



# HHS Public Access

Author manuscript

*IEEE Trans Med Imaging*. Author manuscript; available in PMC 2020 April 01.

Published in final edited form as:

*IEEE Trans Med Imaging*. 2019 April ; 38(4): 919–931. doi:10.1109/TMI.2018.2875814.

## Deep Geodesic Learning for Segmentation and Anatomical Landmarking

**Neslisah Torosdagli,**

Center for Research in Computer Vision at University of Central Florida, Orlando, FL

**Denise K. Liberton,**

Craniofacial Anomalies and Regeneration section, National Institute of Dental and Craniofacial Research (NIDCR), National Institutes of Health (NIH), Bethesda, MD

**Payal Verma,**

Craniofacial Anomalies and Regeneration section, National Institute of Dental and Craniofacial Research (NIDCR), National Institutes of Health (NIH), Bethesda, MD

**Murat Sincan,**

University of South Dakota Sanford School of Medicine, Sioux Falls, SD

**Janice S. Lee, and**

Craniofacial Anomalies and Regeneration section, National Institute of Dental and Craniofacial Research (NIDCR), National Institutes of Health (NIH), Bethesda, MD

**Ulas Bagci [Senior Member, IEEE]**

Center for Research in Computer Vision at University of Central Florida, Orlando, FL

### Abstract

In this study, we propose a novel deep learning framework for anatomy segmentation and automatic landmarking. Specifically, we focus on the challenging problem of mandible segmentation from cone-beam computed tomography (CBCT) scans and identification of 9 anatomical landmarks of the mandible on the geodesic space. The overall approach employs three inter-related steps. In step 1, we propose a deep neural network architecture with carefully designed regularization, and network hyper-parameters to perform image segmentation without the need for data augmentation and complex post-processing refinement. In step 2, we formulate the landmark localization problem directly on the geodesic space for sparsely-spaced anatomical landmarks. In step 3, we utilize a long short-term memory (LSTM) network to identify closely-spaced landmarks, which is rather difficult to obtain using other standard networks. The proposed fully automated method showed superior efficacy compared to the state-of-the-art mandible segmentation and landmarking approaches in craniofacial anomalies and diseased states. We used a very challenging CBCT dataset of 50 patients with a high-degree of craniomaxillofacial (CMF) variability that is realistic in clinical practice. Qualitative visual inspection was conducted for distinct CBCT scans from 250 patients with high anatomical variability. We have also shown the state-of-the-art performance in an independent dataset from MICCAI Head-Neck Challenge (2015).

## Keywords

Mandible Segmentation; Craniomaxillofacial Deformities; Deep Learning; Convolutional Neural Network; Geodesic Mapping; Cone Beam Computed Tomography (CBCT)

---

## I. Introduction

In the United States, there are more than 17 million patients with congenital or developmental deformities of the jaws, face, and skull, also defined as the craniomaxillofacial (CMF) region [1]. Trauma, deformities from tumor ablation, and congenital birth defects are some of the leading causes of CMF deformities [1]. The number of patients who require orthodontic treatment is far beyond this number. Among CMF conditions, the mandible is one of the most frequently deformed or injured regions, with 76% of facial trauma affecting the mandibular region [2].

The ultimate goal of clinicians is to provide accurate and rapid clinical interpretation, which guides appropriate treatment of CMF deformities. Cone-beam computed tomography (CBCT) is the newest conventional imaging modality for the diagnosis and treatment planning of patients with skeletal CMF deformities. Not only do CBCT scanners expose patients to lower doses of radiation compared to spiral CT scanners, but also CBCT scanners are compact, fast and less expensive, which makes them widely available. On the other hand, CBCT scans have much greater noise and artifact presence, leading to challenges in image analysis tasks.

CBCT-based image analysis plays a significant role in diagnosing a disease or deformity, characterizing its severity, planning the treatment options, and estimating the risk of potential interventions. The core image analysis framework involves the detection and measurement of deformities, which requires precise segmentation of CMF bones. Landmarks, which identify anatomically distinct locations on the surface of the segmented bones, are placed and measurements are performed to determine the severity of the deformity compared to traditional 2D norms as well as to assist in treatment and surgical planning. Figure 1 shows nine anatomical landmarks defined on the mandible.

Currently, the landmarks have not evolved from traditional 2D anatomical landmarks for cephalometric analysis though 3D imaging has become more commonplace for clinical application. Landmarking on CT images is tedious and manual or semi-automated and prone to operator variability. Despite some recent elaborative efforts towards making a fully automated and accurate software for segmentation of bones and landmarking for deformation analysis in dental applications [3], [4], the problem remains largely unsolved for global CMF deformity analysis, especially for those who have congenital or developmental deformities for whom the diagnosis and treatment planning are most critically needed.

The main reason for this research gap is high anatomical variability in the shape of these bones due to their deformities in such patient populations. Figure 2 shows some of the known CMF deformities and artifacts, including missing bones (hence missing landmarks) or irregularities from the underlying disease or the surgical treatment (Figures 2a-2b),

varying number of teeth including missing teeth and unerupted deciduous teeth distorting the anatomy (Figures 2c-2d), and surgical interventions such as implants or surgical plates and screws that are necessary to treat the injury or deformity (Figures 2e-2f). Other reasons are image/scanner based artifacts/problems such as noise, truncation, beam hardening, and low resolution. Unlike existing methods focusing on dental applications with relatively small anatomical variations, there is a strong need for creating a general purpose, automated CMF image analysis platform that can help clinicians create a segmentation model and find anatomical landmarks for extremely challenging CMF deformities.

