

OPEN

Automatic vocal tract landmark localization from midsagittal MRI data

 Mohammad Eslami^{ID*}, Christiane Neuschaefer-Rube^{ID} & Antoine Serrurier^{ID*}

The various speech sounds of a language are obtained by varying the shape and position of the articulators surrounding the vocal tract. Analyzing their variations is crucial for understanding speech production, diagnosing speech disorders and planning therapy. Identifying key anatomical landmarks of these structures on medical images is a pre-requisite for any quantitative analysis and the rising amount of data generated in the field calls for an automatic solution. The challenge lies in the high inter- and intra-speaker variability, the mutual interaction between the articulators and the moderate quality of the images. This study addresses this issue for the first time and tackles it by means of Deep Learning. It proposes a dedicated network architecture named *Flat-net* and its performance are evaluated and compared with eleven state-of-the-art methods from the literature. The dataset contains midsagittal anatomical Magnetic Resonance Images for 9 speakers sustaining 62 articulations with 21 annotated anatomical landmarks per image. Results show that the *Flat-net* approach outperforms the former methods, leading to an overall Root Mean Square Error of 3.6 pixels/0.36 cm obtained in a leave-one-out procedure over the speakers. The implementation codes are also shared publicly on GitHub.

In speech, the sounds of a language are produced by varying the shape and position of the organs surrounding the vocal tract. This region is characterized by a high inter- and intra-speaker variability, both in the space and time domains. Analyzing and modeling the shape of the vocal tract articulators is therefore crucial for speech production research^{1–3} and for diagnosis and therapy of related disorders, including speech disorders^{4,5}, velopharyngeal insufficiency⁶ and swallowing dysfunctions⁷. The vocal tract area extends from the glottis to the lips and is surrounded by various structures such as the larynx, the epiglottis, the velum, the tongue, and the upper and lower lips. Despite the high variability, the organisation of these structures is similar among all speakers. Measuring and analyzing their variability implies therefore to know the deformation of these similar structures across speakers and articulations. A pre-requisite to achieve this is to match similar pertinent anatomical features on these structures across speakers and articulations. The key component for this purpose is to identify on each articulation the same pertinent anatomical landmarks characterizing these structures or demarcating them. It constitutes the general framework of the study.

Articulatory speech production studies often rely on midsagittal images of the vocal tract area and Magnetic Resonance Imaging (MRI) constitutes in this approach an essential modality^{8–10}. Identifying landmarks of the vocal tract area on these images has always been done manually or as a byproduct of manual segmentation^{3,11}. If such a manual approach fulfilled the needs until nowadays, the exploding number of data in this field due to the rise of real-time MRI¹² and the recent progress in data science call undoubtedly for an automatic approach. This study aims at solving this issue and is to our knowledge the first study to address it.

This objective takes place in a larger framework in biomedical engineering and computer vision where localizing anatomical landmarks on biomedical images, sometimes referred to as detecting keypoints, has already been considered in other contexts. As an illustration, this has for instance been considered for the aortic valve on Computer Tomography scans¹³, for cephalometric landmarks on lateral cephalograms^{14,15}, for finger joints on X-ray and MRI or for spine landmarks on volumetric Computer Tomography scans^{16,17}. It has also been considered as a key element of more global registration processes, such as between fundus photographs and MRI for the eye¹⁸ and between series of biological microscopic images¹⁹. Regions of interests have also been identified by means of landmark localization, such as for the brain for Alzheimer's disease diagnosis²⁰ or to detect changes in facial temperature²¹. One can finally also mention the use of landmark identification for research on non-human

Clinic for Phoniatrics, Pedaudiology & Communication Disorders, University Hospital and Medical Faculty, RWTH Aachen University, Aachen, Germany. *email: meslami@ukaachen.de; aserrurier@ukaachen.de

animals²². This short review emphasizes the importance of anatomical landmark localization from images for a large variety of applications and this study lengthens this non-exhaustive list with speech production.

Within the existing literature, two fields of application appear more particularly active and connected to our problem. The first is the localization of landmarks for the face, rich of an abundant research from landmark identification on two-dimensional photographs^{23,24} to landmark identification on three-dimensional Ultrasounds for fetuses for instance^{25,26}. A recent and comprehensive review is provided by Wu *et al.*²⁷. It represents a challenging issue due to the high variability of the shapes, poses, occlusions, and lighting conditions. Similarly, localizing the position of the joints of the body on images to estimate the human pose is also a long-standing problem²⁸. It is also a challenging issue in computer vision due to the high variability of the postures, body shapes, actions, clothes and scenes.

The goal and contribution of this study is to propose a fully-automated end-to-end image analysis methods for localizing key anatomical landmarks in the vocal tract area from midsagittal MRI data. As emphasized earlier, the need for such a method in the field is required and has never been attempted so far. It is aimed to be used in the future for new speakers for which no prior data are available. Considering the existing methods of the literature and the recent rise of data science to solve such problems, Deep Learning (DL) appears as the inescapable approach²⁹. Indeed, DL approaches in image processing tasks appear to outperform most of traditional techniques while bringing more robustness to noise³⁰. In image processing, Deep Neural Networks usually take the particular form of Convolutional Neural Networks (CNNs)³¹. This study intends therefore to provide a CNN architecture able to solve our specific problem. This method and the associated network will be referred to as *Flat-net* in this paper. As detailed in the method section, the design of this method is inspired from existing solutions listed earlier. It takes notably advantage of the concept of *heat-maps*³², which are images where the maximum intensities correspond to the landmark coordinates. It transforms interestingly coordinates in images, leading to input and output data of same nature, and appears powerful to deal with landmarks in image processing^{16,32}.

In addition, in an attempt to compare our solution with existing approaches found in the literature for landmark or joint localization in other contexts, eleven state-of-the-art methods have been adapted to our problem and implemented as our competitors: (1) four methods from the facial landmark localization literature, (2) one method from the human pose estimation literature and (3) six generic methods from the literature on landmark localization in medical images and based on *heat-maps*.

One of the challenges of DL approaches lies in the very large amount of data necessary for training. Standardised datasets are publicly available for classical problems such as facial landmark detection or body pose estimation^{33,34}. No benchmark exists for the problem described in this study and datasets in the field of articulatory speech analyses are usually rather limited and characterized by high shape and noise variability³. The dataset considered for this study is in line with this observation and therefore much more limited and heterogeneous than the datasets mentioned above, increasing all the more the challenge. However, preliminary analyses³⁵ taking advantage of DL algorithms for image segmentation carried out on this exact dataset proved the feasibility of such an approach. The code for the methods considered in this study is publicly available, either by ourselves or by other research groups, for research and validation purposes.

