hicketier, J. Byrtus, C. Wybranski, N. G. Hokamp, D. Maintz, and S. Haneder, "Image quality evaluation of dual-layer spectral detector CT of the chest and comparison with conventional CT imaging," *Eur. J. Radiol.*, vol. 93, pp. 52–58, Aug. 2017.



- [36] G. Cybenko, "Approximation by superpositions of a sigmoidal function," *Math. Control, Signals Syst.*, vol. 2, no. 4, pp. 303–314, 1989.
- [37] K. Hornik, "Approximation capabilities of multilayer feedforward networks," *Neural Netw.*, vol. 4, no. 2, pp. 251–257, 1991.
- [38] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, pp. 436–444, May 2015.
- [39] R. T. Ionescu, F. S. Khan, M.-I. Georgescu, and L. Shao, "Object-centric auto-encoders and dummy anomalies for abnormal event detection in video," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 7842–7851.
- [40] A. G. Howard, M. Zhu, B. Chen, D. Kalenichenko, W. Wang, T. Weyand, M. Andreetto, and H. Adam, "MobileNets: Efficient convolutional neural networks for mobile vision applications," 2017, *arXiv:1704.04861*. [Online]. Available: <http://arxiv.org/abs/1704.04861>
- [41] M. Jaderberg, A. Vedaldi, and A. Zisserman, "Speeding up convolutional neural networks with low rank expansions," in *Proc. Brit. Mach. Vis. Conf.*, 2014, pp. 1–12.
- [42] V. Nair and G. E. Hinton, "Rectified linear units improve restricted Boltzmann machines," in *Proc. ICML*, 2010, pp. 807–814.
- [43] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," in *Proc. NIPS*, 2012, pp. 1106–1114.
- [44] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 770–778.
- [45] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *Proc. ICLR*, 2015, pp. 1–15.
- [46] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: From error visibility to structural similarity," *IEEE Trans. Image Process.*, vol. 13, no. 4, pp. 600–612, Apr. 2004.
- [47] Z. Wang and A. C. Bovik, "Mean squared error: Love it or leave it? A new look at signal fidelity measures," *IEEE Signal Process. Mag.*, vol. 26, no. 1, pp. 98–117, Jan. 2009.
- [48] H. R. Sheikh, A. C. Bovik, and G. de Veciana, "An information fidelity criterion for image quality assessment using natural scene statistics," *IEEE Trans. Image Process.*, vol. 14, no. 12, pp. 2117–2128, Dec. 2005.
- [49] C.-Y. Yang, C. Ma, and M.-H. Yang, "Single-image super-resolution: A benchmark," in *Proc. ECCV*, 2014, pp. 372–386.
- [50] T. G. Dietterich, "Approximate statistical tests for comparing supervised classification learning algorithms," *Neural Comput.*, vol. 10, no. 7, pp. 1895–1923, Oct. 1998.
- [51] B. Lim, S. Son, H. Kim, S. Nah, and K. M. Lee, "Enhanced deep residual networks for single image super-resolution," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. Workshops (CVPRW)*, Jul. 2017, pp. 1132–1140.
- [52] C. Popa, N. Verga, M. Patachia, S. Banita, and C. Matei, "Advantages of laser photoacoustic spectroscopy in radiotherapy characterization," *Romanian Rep. Phys.*, vol. 66, pp. 120–126, Jan. 2014.
- [53] D. Sardari and N. Verga, "Calculation of externally applied electric field intensity for disruption of cancer cell proliferation," *Electromagn. Biol. Med.*, vol. 29, nos. 1–12, pp. 26–30, Jan. 2010.



**MARIANA-IULIANA GEORGESCU** received the B.Sc. degree from the Faculty of Mathematics and Computer Science, University of Bucharest, in 2017, and the M.Sc. degree in artificial intelligence from the University of Bucharest, in 2019, where she is currently pursuing the Ph.D. degree with the Faculty of Mathematics and Computer Science. Although she was early in her research career, she has coauthored an article on abnormal event detection in video, accepted at CVPR 2019, and the first author of three papers published at conferences and journals. Her research interests include artificial intelligence, computer vision, machine learning, deep learning, and medical image processing.



**RADU TUDOR IONESCU** (Member, IEEE) graduated from the Faculty of Mathematics and Computer Science, University of Bucharest, Romania, in 2009, and received the M.Sc. Diploma degree in artificial intelligence as a Valedictorian from the University of Bucharest, and the Ph.D. degree from the University of Bucharest, in 2013. He is currently a Full Professor with the University of Bucharest. He received the 2014 Award for Outstanding Doctoral Research in the field of computer science from the Romanian Ad Astra Association. His research interests include machine learning, computer vision, image processing, text mining, and computational biology. He has published over 60 articles at international peer-reviewed conferences and journals, and a research monograph with Springer. He received the Caianiello Best Young Paper Award from ICIAP 2013 for the paper entitled *Kernels for Visual Words Histograms*. He also received the Young Researchers in Science and Engineering Prize organized by Prof. R. Mihalcea for young Romanian researchers in all scientific fields. He participated at several international competitions obtaining top ranks: fourth place in the Facial Expression Recognition Challenge of the WREPL Workshop of ICML 2013, third place in the Native Language Identification Shared Task of the BEA-8 Workshop of NAACL 2013, second place in the Arabic Dialect Identification Shared Task of the VarDial Workshop of COLING 2016, first place in the Arabic Dialect Identification Shared Task of the VarDial Workshop of EACL 2017, first place in the Native Language Identification Shared Task of the BEA-12 Workshop of EMNLP 2017, and first place in the Arabic Dialect Identification Shared Task of the VarDial Workshop of COLING 2018.



**NICOLAE VERGA** is a Senior Doctor in radiotherapy, a Senior Doctor in radiology, and the Head of the Discipline of Oncological Radiotherapy and Medical Imagnostics with the Carol Davila University of Medicine and Pharmacy, Bucharest. During his professional career, he has published several articles and communications, in which he addresses the problem of different types of radiological image processing in order to improve the diagnosis, simulation, and planning in the treatment of different malignant and benign tumors. In daily medical, didactic, and research practice at Coltea Hospital, he evaluates the imagistic investigations of the patients and correlates them with the information provided by the anatomopathological, clinical, and paraclinical examinations of the patients. The average number of cases is about 450 per year. He has participated in the weekly activity of the Commission for Surgical Therapeutic Indications, increasing the number and diversity of cases.

...