The overarching goal of our study is to develop a fully-automated image analysis software for mandible segmentation and anatomical landmarking that can overcome the highly variable clinical phenotype in the CMF region. This program will facilitate the ease of clinical application and permit the quantitative analysis that is currently tedious and prohibitive in 3 *D* cephalometric and geometric morphometrics. To this end, we include a landmarking process as a part of the segmentation algorithm to make geometric measurements more accurate, easier, and faster than manual methods. Our main dataset includes patients with congenital deformities fading to extreme developmental variations in CMF bones. The patient population is highly diverse, consisting of a wide range of ages across both sexes, imposing additional anatomical variability apart from the deformities.

Our proposed algorithm included three inter-connected steps (See Figures 3 and 4 for the overview of the proposed method and a sample processing pipeline for a single CBCT scan). For the first step, we designed a new convolutional neural network (CNN) architecture for mandibular segmentation from 3D CBCT scans. For Step 2, we presented a learning-based geodesic map generation algorithm for each sparsely-spaced anatomical landmark defined on the mandible. For Step 3, inspired by the success of recurrent neural networks (RNN) for capturing temporal information, we demonstrated a long short-term memory (LSTM) based algorithm to capture the relationship between closely-spaced anatomical landmarks as a sequential learning process.

## II. Related Work

The mandible is the lower jaw bone and it is the only mobile bone in the CMF region. It is the largest, the strongest, and the most complex bone in the CMF region that houses the lower teeth as well as canals with blood vessels and nerves. Due to its complex structure and the significant structural variations of patients with CMF disorders, segmentation and landmark localization in the mandibular region is a very challenging problem (See Figure 2). Although, there are efforts with promising performances [3], [5], [6], [7], the literature still lacks a fully-automated, fast, and generalized software solution in response to a wide range of patient ages, deformities, and the imaging artifacts. Hence, the current convention used in clinics is still either manual segmentation and annotations, or semi-automated with software support such as (in alphabetical order) 3dMDvultus (3dMD, Atlanta, Ga), Dolphin Imaging (Dolphin Imaging, Chatsworth, Ca), and InVivoDental (Anatomage, San Jose, Ca).

Over the past decade, there have been significant improvements in mandible segmentation and landmarking using registration-based (atlas-based), model-based, and more recently

machine learning-based approaches [9]. Although, registration-based methods have been reported to achieve relatively high accuracy when shape and appearance information are integrated, these algorithms perform poorly when there are variations due to aging (i.e., pediatrics vs. adults), missing teeth, missing parts of the region of interest, and imaging artifacts [4], [10], [11]. In 2015, Gupta et al. [12] developed a knowledge-based algorithm to localize 20 anatomical landmarks on the CBCT scans. Despite the promising results, in cases of missing lower incisors, mandible fractures, or other anatomical deformities that directly alter the anterior mandible, an error in seed localization can lead to a sub-optimal outcome. In 2016, Zhang et al. [5] digitized CMF landmarks on CBCT scans using a regression forest-based landmark detector. Image segmentation was used as a guidance to address the spatial coherence of landmarks. The authors obtained a mean digitization error less than 2mm for 15 CMF landmarks. In 2017, Zhang et al. [3] improved their method by proposing a joint CMF bone segmentation and landmark digitization framework via a context-guided multi-task fully convolutional neural network (FCN) adopting a U-Net architecture. The spatial context of the landmarks were grasped using 3D displacement maps. An outstanding segmentation accuracy (dice similarity coefficient of  $93.27 \pm 0.97\%$ ) was obtained along with a mean digitization error of less than 1.5 mm for identifying 15 CMF landmarks. Despite these promising performances, the study had the limitation of working on small number of landmarks due to memory constraints. That is, if there are  $N_l$  landmarks and each patient's 3D scan is composed of  $V$  voxels, 3D displacement maps use  $3 \times N_l \times V$  memory as input to the 2<sup>nd</sup> U-Net. Furthermore, most of the slices in the displacement maps were the same due to planar spatial positions, leading to inefficiency because of redundant information.

In a more conventional way, Urschler et al. [7] combined image appearance information and geometric landmark configuration into a unified random forest framework, and performed a coordinate descent optimization procedure that iteratively refines landmark locations jointly. The authors achieved a high performance on MRI data with only a small percentage of outliers.

### A. Our Contribution

To date, little research has been carried out involving deep learning-based segmentation of CMF bones and landmarking. Herein, we demonstrate in-depth mandible segmentation and landmarking in a fully automated way, and we propose novel techniques that enhance accuracy and efficiency to improve the state-of-the-art approaches in the setting of high degree of anatomical variability. The latter function is highly critical as previous methods have been developed based on optimized and normal patient cohorts, but the limitations of these methods are evident in diseased and pathological cohorts and fallshort of clinical utilization and application. Specifically, our contributions can be summarized as follows:

- Our proposed method is unique because we propose a fully automated system with a geodesic map of bones automatically injected into the deep learning settings unlike the state-of-the-art deep approaches where landmarks are annotated in Euclidean space [3], [13].

- While other works learn landmark locations using only spatial location of the landmarks along with context information, we learn the sparsely-spaced landmark relationship on the same bone by utilizing a U-Net based landmark detection algorithm. Then, an LSTM based algorithm is developed to identify closely-spaced landmarks. We consider the landmarks as states of the LSTM, and operate on the geodesic space. This approach is not only realistic, but also computationally more feasible.
- We present in-depth analysis of architecture parameters such as the effect of growth rate in segmentation, the different pooling functions both in detection and segmentation tasks, and the harmony of dropout regularization with pooling functions.
- Our dataset includes highly variable bone deformities along with other challenges of the CBCT scans. For an extremely challenging dataset, the proposed geodesic deep learning algorithm is shown to be robust by successfully segmenting the mandible bones and providing highly accurate anatomical landmarks.