The rest of the paper is organized as follows: in section *Methods*, the dataset, the anatomical landmarks, the methods, and the evaluation schemes are presented; the section *Results* reports the experimental results; finally, the methodology and the results are discussed in the section *Discussion*, together with the perspectives.

Methods

Data. The study considers static midsagittal MRI recorded between 2002 and 2011 from 9 French speakers (5 males, 4 females), referred to as *subjects* in this study, sustaining 62 different articulatory positions, also referred to as *classes* in the context of machine learning, designed to be representative of the French phonemic repertoire^{3,36}. This study does not include any human experiments more than the use of non-invasive MRI data collection mentioned above. All the methods and data acquisition were carried out in accordance with the guidelines and regulations of the local ethic committee called CPP, 'Comité de Protection des Personnes'³⁷ (English translation: Committee for the Protection of People) and the recording protocols were approved by this ethic committee. All subjects were older than 18 and an informed consent was obtained from all of them. The images have been recorded either on a 1.5 or on a 3 Tesla MRI scanner and have a field of view of $256 \times 256 \text{ mm}^2$ and a resolution of 1 mm per pixel. Note that two speakers have been discarded in comparison to Serrurier *et al.*³ and Valdés Vargas³⁶ due to the significantly lower quality of the images, the different fields of view and the different sizes of the images.

21 anatomical landmarks relevant to the study of the speech articulations have been identified. They represent either characteristic landmarks of the speech articulators, such as the tip of the tongue, or the junction between two articulators. They are listed in Table 1 and illustrated for one articulation of one subject in Fig. 1. They have been manually identified on all images of the dataset by an expert. Note that the upper and lower teeth landmarks (UT and LT) denote dental structures and, as such, are not distinguishable from the air on MRI data. They have been determined for each subject by contrast with soft tissues on an articulation acquired on purpose and reported on the other images using their relative position with the hard palate. Please refer to Serrurier *et al.*³ for further information regarding this procedure as well as the data collection and processing.

Challenges. Localizing landmarks of the vocal tract area on midsagittal MRI images presents particular characteristics. Figure 2 shows a few articulations from different subjects illustrating the diversity of the dataset. The main challenges are summarized in the following list.

- The shapes and positions of the vocal tract articulators are characterized by a high variability, due to the variety of the speech task and to the different morphologies and articulatory strategies of the speakers to perform a same task.

Abbreviation	Name	Description
ANS	Anterior Nasal Spine	Anterior nasal spine
EG	Epiglottis- Glottis	Junction between the epiglottis and the glottis
ET	Epiglottis Tip	Tip of the epiglottis
LC	Lip-Chin	Labiomental groove
LLSV	Lower Lip Skin Vermillion	Vermillion border of the lower lip
LLV	Lower Lip Vermillion	Junction between the wet and dry vermillion of the lower lip
LT	Lower Teeth	Upper point of the lower incisors
N	Nose	Most anterior point of the tip of the nose
NM	Neck-Mandible	Junction between the horizontal submandibular line and vertical neck line
NP	Nose-Philtrum	Junction between the philtrum and the external nose
NPX	Nasopharynx	Upper point of the nasopharynx
PL	Pharynx-Larynx	Junction between the pharyngeal wall and the posterior supraglottic region
PNS	Posterior Nasal Spine	Posterior nasal spine
TE	Tongue-Epiglottis	Junction between the tongue and the epiglottis
TJ	Tongue-Jaw	Junction between the tongue and the jaw
TS	Tongue Sub	Most posterior point of the sublingual cavity
TT	Tongue Tip	Tip of the tongue
ULPV	Upper Lip Philtrum Vermillion	Vermillion border of the upper lip
ULV	Upper Lip Vermillion	Junction between the wet and dry vermillion of the upper lip
UT	Upper Teeth	Lower point of the upper incisors
VT	Velum Tip	Tip of the velum

Table 1. List of the landmarks of interest for vocal tract area MRI image analysis.

- Some articulators such as the tongue, velum, lips and epiglottis present a high variability in the vocal tract area, leading to very different locations in the dataset associated to the landmarks VT, TT, TS, LLSV, ULPV, ET and TE.
- Some tissues may touch each others for certain articulations, leading to hardly distinguishable landmarks at these locations, such as TE, TS and ET.
- The larynx area appears very difficult to capture on midsagittal MRI data, leading to a confusing area on the images and hardly identifiable landmarks, such as PL and EG.
- Some articulators may occasionally show very different shapes than in the large majority of the articulations, such as the velum rolled up against the tongue for a few articulations³⁶, leading to unusual location of the landmark VT.
- Two important landmarks for speech production analyses are the teeth landmarks UT and LT, which are not directly visible on the images, as mentioned earlier.
- The images are recorded at different times with different scanners, leading to variable quality and noise levels. Similarly, the quality and noise level may not be homogeneous within a single image.
- Despite the high variability, the size of the available dataset, *i.e.* 9 subjects with 62 articulations, hence 558 articulations, is rather limited in comparison to those usually used for landmark localization in the literature, *e.g.* over 10 K images for mentioned refs.^{23,34,38}.

Methods. In order to compare the method proposed in this paper with the current state-of-the-art in the field, eleven methods taken from the literature have been adapted to our problem and implemented: (1) five methods from the literature related to facial landmark detection and body pose estimation and (2) six general methods for landmark localization and based on the concept of *heat-maps*. Our method, referred to as *Flat-net*, is inspired by these methods and presented at the end.

Methods from facial landmark detection and body pose estimation. The five methods considered in this section are the *dlib*, *HyperFace*, *Deep Alignment Network*, *Shape Fitting by Deep-Regression* and *Multi-Context Attention Model* methods and are presented as follows.

dlib. The algorithm available as part of the *dlib* library is an implementation of the ensemble of regression trees presented in 2014 by Kazemi and Sullivan³⁹. This technique takes advantage of simple features with fast computing capacities, *e.g.* the pixels' intensity differences, to directly estimate the landmark locations. These locations are subsequently refined with an iterative process made of a cascade of regressors and using gradient boosting. Note that the *dlib* method is the only method considered in this study not based on DL.

HyperFace. The *HyperFace* method makes use of an end-to-end DL network for simultaneous face detection, landmark localization, pose estimation and gender recognition²³. It exploits the intermediate layers of a deep CNN, such as the *ResNet-101*⁴⁰, by connecting together the intermediate feature maps to further predict the various desired outputs. The last layers of the *HyperFace* network are fully connected layers.