### III Methods

The proposed system comprises three steps (see Figure 3 for the overview). Step 1 includes a newly proposed segmentation network for mandible based on a unified algorithm combining U-Net and DenseNET with carefully designed network architecture parameters and a large number of layers. In Step 2, we propose a U-Net based geodesic learning architecture to learn true and more accurate spatial relationships of anatomical landmarks on the segmented mandible. Finally, in Step 3, we identify closely-spaced landmark locations by a classification framework where we utilize an LSTM network.

#### A. Step 1: Segmentation Network

CNN based approaches such as U-Net [8], fully convolutional network (FCN) [14], and encoder-decoder CNNs [15] have achieved increasing success in image segmentation. These methods share the same spirit of obtaining images at different resolutions by consecutive downsampling and upsampling to make pixel level predictions. Despite the significant progress made by such standard approaches towards segmentation, they often fail to converge in training when faced with objects with high variations in shape and/or texture, and complexities in the structure. Another challenge is the optimization of massive amount of hyper-parameters in deep nets. Inspired by the recently introduced notion of densely connected networks (DenseNET) for object recognition [16], a new network architecture was presented by Jégou et al. [17] for semantic segmentation of natural images, called Fully Convolutional DenseNET (or *Tiramisu* in short). In this study, we adapted this Tiramisu network for medical image segmentation domain through significant modifications:

- (1) We set all default pooling functions (often they are defined as max pooling) with average pooling to increase pixel-level predictions. Although pooling functions in the literature have been reported to perform similarly in various tasks, we hypothesized that average pooling was more suitable for pixel level predictions

because average pooling identifies the extent of an object while max-pooling only identifies the discriminative part.

- (2) We explored the role of dropout regularization on segmentation performance with respect to the commonly used batch normalization (BN) and pooling functions. Literature provides mixed evidence for the role of these regularizers.
- (3) We investigated the effect of growth rate (of dense block(s)) on the segmentation performance. While a relatively small growth rate has been found successful in various computer vision tasks, the growth rate of dense blocks is often fixed and its optimal choice for segmentation task has not been explored yet.
- (4) We examined appropriate regularization as well as network architecture parameters, including number of layers, to avoid the use of post-processing methods such as CRF (conditional random field). It is common in many CNN-based segmentation methods to use such algorithms so that the model predictions are further refined because the segmentation accuracy is below an expected range.

Figure 5 illustrates the Tiramisu network architecture (a) and the content of a dense block (b), respectively. Tiramisu network is extremely deep, including 103 layers [17] as compared to the U-Net which has only 19 layers in our implementation. The input of the Tiramisu network was the 2D sagittal slices of the CBCT scan of patients with CMF deformities, and the output was the binary 2D sagittal slices with the mandible segmented (see Figure 4 for an example workflow of a 2D slice). The architecture consisted of 11 dense blocks with 103 convolutional layers. Each dense block contained a variable length of layers and the growth rate was set specifically for each dense block based on extensive experimental results and comparison. The network was composed of approximately 9M trainable parameters. We trained the revised Tiramisu from scratch without the need for data augmentation and complex post-processing. Details of the network parameters are given in Tables I and II.

## B. Step 2: Geodesic Learning for Landmarking

We approach the problem of anatomical landmarking (landmark detection) as a learning problem. The state-of-the-art method in the literature adopts a U-Net architecture to learn the locations of the anatomical landmarks [3]. For a given 3D CBCT scan  $X$  and a landmark  $l$ , authors [3] created three displacement maps  $D^{l,x}$ ,  $D^{l,y}$ ,  $D^{l,z}$  corresponding to  $x, y$ , and  $z$  axes [3]. That is, if there are  $N_l$  landmarks,  $N_l \times 3$  displacement maps are generated. Displacement maps, also called *heatmaps*, were created using a simple Euclidean metric measuring the distance of a landmark to a reference point ((0,0) index of image). Although the method is simple to implement and efficient within the multi-task learning platform, it does not incorporate information about the object of interest (mandible) and works on the image space. In addition, the method generates a large number of heatmaps when the number of landmarks is high. Lastly, the method operates directly on the Euclidean space and it does not capture the underlying data distribution, which is non-Euclidean in nature.

To alleviate these problems and to solve the landmarking problem directly on the shape space, we propose to use a Geodesic Distance Transform to learn the relationship of

landmarks directly on the shape space (mandible surface). To this end, we first apply linear time distance transform (LTDT) [18] to the segmented mandible images (i.e., binary) and generate signed distance maps. Assuming  $\mathbf{I}$  is a 3D-segmented binary image (mandible) obtained at Step 1 from a given CBCT scan  $\mathbf{X}$  in the domain  $\Omega = \{1, \dots, n\} \times \{1, \dots, m\}$ , Mandible  $M$  is represented by all white voxels ( $I(v) = 1$ ), while Mandible complement (background)  $M^C$  is represented by all black voxels ( $I(v) = 0$ ) [19]:

$$\begin{aligned} M &= \{v \in \Omega \mid I(v) = 1\}, \\ M^C &= \{v \in \Omega \mid I(v) = 0\}. \end{aligned} \quad (1)$$

LTDT represents a map such that each voxel  $v$  is the smallest Euclidean distance from this voxel to the  $M^C$ :

$$LTDT(v) = \min \{ \text{dist}(v, q) \mid q \in M^C \}. \quad (2)$$

Then, the signed LTDT, namely sLTDT, of  $\mathbf{I}$  for a voxel  $v$  can be represented as:

$$sLTDT(v) = \begin{cases} LTDT(v) & \text{if } v \in M, \\ -\min \{ \text{dist}(v, q) \mid q \in M \} & \text{if } v \in M^C. \end{cases} \quad (3)$$

For each landmark  $l$ , we generate a geodesic distance map  $D_l^G$ . To do so, we find the shortest distance between landmark  $l$  and each voxel  $v$  as:

$$D_l^G(v) = \begin{cases} \min \pi(l, v) & \text{if } v \in M, \\ \inf & \text{if } v \in M^C, \end{cases} \quad (4)$$

where  $\pi$  indicates all possible paths from the landmark  $l$  to the voxel  $v$  ( $v \in \mathbf{M}$ ). Since the shortest distance between two points is found on the surface, it is called geodesic distance [20], [21] as a convention. To find the shortest path  $\pi$ , we applied Dijkstra's shortest path algorithm. For each landmark  $l$ , we generated one geodesic map as  $D_l^G$ . For multiple landmarks, as is the case in our problem, we simply combined the geodesic maps to generate one final heatmap, which includes location information for all landmarks. Final geodesic map for all landmarks was obtained through hard minimum function  $D_l^G = \min(D_l^G \circ D_{l_2}^G \circ \dots \circ D_{l_n}^G)$ , where  $\circ$  indicates pixel-wise comparison of all maps. In other words, the final geodesic map  $D_l^G$  includes  $n$  extrema (minimum) identifying the locations of the  $n$  landmarks.

To learn the relationship of  $n$  landmark points on the mandible surface, we designed a landmark localization network, based on the Zhang's U-Net architecture [3]. Tiramisu



network could perhaps be used for the same purpose. However, the data was simplified in landmark localization due to geodesic distance mapping, and Zhang’s U-Net uses only 10% of the overall parameter space for landmark localization. The improved Zhang’s U-Net accepts  $2D$  slices of the signed distance transform of the segmented mandible ( $\mathbf{I}$ ) as the input, and produces the  $2D$  geodesic map ( $D_I^G$ ) revealing the location of  $N_I$  landmarks as the output. The details of the landmark localization architecture (improved version of the Zhang’s U-Net) with 19 layers and parameters are given in Tables III and IV, respectively. Briefly, the encoder path of the U-Net was composed of 3 levels. Each level consisted of (multiple) application(s) of convolutional nodes:  $5 \times 5$  convolutions, batch normalization (BN), rectified linear unit (ReLU), and dropout. Between each level max pooling, downsampling with a stride of 2, was performed. Similar to the encoder path, the decoder path was also composed of 3 levels. In contrast to encoder path, dropout was not applied in the decoder path. Between the levels in the decoder path, upsampling operation was applied. To emphasize the high-resolution features that may be lost in the encoder path, copy operation was used in the decoder path. Copy operation, as the name implies, concatenated the features at the same  $2D$  resolution levels from the encoder path to the decoder path.

We have chosen the optimization algorithm as RM-SProp [22] due to its fast convergence and adaptive nature. The initial learning rate was set to  $1e-3$  with an exponential decay of 0.995 after each epoch (Table IV). At the end of the decoder path, softmax cross entropy was applied as a loss function because mean squared error (MSE) caused serious convergence issues. We quantized the geodesic map in the range  $[0 - 20]$ , where the limit 20 was set empirically. The network was composed of  $\approx 1M$  trainable parameters. Compared to the Zhang’s U-Net [3], in our improved implementation, in addition to the  $5 \times 5$  convolutions, on the expanding path at level 2, we kept the symmetry in the number of features obtained as in the contracting path. These alterations made sure Zhang’s U-Net to work without failures.

### C. Step 3: Localization of Closely-Spaced Landmarks

Fusion of geodesic maps through pixel-wise hard-coded minimum function is reliable when landmarks are sufficiently distant from each other. In other words, if landmarks are very close to each other, then the combined geodesic map  $D_I^G$  may have instabilities in locating its extrema points. In particular for our case, it was not possible to localize specifically “Menton” and other mid-sagittal closely-spaced landmarks in a clinically acceptable error range (i.e.,  $\leq 3mm$ ). In order to avoid such scenarios, we divided the landmarking process into two distinct cases: learning closely-spaced and sparsely-spaced landmarks separately. First, we classified the mandible landmarks into sparsely and closely-spaced sets. Sparsely-spaced landmarks ( $N=5$ ) were defined in the inferior, superior-posterior-left, superior-posterior-right, superior-anterior-left, and superior-anterior-right regions. Closely-spaced landmarks ( $N=4$ ) were defined as the ones that were closely tied together (Infradentale ( $I_d$ ), B point ( $B$ ), Pogonion ( $P_g$ ), and Gnathion ( $G_n$ )).