Deep alignment network. The *Deep Alignment Network (DAN)* is a method based on a DL to localize facial landmarks²⁴. It consists of multiple stages of CNNs, where each stage improves the locations of the facial landmarks estimated by the previous stage. A key element of the system is the use of *heat-maps* within each stage.




Figure 1. MRI superimposed with the 21 anatomy landmarks of the study.



Figure 2. Five images of the database illustrating the variability.

In their approach, a *heat-map* is defined as an image with highest intensity values at the exact locations of all the considered landmarks and decreasing intensities around as a function of the distance to the nearest landmark. The last two layers of each stage are fully connected layers.

Shape fitting by deep-regression. The *Shape Fitting Deep-Regression (SFD)* method has been proposed recently. It combines CNNs with model-based fitting algorithms of the landmark positions such as obtained by Principal Component Analysis (PCA)⁴¹. The PCA is included in the network using a dedicated layer type. This network is characterized by a fast computing performance and can process until several hundreds of frames per second.

Multi-context attention model. The *Multi-Context Attention Model (MCAM)* method is an extended version of the DL *stacked hourglass networks*⁴² designed for human pose estimation and body joint localization⁴³. It generates *heat-maps* describing the body joint locations by using multiple resolutions, conditional random fields and an original layer type combining various convolutional layers together. For training, the ground truth *heat-maps* are generated by 2-D Gaussians centered on the joint locations. The generated *heat-maps* contain all joint locations together and are further split into partial *heat-maps* for each body joint by means of an extra spatial classifier. Since the network is designed to localize 16 body joints, two of these networks are necessary in practice in the current study to localize the 21 landmarks.

Heat-map-based methods. Previous methods aim at detecting landmark coordinates from images. One limitation of these approaches come from the different nature of the input and output data. An alternative to overcome this is to consider the landmark coordinates as images via *heat-maps* in output. One of the pioneering attempt to implement such approach can be attributed to Pfister *et al.*³² for landmark coordinate detection for body pose estimation.

In this approach, a single landmark is described as a full image, the *heat-map*, with a maximal intensity on the landmark location. Several landmarks can be considered on the same *heat-map* or can be described as several channels of an image (or tensor), leading to the *heat-maps in channels* (one *heat-map* per landmark).

Localizing the landmarks on an image consists therefore in generating the associated *heat-maps*, instead of the vectors of the landmarks' coordinates as for the previous studies. The landmark locations are then straightforwardly derived from the *heat-maps* as the points with maximal intensity in each channel. This kind of approach




Figure 3. Superposition of 3 *heat-maps* for 3 different landmarks, represented as the 3 channels of a single RGB image (right), and superimposed on the corresponding input MRI image (left).

appears particularly appropriate for medical image processing^{16,17}. Note that this concept of *heat-map* is related to the concept of *heat-maps* mentioned in the DAN and MCAM methods. However, in these methods, the *heat-maps* contain all landmark/joint locations together and are not directly provided as output of the networks.

Practically in using *heat-maps in channels*, for an input image of size $N \times N$, the network generates L output *heat-map* images (i.e. tensor $H \in \mathbb{R}^{N \times N \times L}$) where the l^{th} *heat-map* ($H_l \in \mathbb{R}^{N \times N \times 1}$) has the maximum intensity value at the location of the estimated l^{th} landmark. The normalized target *heat-maps* are obtained via a Gaussian hat with a σ ($=10$ in our experiment) pixels and a maximum of 1 around each landmark location. Figure 3 shows an example of the combination of 3 *heat-maps* created for 3 landmarks and displayed as the 3 channels of a single RGB image. The predicted location of the l^{th} landmark corresponds to the location of the maximum value in the l^{th} predicted *heat-map*:

$$(x_p, y_p)_l = \operatorname{argmax}_{x,y} H_l(x, y), \quad l = 1, \dots, L \quad (1)$$

Six methods based on the generation of *heat-maps* are considered in this section. The first five methods make use of the implementation proposed by Payer *et al.*¹⁶ and for which the code is publicly available. Payer *et al.* investigate several DL approaches for landmark localization on medical images based on the generation of *heat-maps*. They implement existing networks such as *Downsampling-net*, *ConvOnly-net*, and *U-net*. They also propose a new architecture called *SpatialConfiguration-net* and its extension with embedded *U-net*^{16,17}. These five methods show promising results for localizing diverse landmarks on various biomedical image modalities as mentioned in the introduction^{16,17} (finger landmarks on X-Ray and MRI, cranium landmarks on lateral cephalogram X-Ray and spine landmarks on Computer Tomography scans). These methods are briefly described below, please refer to the dedicated references^{16,17} for further details. The last method presented in this section, the *pix2pix's generator network* covers our own implementation of the concept of *heat-maps* into the existing *pix2pix* network.

Downsampling-net. The *Downsampling-net* (*DS-net*) methods alternates convolution and pooling layers in the network. The downsampling approach ensures covering large image areas with small kernel sizes but poor accuracy in localization has to be expected due to the associated low resolution of the resulting *heat-maps*¹⁶. This network with its default hyper-parameters did not converge in our experiment.

ConvOnly-net. The *ConvOnly-net* (*conv-net*) method aims at overcoming the low resolution of the generated *heat-maps* mentioned above by discarding all pooling and strided convolution layers¹⁶. This network results in six convolution layers.

U-net. The *U-net*⁴⁴ is a well-known architecture for medical image processing, more specifically for segmentation. The *U-net* includes convolutional layers with skip connections and pooling layers in an auto-encoder architecture, making efficient the processing of large images even with small kernel size. The implementation proposed by Payer *et al.*¹⁶ and used in this study is a slightly modified version, where for instance the maximum pooling has been replaced by an average pooling.

SpatialConfiguration-net. The *SpatialConfiguration-net* (*SCN*) method combines interestingly the local appearance of the landmarks with their spatial locations in reference to all other landmarks. It consists of three blocks. The first block is made of three convolutional layers with small kernel sizes, resulting in local appearance *heat-maps*. The second and third blocks refine the outputs of the first block by considering the spatial configurations and by combining them with the initial outputs to discard the unrealistic results¹⁶.