Note that these anatomical landmarks often reside on the same sagittal plane in the same order according to the midpoint of the lower-jaw incisors. We propose to capture this order dependence by using an LSTM architecture in the sagittal axis of the images containing the



landmark “Menton” (Figure 6). The rationale behind this choice was that LSTM network is a type of RNN introduced by Hochreiter et al. [23] in 1997, modeling the temporal information of the data effectively. Although the imaging data that we used does not include temporal information in the standard meaning, we modeled the landmark relationship as a temporal information due to their close positioning in the same plane. This phenomenon is illustrated in Figure 6. The input data to the LSTM network was a  $64 \times 64$  mandible binary boundary image of the sagittal plane of the landmark  $Me$ , and the output was a vector of 0’s and 1’s: while 0 refers to non-landmark location, 1 refers to a landmark location in the sagittal axis. Figure 7 shows further details of the LSTM network and content of a sample LSTM block that we used for effective learning of closely-spaced landmarks.

To generate the training data, the sagittal slice containing the closely-spaced landmarks “Menton”, “Gnathion”, “Pogonion”, “B-point” and “Infradentale” was scaled into a binary boundary image of size  $64 \times 64$ . The 5 landmark locations (marked by red circles in Figure 6) on this boundary image were parameterized as  $(x,y)$ , where  $y$  is the row number in the range 0 to 64, and  $x$  is the white boundary column number of the corresponding row  $y$ .

LSTM network was composed of 64 cells (Figure 7), and each cell in the LSTM network consisted of 512 units. The training images were row-wise input to the LSTM network such that  $n^{th}$  row was input to the corresponding  $n^{th}$  cell of the network. The output of each cell was multiplied by  $512 \times 2$  weight and  $1 \times 2$  bias was added. The resultant  $1 \times 2$  tensors at each cell were concatenated and softmax cross entropy was applied as a loss function.

#### D. Training Framework: End-to-end vs. Sequential vs Mixed

Since the proposed learning system is complex, it is worth to explore whether gradient-descent learning system can be applied to the system as a whole (called *end-to-end*). For this purpose, first, we evaluated the performances of each network individually, so named *sequential training* followed by an engineering approach for concatenation of the three networks. Since end-to-end learning systems require all modules of the complex system to be differentiable, our proposed system was not fully eligible for this learning type. It is because the 3<sup>rd</sup> module (LSTM network for closely-spaced landmark localization) had differentiability issues for the given loss function. Therefore, we trained the first and second modules in an *end-to-end* manner while integrating the third network module into this system sequentially. In summary, we devised two alternative methods to solve our overall goal: in the first solution, the overall system was considered as a sequential system. In the second solution, the first two modules of the system were trained in an end-to-end manner with the inclusion of the third module as a sequential block. Owing to the usage of sequential and end-to-end frameworks together, we named the second solution as “mixed”.

Although end-to-end networks are conceptually and mathematically beautiful, it has a strict condition that each module should be differentiable with respect to the loss function so that a positive impact can be obtained on the final objective. However, as stated in [24] and [25], when some modules are not differentiable (as the third module of our proposed method), or when the system is too complex with sparse modules, the overall results may be inferior compared to the sequential method. Due to the differentiability issue in the third module, our system falls into this category. That is, the input to the 3<sup>rd</sup> module was the 2D sagittal slice

containing the anatomical landmark “Menton”. Since not every output slice in the training could be used for the 3<sup>rd</sup> module, differentiability was lost. In addition, we observed that unless the first two modules were close to the converging state, it was not possible to locate “Menton” more precisely than a random guess. Due to the requirement of convergence within this module, eventually, it was not possible to apply LSTM training in a truly end-to-end manner.

## IV. Experiments and Results

### A. Data description:

Anonymized CBCT scans of 50 patients (30 female and 20 male, mean age = 22.4 years, standard deviation = 9.6 years) were included in our analysis through an IRB-approved protocol and data sharing agreement between UCF and NIH. These patients had craniofacial congenital birth defects, developmental growth anomalies, trauma to the CMF, surgical intervention, and included pediatric and adult patients. All images were obtained on a CB MercuRay CBCT system (Hitachi Medical Corporation, Tokyo, Japan). The 12-inch field of view was required for this study to capture the entire length of the airway and was scanned at 10 mA and 100 Kvp. The equivalent radiation dosage for each scan was approximately 300 mSv. After the study had begun, the machine was modified to accommodate 2 mA for the same 12-inch field of view, thus lowering the equivalent radiation dosage for each scan to approximately 132.3 mSv. Each patient’s scan was re-sampled from  $512 \times 512 \times 512$  to  $256 \times 256 \times 512$  to reduce computational cost. In-plane resolution of the scans was noted either as  $0.754\text{mm} \times 0.754\text{mm} \times 0.377\text{mm}$  or  $0.584\text{mm} \times 0.584\text{mm} \times 0.292\text{mm}$ . Apart from highly diverse nature of this data set, the following image-based variations have also been confirmed: aliasing artifacts due to braces, metal alloy surgical implants (screws and plates), dental fillings, and missing bones or teeth.

Additionally, we tested and evaluated our algorithm(s) using the MICCAI Head-Neck Challenge 2015 dataset [26]. MIC-CAI Head-Neck Challenge 2015 dataset was composed of manually annotated CT scans of 48 patients from the Radiation Therapy Oncology Group (RTOG) 0522 study (a multi-institutional clinical trial [27]). For all data, the reconstruction matrix was  $512 \times 512$  pixels. The in-plane pixel spacing was isotropic, and varied between  $0.76\text{mm} \times 0.76\text{mm}$  and  $1.27\text{mm} \times 1.27\text{mm}$ . The range of the number of slices of the scans were 110–190. The spacing in the z-direction was between  $1.25\text{mm}$  and  $3\text{mm}$  [26]. In the challenge, there were three test results provided, where test data part 1 (off-site data) and part 2 (on-site data) did not have publicly available manual annotations to compare to our performances. Hence, we compared our test results to the the cross-validation results as provided in [28].