SpatialConfiguration-net with embedded U-net. The *SpatialConfiguration-net with embedded U-net* (*SCN(U-net)*) method is an extension of the *SCN* method which embeds a *U-net* network into the local appearance extraction block^{16,17}.

pix2pix's generator network. This architecture exploits the generator component of the *pix2pix* network⁴⁵, referred to in the following as *p2p-GN*, standing for *pix2pix's generator network*. It is based on a hourglass-shaped CNN with skip connections. This network is specifically designed to analyze and generate images, hence particularly adapted in our case for the generation of *heat-maps* from MRI. It has already proved to be very efficient for




Figure 4. Architecture of one of the *Flat-net* networks producing 5 *heat-maps* corresponding to 5 landmarks.

such tasks⁴⁶. In the current study, this network is adapted to generate particular types of images, the *heat-maps*. The loss function as well as the hyper-parameters are those reported by⁴⁵. Experimental analyses showed that this network produces better results when it does not generate more than 3 *heat-maps* at the same time. For this reason the prediction of the 21 *heat-maps* are split into 7 different networks. The additional number of weights resulting from the use of 7 networks in this approach is discussed in section *Results*.

Proposed network architecture. This section details our own solution to solve the problem. It implements the concept of *heat-maps-in-channels* described in the previous section and proposes a dedicated network for the generation of such *heat-maps*.

Flat-net. The proposed architecture *Flat-net* is presented in Fig. 4. As indicated by its name, it does not contain any pooling, down and up sampling nor fully connected layers but considers various kernel sizes and dilation rates. It is designed to explore in the first layers the image at different resolutions, so as to deal with the different morphologies, and combine the resulting feature maps in a second step to output the desired *heat-maps*. In layers L1 and L2, convolutional filters of kernel size 9×9 are applied consecutively using 5 different dilation rates. The generated feature maps are then concatenated in layer L3. It is followed by a convolutional layer of kernel size 5×5 and three consecutive convolutional layers of kernel size 1×1 . The activation functions for all convolutional layers are *Relu*, except the last layer (L7) using *tanh*. The number of filters used for each layer is indicated in the Fig. 4. The loss is measured as the mean absolute error between the predicted and the desired *heat-maps*. Since there is no pooling in this architecture, the concatenation of the layer L3 leads to a very large size tensor, causing practical memory issues. To solve this problem, similarly to the *p2p-GN* approach, the prediction of the 21 *heat-maps* are split in practice into 5 different networks. Again, the additional number of weights resulting from the use of 5 networks in this approach is discussed in section *Results*.

Implementation and evaluation. The 9×62 input grayscale images are converted into grayscale RGB by simple channel repetition to comply with the input format of the networks. Considering the relatively limited number of data, the dataset is augmented⁴⁷ via 10 different arbitrary methods as follows: (1) Addition of a Gaussian intensity noise of mean equal to 0 and variance to 12.75, (2) Blurring with a Gaussian of variance $\sigma = 5$ pixels, (3–4) Rotation of $+10$ and -5 degrees, (5–6) Translation of $(+30, +10)$ pixels and $(+40, -10)$ pixels, (7) Rotation of -5 degrees followed translation of $(+30, +10)$ pixels, (8) Zooming out of scale 0.8, (9) Translation of $(+30, +10)$ pixels followed by zooming in of scale 1.2, and (10) Translation of $(+40, +20)$ pixels followed by zooming out of scale 0.9 plus blurring with a Gaussian of variance $\sigma = 3$ pixels. By this method, the dataset is artificially augmented from $9 \times 62 = 558$ to $11 \times 9 \times 62 = 6138$ images.

The errors of prediction of the landmark coordinates are evaluated by means of Euclidean distance and Root Mean Squared Error (RMSE). The Euclidean distances are expressed in pixels to comply with the existing results in the literature of the domain while the RMSE is expressed in centimeters to provide comprehensive and interpretable results for speech analyses purposes. In addition, the percentage of samples (*i.e.* test images) with landmarks presenting distance errors higher than 5 pixels (*i.e.* outliers), are reported. Moreover, the statistical significance of the difference between the means of the distance errors obtained for our methods on the one hand and all the other methods on the other hand is evaluated by means of *paired t-test*.

For the l^{th} landmark, if $(x_g, y_g)_l$ and $(x_p, y_p)_l$ denote respectively the ground-truth and the predicted coordinates, the Euclidean distance d_l is calculated as follows:

$$d_l = \sqrt{(x_g - x_p)_l^2 + (y_g - y_p)_l^2} \quad (2)$$

Using similar notations, the RMSE is calculated as follows:

$$RMSE = \sqrt{\frac{1}{Q} \sum_{q=1}^Q (x_g - x_p)_q^2 + (y_g - y_p)_q^2} \quad (3)$$

where Q is the number of considered elements, *e.g.* all the landmarks of all images, and q is the corresponding index.

Most of the landmark localization methods are evaluated in the literature by a hold-out scheme, *i.e.* by splitting the data into train/test sets, for instance 21,997/1,000 for the *HyperFace* method²³, 2,000/330 for the *dlib* method³⁹ or 3148/600 for the *DAN* method²⁴. In the current application, the ultimate objective is however to localize landmarks on new speakers, *i.e.* on speakers where no data were available before. To comply both with the literature benchmarks and the specificity of the study, the performances are evaluated via two schemes: (1) the randomized 10-fold cross-validation (CV) and (2) the leave-one-subject-out cross-validation (LoSo). Note in addition that 5% of the training data in each training session are randomly set aside in advance for validation purposes and tracking the learning curves. The learning would stop when the loss curve of the validation set reaches plateau.

In the CV scheme, the augmented dataset is randomly split into 10 groups, 9 being used for training purposes, *i.e.* 5,525 samples (after rounding), and 1 for test purposes, *i.e.* 613 samples (after rounding). The training and evaluation is then repeated 10 times until each single group has been used as the test set. Each sample is therefore being used as a test sample at some point during the process. Note that in this evaluation scheme, the train and test sets may share data of same subjects and/or of same articulations, making the two sets not completely independent. However, in accordance with the literature in the domain, a very large dataset with many more subjects and articulations could be perfectly evaluated through this scheme and is therefore considered in this study. Above all, this scheme is considered in our evaluation to assess the validity of the hyper-parameters reported by the methods for our environment.

In the LoSo scheme, the 62 images of one arbitrary subject are set aside to serve as the test set. The remaining images are then augmented, leading to $8 \times 62 \times 11 = 5,456$ images, and used for the training. In other words, the network is not trained with data from the test subject. The training and evaluation is then repeated 9 times until each subject has been used as the test set. Each sample is therefore being used as a test sample at some point during the process on a model trained on the other subjects. This scheme is much stricter and challenging than the CV scheme as the trained network does not contain any information regarding the tested subject. Note however that the train and test sets may still share data of same articulations (but not speakers). This point will be revisited in the discussion.

The results for the two evaluation schemes are presented in the section *Results*. All of the hyper-parameters of the methods taken from the literature are set to their default values mentioned in their corresponding studies. The training machine was made of an Intel Xeon w-2145 (3.70 GHz) CPU and a NVIDIA Tesla P100-SXM2-16GB GPU. Except for *dlib*, all the methods are trained on GPU. All the implementation are available online on GitHub (<https://github.com/mohaEs> and <https://github.com/christianpayer/MedicalDataAugmentationTool>).

Results

As noted in the section *Methods*, one of the 12 methods considered in this study, the *DS-net*, did not converge. For this reason, unless explicitly mentioned, the results cover only the remaining 11 methods. An overall comparison of the performances of the eleven methods are provided in Fig. 5 for both the CV and the LoSo schemes. It displays in box plots the Euclidean distances between the predicted and true landmark locations. For each box, the central mark indicates the median while the bottom and top edges indicate respectively the 25th and 75th percentiles. The whiskers extend to the most extreme data points not considered as outliers, the outliers being plotted individually using the '+' symbol.

Regarding the CV scheme, all the methods show good accuracy, with boxes below 2.5 pixels, except the *MCAM* method. This might be ascribed to the design of the method, optimized for joint localization and not anatomical landmarks as in the current case. Nevertheless, the method still shows decent results, with 75% of the distance errors being smaller than 5 pixels. On the other side, the *dlib* method presents the best results, possibly due to its boosting approach.

By attempting to predict landmarks on a subject not used to train the models, the LoSo evaluation scheme is more constraining and presents logically deteriorated – but more pertinent – results in comparison to the CV evaluation scheme. The results for the *HyperFace*, *DAN*, *SFD* and *MCAM* methods appear in particular significantly deteriorated. On the contrary, the deterioration appears more limited for the other methods and still lead to fairly good accuracy, with boxes remaining below 3.5 pixels.

The results in terms of RMSE and per landmark are provided in Fig. 6. It confirms the lower accuracy already noted for the LoSo scheme in comparison to the CV scheme. It also shows that the four methods *HyperFace*, *DAN*, *SFD* and *MCAM* estimate in the LoSo scheme many landmarks for more than 50% of the images with an error larger than 0.5 cm (5 pixels), rather problematic for speech production studies. The *conv-net*, *U-net*, *SCN* and *SCN(U-net)* methods display decent results in general but present serious errors on the PL and ET landmarks. On the contrary, the three methods *dlib*, *p2p-GN* and *Flat-net* still show acceptable results, with almost all landmarks for more than 70% of the images having an error lower than 0.5 cm. For these methods, the landmark EG is the most challenging one. It means that our method, the *Flat-net*, performs at least as good as the two best methods taken from the literature, namely the *dlib* and *p2p-GN* methods. They provide fairly good results in the




Figure 5. Box plots of the overall Euclidean distances between the predicted and true landmark locations for the 11 methods considered in the results, for the CV (top) and LoSo (bottom) evaluation schemes. Note the y-axis scale, linear in two steps to zoom in the 0–10 pixels range.

LoSo scheme despite the limited dataset. The other methods do not appear suitable to handle satisfactorily the data and problem presented in this study. A larger dataset with significantly more subjects and articulations might solve this issue.

A summary of the results and characteristics of the twelve methods of the study is reported in Table 2. Note that the storage space is reported for information but does not play in our eyes a critical role. Moreover, it depends on the data format and of the extent of meta-data stored. Also, the training time depends on the learning rate, the type of optimizer and the deep learning framework (Pytorch vs. Tensorflow, etc.). In overall, the best results in LoSo scheme are reported for our method, the *Flat-net* method, with a RMSE of 0.36 cm, slightly better than for the *dlib* (0.39) and *p2p-GN* (0.41) methods. Among all methods, the DL-based methods require significantly more training time and storage space. The *p2p-GN* method in particular has the largest number of weights to train and requires the largest space for storage. The *dlib* method is on the contrary rapidly trained, even on CPU, and requires the smallest space for storage.

The results of the t-tests conducted between the distance errors obtained by our method, the *Flat-net* method, also presenting interestingly the best performance, and the errors obtained by the other methods are presented in Table 3. It confirms a statistical significance between the means of the distance errors, all of them obtained with $p = 0.000$ except $p < 0.04$ for the *SCN* method.

The overall RMSEs per subject in LoSo scheme for the eleven methods considered in the results are displayed in Fig. 7. The *p2p-GN* and *Flat-net* methods tend to show more homogeneous and lower errors across subjects, suggesting that they are the two more accurate and robust methods to predict the landmarks for new and unseen subjects. Finally, some examples of practical results for the three best and the two worst methods are illustrated in Fig. 8.

Discussion

The present paper described an original method to localize anatomical landmarks of the vocal tract area on MRI images and compared its performance to 11 state-of-the-art methods of the literature. A dataset of mid-sagittal MRI from 9 speakers sustaining 62 articulations and annotated with the location of 21 landmarks has been considered. The methods have been evaluated through two schemes, a randomized 10-fold scheme and a leave-one-speaker-out scheme, considered as more challenging. Experimental results show the ability of all methods to cope with the problem in the 10-fold scheme but divergence of performance appear in the more challenging leave-one-speaker-out scheme. In general, our method, the *Flat-net* method, outperforms the other methods to solve the specific problem of the study and is only approached by two other methods, namely *dlib* and *p2p-GN*. Interestingly, these two methods include the only method not based on DL (*dlib*) as well as the method that has been adapted the most to fit at best our problem (*p2p-GN*). Note also that the *heat-map*-based methods present in general significantly better results than the other methods, supporting this approach to tackle the problem of landmark localization on medical images. This approach leads to networks without fully connected layers, usually




Figure 6. RMSE and percentage of samples with errors higher than 5 mm for the 11 methods considered in the results for the CV scheme (top) and the LoSo scheme (bottom). (Left) RMSE matrix for each landmark. (Middle) Matrix of the percentage of samples with errors > 5 mm for each landmark. (Right) Overall RMSEs and percentages.