**Training deep networks:** We have trained our deep networks with 50 patients’ volumetric CBCT scans in a 5-fold cross validation experimental design. Since each patient’s scan includes 512 slices (i.e., 2D images with  $256 \times 256$  pixels inplane), we had a total of 25,600 images to train and test. In each training experiment, we have used 20,480 2D images to train the network while the remaining slices (5,120) were used for testing. This

procedure was repeated for each fold of the data, and average of the overall scores were presented in the following subsections.

## B. Evaluation metrics and annotations:

Three expert interpreters annotated the data (one from the NIH team, two from the UCF team) independently. Interobserver agreement values were computed based on these three annotations. Later, second and third experts (from the UCF team) repeated their manual annotations (after one month period of their initial annotations) for intra-observer evaluations. Experts used freely available 3D Slicer software for the annotations. Annotated landmarks were saved in the same format of the original images, where landmark positions in a neighborhood of  $3 \times 3 \times 3$  were marked according to the landmark ID while the background pixels were marked as 0.

A simple median filtering was used to minimize noise in the scans. No other particular preprocessing algorithm was used. Experiments were performed through a 5-fold crossvalidation method. Intersection of Union (IoU) metric was used to evaluate object detection performance. For evaluating segmentation, we used the standard DSC (dice similarity coefficient), Sensitivity, Specificity, and HD (Hausdorff Distance) (100% percentile). As a convention, high DSC, sensitivity, specificity and low HD indicate a good performance. The accuracy of the landmark localization was evaluated using the detection error in pixel space within a  $3 \times 3 \times 3$  bounding box. Inter-observer agreement rate was found to be 91.69% for segmentation (via DSC).

## C. Evaluation of Segmentation

The proposed segmentation framework achieved highly accurate segmentation results despite the large variations in the imaging data due to severe CMF deformities. Table V summarizes the segmentation evaluation metrics and number of parameters used for the proposed and the compared networks. The proposed segmentation network outperformed the state-of-the-art U-Net [8]. Specifically, we have improved the success of the baseline U-Net framework by increasing the number of layers into 19. In terms of the dice similarity metric, both improved Zhang's U-Net and the proposed segmentation network were statistically significantly better than the baseline U-Net ( $P=0.02$ , t-test). In summary, (i) there is no statistically significant difference noted between our proposed method and the manual segmentation method ( $P=0.77$ ); (ii) there is a statistically significant difference between our proposed method and the baseline U-Net ( $P=0.02$ ); (iii) there is no statistically significant difference noted between the proposed method and our improvement over the Zhangs U-Net ( $P=0.28$ ). It is also worth to note that the proposed Tiramisu network performed more robustly in training, converging faster than the improved Zhang's U-Net despite the larger number of parameters in the Tiramisu.

We evaluated the segmentation performances on different datasets and training styles (sequential vs. mixed learning) and summarized the results in Table VI. With the MICCAI Head-Neck Challenge 2015 dataset, we obtained a dice accuracy of 93.86% compared to 90% [28]. High accuracies of the MICCAI Head-Neck Challenge 2015 and the NIH datasets imply the robustness of the Tiramisu segmentation network. It should be noted that MICCAI

Head-Neck Challenge 2015 dataset contains mainly scans with imaging artifacts as well as different diseases. Closer inspection of Table VI also shows that a simple post-processing step such as “Connected Component Analysis” and “3D fill” were important to decrease the number of the false positives and false negatives in the challenge dataset. The slightly lower performances of mixed training with Tiramisu network for both segmentation and landmark localization can be explained by the increased number of parameters but insufficient dataset size to derive learning procedure as a whole. Sequential learning was sufficient to obtain good results in segmentation, though.

#### D. Evaluation of Landmark Localization

Ground truth annotations (manual segmentation and anatomical landmarking) were performed by three experts independently. Inter-observer agreement rate was 91.69%. Figure 8c presents per landmark and overall expert reading variations of landmarking in the pixel space. We observed that there was an average 3 pixel errors among the experts. Hence, any landmarking algorithm leading to error within 3 pixel range can be considered a clinically acceptable level of success. Figures 8a and 8b summarize the proposed algorithm’s landmark localization errors in the pixel space and the volume space, respectively.

The mean and median volume space errors for each landmark are presented at Table VII. The errors in the pixel space (Figure 8a) were less than 3 pixels for all 9 landmarks, indicating that our method is highly accurate and can be used for clinical applications as it results in less variations than the inter-observer variation rate as explained earlier (Figure 8c).

Figure 9 presents three experimental results when there was high morphological variation and deformity. In Figure 9a, due to the genioplasty with chin advancement and rigid fixation, there is a protuberance on the mandible distorting the normal anatomy. In Figure 9b, condyle-ramus unit is absent on the left side of the mandible due to a congenital birth defect. The Geodesic Landmark Localization network successfully detected 4 landmarks. Note that the fifth landmark was on the missing bone, and it was not located as an outcome of the landmarking process. This is one of the strengths of the proposed method. In Figure 9c, the patient had bilateral surgical implants along the ascending ramus (bicortical positional screws), and bilateral condyle and coronoid processes are fixed with these implants. The landmarking process was successful in this challenging case too.