	RMSE in CV (cm)	RMSE in LoSo (cm)	Distance error in CV (pixels)	Distance error in LoSo (pixels)	Number of weights	Number of epochs	Training time (min)	Storage space (MB)
dlib	0.08	0.41	0.59 ± 0.59	2.94 ± 2.74	—	—	55 (cpu)	45
HyperFace	0.18	0.96	1.49 ± 0.99	8.02 ± 4.91	35,168,006	200	50	402
DAN	0.19	0.69	1.66 ± 0.91	5.81 ± 3.53	$2 \times 23,104,092$	2×120	2×130	2×280
SFD	0.36	0.64	2.85 ± 2.21	5.23 ± 3.51	6,248,029	150	—	25
MCAM	0.45	0.96	3.96 ± 1.96	8.47 ± 4.18	14,500,480	60	150	210
DS-net	—	—	—	—	2,056,741	—	—	16
conv-net	0.18	0.65	1.39 ± 1.14	3.34 ± 5.47	9,933,349	40	117	78
U-net	0.17	0.91	1.29 ± 1.17	3.07 ± 8.38	2,662,181	80	198	21
SCN	0.20	0.54	1.52 ± 1.38	2.66 ± 4.61	1,240,958	100	225	10
SCN(U-net)	0.19	0.73	1.51 ± 1.21	2.92 ± 6.55	3,718,602	60	175	30
p2p-GN	0.12	0.39	0.89 ± 0.85	2.93 ± 2.51	$7 \times 54,420,483$	7×60	7×55	7×440
Flat-net	0.16	0.36	1.12 ± 1.13	2.39 ± 2.47	$5 \times 2,958,533$	5×30	5×88	5×34

Table 2. Summary of the results and characteristics of the twelve methods of the study. The lowest occurrence for the four left columns is emphasized in bold. The sign ‘—’ means that the information is not relevant or not available. Remember that the *DS-net* method did not converge in our study.

	dlib	HyperFace	DAN	SFD	MCAM	Conv-net	U-net	SCN	SCN (U-net)	p2p-GN
CV	0	2.2e-150	0	0	0	2.1e-72	1.7e-29	3.6e-125	5.3e-137	5.5e-70
LoSo	9.1e-28	0	0	0	0	5.3e-44	3.9e-10	0.037	4.6e-08	1.7e-28

Table 3. p-values of paired t-tests for distance errors between the *Flat-net* and the remaining 10 methods considered in the results for the CV and LoSo schemes.

used to transform the feature maps into vectors of landmark locations in output. This may result in an architecture possibly less prone to error propagation, especially for such a limited dataset. Indeed, although *heat-maps* are also somehow part of the *DAN* and *MCAM* methods, the design in channels together with the use of adapted networks able to output directly these *heat-maps* proved successful for all the other methods.

The good performance of the *dlib* method might be ascribed to the use of the boosting approach combined with the analysis of the input image in small regions by applying windows, possibly reducing the sensitivity to the variability of other regions of the image. In overall, our *Flat-net* method, the only method entirely developed in this study, tends to be more robust and to present better performances. One could object that the better performance of the *Flat-net* network may simply come from its hyper-parameters optimized for the current problem. The results obtained in the CV evaluation scheme suggest however that all eleven methods considered in the results – except the *MCAM* method to a certain extent – can perform the task with success, discarding this objection. These results support the approach using *heat-maps* in channels and networks without fully connects layers to localize landmarks of the vocal tract area on MRI data.




Figure 7. Bar plots of the overall RMSEs per tested subject in LoSo scheme for the 11 methods considered in the results.

Method.	Test image #1 in CV	Test image #1 in LoSo	Test image #2 in LoSo	Test image #3 in LoSo
<i>dlib</i>				
<i>HyperFace</i>				
<i>MCAM</i>				
<i>p2p-GN</i>				
<i>Flat-net</i>				

Figure 8. MRI images of three different subjects and articulations zoomed in the region of the vocal tract and superimposed with the 21 ground truth (green crosses) and predicted (red crosses) landmarks for the best three methods (*dlib*, *p2p-GN* and *Flat-net*) and the two worst methods (*HyperFace* and *MCAM*) reported for the LoSo scheme. The first test image is also presented with results of the CV scheme.

Regarding the localization of the landmarks, the overall RMSEs in LoSo scheme for the three best methods identified above and for a subset of 13 out of 21 landmarks are displayed in Fig. 9: the 11 landmarks mentioned in the section *Methods* plus the landmarks N on the one hand and NM and EG on the other hand, presenting respectively the best and worst results. Note that these results are a zoom of the results presented in Fig. 6a. All these landmarks are more accurately located by the *Flat-net* method, occasionally at the same precision than other methods. Except for NM, PL and EG, the RMSEs stay below 0.35 cm, a fairly good result considering the challenge associated with these landmarks and comparable to the overall RMSE achieved by the method. The landmarks UT and LT in particular, not visible on MRI and giving many problems in articulatory speech studies⁴⁸, are estimated with a respective accuracy of 0.25 cm and 0.3 cm for new speakers without additional *a priori* information. Similarly,




Figure 9. Bar plots of the overall RMSEs in LoSo scheme for the 3 methods *dlib*, *p2p-GN* and *Flat-net* for a subset of 13 out of 21 landmarks.

the landmark TS, important for speech articulations and challenging to identify⁴⁹, is estimated with a precision of 0.25 cm. This emphasizes the robustness of the chosen approach through deep learning and the potential of these results for speech studies. The worst results are achieved for the landmarks NM, outside the range of the vocal tract area, and EG and PL, in the larynx area. Although the larynx plays an important role in speech production, it remains at the margin of the vocal tract area and does not play a central role in articulatory speech studies³.

It should also be noted that some landmarks do not exhibit very salient characteristics, such as the junctions between the wet and dry vermillion of the lips (ULV and LLV) or the sublingual cavity posterior point (TS) when the tip of the tongue is in a low position. Similarly, some regions tend to be hardly interpretable in terms of anatomy, such the anterior part of the larynx associated with the landmark EG. Annotating manually their exact location on the images was a challenging task at the first point, questioning the quality of the ground-truth data and possibly explaining the lower accuracy achieved for EG. Furthermore, a detailed analysis of the results such as presented in Fig. 8 reveals that the location of some landmarks may appear occasionally more accurate in output of the presented methods than in the original so-called ground-truth. This is a well known effect of machine learning methods, and DL methods in particular, which tend to avoid encoding noise and outliers by means of regularization techniques^{50,51}. In general, a larger dataset labelled by several experts may limit the impact of the uncertainties in the ground-truth data and reinforce the robustness the DL methods and their resistance to noise.

The methods have been evaluated by means of two schemes, the CV and LoSo schemes. Strictly speaking, the most rigorous scheme would have been to leave both subject and articulation out, *i.e.* leave-one-subject-one-class-out (LoSoCo), to ensure that the network does not contain any information regarding the new tested image. In this scheme, all the articulations of one subject and one specific articulation for all subjects would be discarded in the training and the same specific articulation of the left subject would be tested. This would lead to $62 \times 9 = 558$ training sessions for each of the twelve methods. According to the times reported in Table 2, it would take more than a year, making this evaluation unrealistic in practice. However, the challenge of the problem lies rather in the estimation of the landmark locations for a new speaker than for a new articulation. Indeed, the corpus of 62 articulations can be considered large enough and representative of the French phonemic repertoire so that one articulation could fairly well be estimated from the 61 others³. For this reason, the LoSo scheme appears as a valid approximation to evaluate the methods on our problem.

In summary, the method proposed in this study (*Flat-net*) outperforms the state-of-the-art methods. It supports the description of landmarks locations in terms of *heat-maps in channels* and the generation of these *heat-maps* by means of DL networks without fully connected layers for such a variable and limited dataset. Future works may include the combination of successful features from the *dlib* method, showing very promising results, with DL approaches to create more robust methods, such as for instance the methods using deep forest networks^{52,53}. In addition, since the accuracy of the *heat-map*-based methods for all landmarks are not same, it seems reasonable to combine them to form so-called machine learning ensemble methods. For example, the *p2p-GN* and *Flat-net* methods present good accuracy for the landmarks ET, PL, TE and TS but not for the landmarks EG and NPX while on the contrary the *SCN* method presents good accuracy for the landmarks EG and NPX but not for the landmarks ET, PL, TE and TS. Combining these methods to take advantage of their relative assets may therefore lead to higher accuracy in the localization of vocal tract area landmarks from MRI data together with more robustness. Furthermore, considering the recent rise of real-time MRI for speech production studies¹², the next steps will be to adapt this technique to real-time MRI data.

Data availability

The source codes of the methods are available at <https://github.com/mohaEs>. The trained models are available from the author A.S. on reasonable request. Authors have not a permission to share the data.

Received: 3 October 2019; Accepted: 9 January 2020;

Published online: 30 January 2020

References

1. Harshman, R., Ladefoged, P. & Goldstein, L. Factor analysis of tongue shapes. *The J. Acoust. Soc. Am.* **62**, 693–707 (1977).
2. Beautemps, D., Badin, P. & Bailly, G. Linear degrees of freedom in speech production: Analysis of cineradio- and labio-film data and articulatory-acoustic modeling. *The J. Acoust. Soc. Am.* **109**, 2165–2180 (2001).