We also evaluated the impact of segmentation accuracy on the landmark localization error (Figure 10). In this evaluation, we first grouped the testing scans into 2 groups according to their dice values as lower and higher segmentation accuracies (i.e.,  $\leq 90\%$  as lower,  $> 90\%$  as higher). Next, we compared the landmark localization errors in pixel space for these two groups. In Figure 10, the landmarking process was robust to changes in segmentation accuracy, and never reached more than 3 pixels errors. It should also be noted that the mean and median segmentation accuracy were still very high in our experiments, leading to successful landmark localization even at the low end of the dice values. Overall, the landmark localization was robust to the segmentation step and a potential (visible) error may happen only when the Menton (closely-spaced landmark) is located incorrectly due to a potential segmentation error.

Table VII summarizes the average and median errors of localized landmarks in millimeters with respect to different regularization methods, in particular pooling strategies. We observed that max pooling consistently outperformed other regularization methods. Unlike the segmentation problem, where average pooling was most effective in pixel level predictions, landmarking was driven by discriminative features, enhanced by max pool operation. All average and median errors of the landmark localizations were within the clinically acceptable limits (less than 3 pixels).

## V. Discussion and Conclusion

Overall, the proposed networks (Tiramisu and improved Zhang's U-Net) have enjoyed fast convergence (around 20 epochs) and high accuracy in a very challenging CBCT dataset. Tiramisu was observed to have better converging and training ability compared to improved Zhang's U-Net. For landmark localization, improved Zhang's U-Net in the geodesic space has performed comparably to the validated operator manual landmarking (e.g., median displacement error of 0 mm in most landmarks). We also addressed some of the poorly understood concepts in deep network architecture (particularly designed for medical image analysis applications) such as the use of dropout and pooling functions for regularization, activation functions for modeling non-linearity, and growth rate for information flow in densely connected layers (see Appendix A).

Fully convolution network (FCN) [14] has significantly changed the landscape of semantic image segmentation frameworks. Based on the FCN, Ronneberger et al. [8] introduced the U-Net which became the baseline for the current medical image segmentation tasks. The literature for particular medical image segmentation applications based on U-Net is vast; employing the encoder-decoder structure, dense connections, skip connections, residual blocks, and other types of architectural additions to improve segmentation accuracy for particular medical imaging applications. One major drawback of the U-Net framework is the inefficiency introduced by the significantly higher number of parameters to learn [29]. Hence, there is an anticipation for improvements in the efficiency and robustness of the U-Net type of architecture in the medical imaging field in the near future. One example of such studies, called Capsules [29], may be a good future alternative to what we propose herein.

In our study, we have focused on individual aspects of segmentation and landmarking, and have proposed novel architectural designs to address problems in both processes that have not been corrected in currently available systems. The natural extension of our work will be to formulate segmentation and landmarking problem within the multi-task learning algorithm, similar to the one proposed by Zhang et al. [3].

There are some limitations to our proposed method. Due to extensive memory and hardware requirement, we used pseudo-3D image analysis instead of fully 3D. A possible extension of our study will be to work on completely 3D space once hardware and memory supports are available. Another limitation of our work is to have a two-cascaded system for landmark localization instead of one because the hardcoded minimum function that we used for combining geodesic distances created additional artificial landmarks between those closely distributed landmarks. To overcome this problem, we showed a practical and novel use of

LSTM-based algorithm to learn the locations of closely-spaced landmarks and avoided such problems. Exploration of different functions other than hard-coded minimum for closely-spaced landmark localization is subject to further theoretical investigation in geodesic distance maps.

Future studies will include utilization for large cohort landmarking and analysis to establish normative craniofacial datasets. This fully automated method will enhance high throughput analysis of large, population-based cohorts. Additionally, studies on rare craniofacial disorders that often have anatomical variation will greatly benefit from the highly accurate landmark localization process.

## Appendix

### A. Evaluation of the Segmentation Network Parameters

#### 1) Effect of pooling functions:

After extensive experimental comparisons, we found that average pooling acts as a robust regularizer compared to other pooling functions such as max pooling and stochastic pooling (Table IX).

#### 2) Disharmony between BN and dropout:

We found that when BN is used in the network for segmentation purpose, the use of dropout is often detrimental except for only a drop rate of 20%. Similarly, we found that average pooling was the most robust pooling function compared to others when BN and dropout were used together.

#### 3) The role of growth rate in dense blocks:

Tiramisu network with 103 layers (growth rate of 16) has a proven success in the computer vision tasks. In our experiments, we observed that a Tiramisu network with a growth rate of 24 and drop rate of 0.2 produces the best accuracies instead of growth rate of 16 as in computer vision tasks (See Table X). Further, when no dropout is used (drop rate is 0), the growth rate performance inverses (See Table VIII), implying the regularizing impact of employing dropout on the neural networks.

**4) The choice of activation functions:** Although there have been many hand-designed activation functions proposed for deep networks, ReLU (rectified linear unit) became an almost standard choice for most CNNs. The main reason is due to its significant effect on the training dynamics and high task performances. More recently, another activation function, called “Swish” [30], was proposed. Unlike other activation functions, Swish was automatically determined based on a combination of exhaustive and reinforcement learning-based search. Authors showed that Swish tend to perform better than ReLU for very deep models. Since the proposed Tiramisu has 103 layers, we replaced all ReLU functions with Swish, which is a weighted sigmoid function  $f(x) = x \cdot \text{sigmoid}(\beta x)$ , and explored the network behaviors. We summarized the network performance in Table



VIII. Overall, we did not observe significant differences between ReLU and Swish, but ReLU led into slightly better results in all sub-experiments.