3. Serrurier, A., Badin, P., Lamalle, L. & Neuschaefer-Rube, C. Characterization of inter-speaker articulatory variability: a two-level multi-speaker modelling approach based on MRI data. *The J. Acoust. Soc. Am.* **145**, 2149–2170, <https://doi.org/10.1121/1.5096631> (2019).
4. Yamasaki, R. *et al.* Vocal tract adjustments of dysphonic and non-dysphonic women pre-and post-flexible resonance tube in water exercise: a quantitative mri study. *J. Voice* **31**, 442–454 (2017).
5. Guzman, M. *et al.* Computerized tomography measures during and after artificial lengthening of the vocal tract in subjects with voice disorders. *J. voice* **31**, 124–e1 (2017).
6. Freitas, A. C., Wylezinska, M., Birch, M. J., Petersen, S. E. & Miquel, M. E. Comparison of cartesian and non-cartesian real-time mri sequences at 1.5 t to assess velar motion and velopharyngeal closure during speech. *PLoS one* **11**, e0153322 (2016).
7. De Alarcón, A., Prager, J., Rutter, M. & Wootten, C. T. Laryngeal and tracheal airway disorders. In *Kendig's Disorders of the Respiratory Tract in Children*, 1118–1124 (Elsevier, 2019).
8. Labrunie, M. *et al.* Automatic segmentation of speech articulators from real-time midsagittal mri based on supervised learning. *Speech Commun.* **99**, 27–46 (2018).
9. Narayanan, S., Nayak, K., Lee, S., Sethy, A. & Byrd, D. An approach to real-time magnetic resonance imaging for speech production. *The J. Acoust. Soc. Am.* **115**, 1771–1776, <https://doi.org/10.1121/1.1652588> (2004).
10. Story, B. H. Synergistic modes of vocal tract articulation for american english vowels. *The J. Acoust. Soc. Am.* **118**, 3834–3859, <https://doi.org/10.1121/1.2118367> (2005).
11. Echternach, M., Burk, F., Burdumy, M., Traser, L. & Richter, B. Morphometric differences of vocal tract articulators in different loudness conditions in singing. *PLOS ONE* **11**, 1–17, <https://doi.org/10.1371/journal.pone.0153792> (2016).
12. Ramanarayanan, V. *et al.* Analysis of speech production real-time mri. *Comput. Speech & Lang.* **52**, 1–22 (2018).
13. Zheng, Y. *et al.* Automatic aorta segmentation and valve landmark detection in c-arm ct for transcatheter aortic valve implantation. *IEEE transactions on medical imaging* **31**, 2307–2321 (2012).
14. Lindner, C. *et al.* Fully automatic system for accurate localisation and analysis of cephalometric landmarks in lateral cephalograms. *Sci. reports* **6**, 33581 (2016).
15. Vandaele, R. *et al.* Landmark detection in 2d bioimages for geometric morphometrics: a multi-resolution tree-based approach. *Sci. reports* **8**, 538 (2018).
16. Payer, C., Štern, D., Bischof, H. & Urschler, M. Regressing heatmaps for multiple landmark localization using cnns. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 230–238 (Springer, 2016).
17. Payer, C., Štern, D., Bischof, H. & Urschler, M. Integrating spatial configuration into heatmap regression based cnns for landmark localization. *Med. Image Analysis* **54**, 207–219 (2019).
18. De Zanet, S. I. *et al.* Landmark detection for fusion of fundus and mri toward a patient-specific multimodal eye model. *IEEE transactions on biomedical engineering* **62**, 532–540 (2014).
19. Wang, C.-W., Ka, S.-M. & Chen, A. Robust image registration of biological microscopic images. *Sci. reports* **4**, 6050 (2014).
20. Liu, M., Zhang, J., Adeli, E. & Shen, D. Joint classification and regression via deep multi-task multi-channel learning for alzheimer's disease diagnosis. *IEEE Transactions on Biomed. Eng.* **66**, 1195–1206 (2018).
21. Sonkusare, S. *et al.* Detecting changes in facial temperature induced by a sudden auditory stimulus based on deep learning-assisted face tracking. *Sci. reports* **9**, 4729 (2019).
22. Finka, L. R. *et al.* Geometric morphometrics for the study of facial expressions in non-human animals, using the domestic cat as an exemplar. *Sci. reports* **9**, 9883 (2019).
23. Ranjan, R., Patel, V. M. & Chellappa, R. Hyperface: A deep multi-task learning framework for face detection, landmark localization, pose estimation, and gender recognition. *IEEE Transactions on Pattern Analysis Mach. Intell.* **41**, 121–135 (2019).
24. Kowalski, M., Naruniec, J. & Trzcinski, T. Deep alignment network: A convolutional neural network for robust face alignment. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 88–97 (2017).
25. Vezzetti, E., Speranza, D., Marcolin, F., Fracastoro, G. & Buscicchio, G. Exploiting 3d ultrasound for fetal diagnostic purpose through facial landmarking. *Image Analysis & Stereol.* **33**, 167–188 (2014).
26. Vezzetti, E., Speranza, D., Marcolin, F. & Fracastoro, G. Diagnosing cleft lip pathology in 3d ultrasound: a landmarking-based approach. *Image Analysis & Stereol.* **35**, 53–65 (2016).
27. Wu, Y. & Ji, Q. Facial landmark detection: A literature survey. *Int. J. Comp. Vis.* **127**, 115–142 (2019).
28. Gong, W. *et al.* Human pose estimation from monocular images: A comprehensive survey. *Sensors* **16**, 1966 (2016).
29. Pouyanfar, S. *et al.* A survey on deep learning: Algorithms, techniques, and applications. *ACM Comput. Surv. (CSUR)* **51**, 92 (2019).
30. Voulodimos, A., Doulamis, N., Doulamis, A. & Protopapadakis, E. Deep learning for computer vision: A brief review. *Comput. intelligence and neuroscience* **2018** (2018).
31. Gu, J. *et al.* Recent advances in convolutional neural networks. *Pattern Recognit.* **77**, 354–377 (2018).
32. Pfister, T., Charles, J. & Zisserman, A. Flowing convnets for human pose estimation in videos. In *Proceedings of the IEEE International Conference on Computer Vision*, 1913–1921 (2015).
33. Sagonas, C., Tzimiropoulos, G., Zafeiriou, S. & Pantic, M. 300 faces in-the-wild challenge: The first facial landmark localization challenge. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, 397–403 (2013).
34. Andriluka, M., Pishchulin, L., Gehler, P. & Schiele, B. 2d human pose estimation: New benchmark and state of the art analysis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 3686–3693 (2014).
35. Eslami, M., Neuschaefer-Rube, C. & Serrurier, A. Automatic vocal tract segmentation based on conditional generative adversarial neural network. *Studientexte zur Sprachkommunikation: Elektronische Sprachsignalverarbeitung 2019* 263–270 (2019).
36. Valdés Vargas, J. A. *Adaptation of orofacial clones to the morphology and control strategies of target speakers for speech articulation*. Ph.D. thesis, Université de Grenoble (2013).
37. Comité de protection des personnes (cpp), <https://www.iledefrance.ars.sante.fr/comites-de-protection-des-personnes-cpp>. [Online; accessed 3-Oct-2019].
38. Honari, S., Yosinski, J., Vincent, P. & Pal, C. Recombinator networks: Learning coarse-to-fine feature aggregation. In *Computer Vision and Pattern Recognition (CVPR)*, 2016 *IEEE Conference on* (IEEE, 2016).
39. Kazemi, V. & Sullivan, J. One millisecond face alignment with an ensemble of regression trees. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1867–1874 (2014).
40. He, K., Zhang, X., Ren, S. & Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778 (2016).
41. Kopaczka, M., Schock, J. & Merhof, D. Super-realtime facial landmark detection and shape fitting by deep regression of shape model parameters. *arXiv preprint arXiv:1902.03459* (2019).
42. Yang, J., Liu, Q. & Zhang, K. Stacked hourglass network for robust facial landmark localisation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 79–87 (2017).
43. Chu, X. *et al.* Multi-context attention for human pose estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 1831–1840 (2017).
44. Ronneberger, O., Fischer, P. & Brox, T. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, 234–241 (Springer, 2015).
45. Isola, P., Zhu, J.-Y., Zhou, T. & Efros, A. A. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1125–1134 (2017).

46. Wang, T.-C. *et al.* High-resolution image synthesis and semantic manipulation with conditional gans. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 8798–8807 (2018).
47. Bloice, M. D., Stocker, C. & Holzinger, A. Augmentor: an image augmentation library for machine learning. *arXiv preprint arXiv:1708.04680* (2017).
48. Takemoto, H., Kitamura, T., Nishimoto, H. & Honda, K. A method of teeth superimposition on MRI data for accurate measurement of vocal tract shape and dimensions. *Acoust. Sci. Technol.* **25**, 468–474 (2004).
49. Ananthakrishnan, G., Badin, P., Valdés Vargas, J. A. & Engwall, O. Predicting unseen articulations from multi-speaker articulatory models. In *Proceedings of Interspeech 2010* (Makuhari, Japan, 2010).
50. Zheng, Q., Yang, M., Yang, J., Zhang, Q. & Zhang, X. Improvement of generalization ability of deep cnn via implicit regularization in two-stage training process. *IEEE Access* **6**, 15844–15869 (2018).
51. Goodfellow, I., Bengio, Y. & Courville, A. *Deep learning*. (MIT press, 2016).
52. Zhou, Z.-H. & Feng, J. Deep forest: Towards an alternative to deep neural networks. *arXiv preprint arXiv:1702.08835* (2017).
53. Orlando, J. I., Prokofyeva, E., del Fresno, M. & Blaschko, M. B. An ensemble deep learning based approach for red lesion detection in fundus images. *Comput. methods and programs in biomedicine* **153**, 115–127 (2018).

Acknowledgements

The authors are very grateful to P. Badin for providing the data, L. Lamalle for recording them, and J.-A. Valdés Vargas and G. Ananthakrishnan for performing the majority of the initial landmark labelling. This research project is supported by the START-Program of the Faculty of Medicine, RWTH Aachen. The data component of this work has been partially funded by the French ANR (grant ANR-08-EMER-001-02 ‘ARTIS’).

Author contributions

M.E. and A.S. conceived the study, processed the data and wrote the manuscript, M.E. designed the networks and conducted the experiment, M.E., C.N.-R. and A.S. analyzed the results and reviewed the manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Correspondence and requests for materials should be addressed to M.E. or A.S.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher’s note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2020