## B. Qualitative Evaluation

Our total CBCT dataset is composed of 250 patient CBCT images provided by our collaborators at the NIDCR/NIH. Only 50 of them were manually annotated by the three experts. To measure the performance of the algorithm on all available scans, by following the routine radiologic evaluation of the scans, two experts visually scored the segmentation results in the range from 0 to 4, where 1 is unacceptable, 2 is borderline, 3 is acceptable at clinical level, and 4 is superior (excellent) (Figures 11 and 12). When the scan is completely distorted or mandible does not exist in its entirety in the scan, it is not possible to automatically segment mandible, hence a score of 0 is given.

Scores of 3 and 4 represent clinically acceptable segmentation, where score 3 may correspond to minor deformations in the segmentation. The left top part of the Mandible (Figure 12-a6) was missing, for instance, but it was still precisely segmented. The mandible (Figure 12-c2), for another example, was composed of two separate parts, and the algorithm detected the larger portion of the mandible. Further analysis showed that the scans that were scored as 1 or 2 were typically the ones with serious anatomical deformations.

Approximately 5% of the segmentation results were scored as 1 by both experts **A** and **B** (Figure 11).

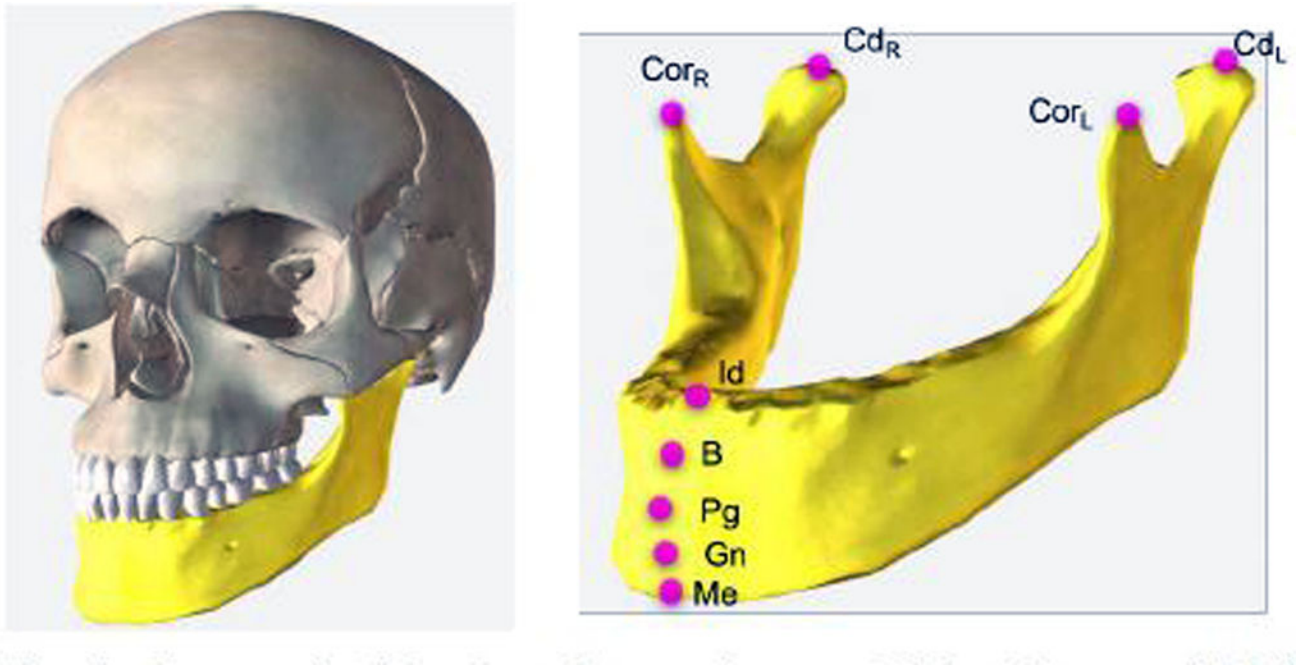
## References

- [1]. Xia JJ, Gateno J, and Teichgraber JF, "A paradigm shift in orthognathic surgery: A special series part i," *Journal of Oral and Maxillofacial Surgery*, vol. 67, no. 10, pp. 2093–2106, 2009. [PubMed: 19761903]
- [2]. Armond ACV, Martins C, Glria J, Galvo E, dos Santos C, and Falci S, "Influence of third molars in mandibular fractures. part 1: mandibular anglea meta-analysis," *International Journal of Oral and Maxillofacial Surgery*, vol. 46, no. 6, pp. 716–729, 2017. [PubMed: 28291569]
- [3]. Zhang J, Liu M, Wang L, Chen S, Yuan P, Li J, Shen SG-F, Tang Z, Chen K-C, Xia JJ, and Shen D, *Joint Craniomaxillofacial Bone Segmentation and Landmark Digitization by Context-Guided Fully Convolutional Networks*. Cham: Springer International Publishing, 2017, pp. 720–728.
- [4]. Anuwongnukroh N, Dechkunakorn S, Damrongsri S, Nilwarat C, Pudpong N, Radomsutthisarn W, and Kangern S, "Accuracy of automatic cephalometric software on landmark identification," *IOP Conference Series: Materials Science and Engineering*, vol. 265, no. 1, p. 012028, 2017.
- [5]. Zhang J, Gao Y, Wang L, Tang Z, Xia JJ, and Shen D, "Automatic craniomaxillofacial landmark digitization via segmentation-guided partially-joint regression forest model and multiscale statistical features," *IEEE Transactions on Biomedical Engineering*, vol. 63, no. 9, pp. 1820–1829, 9 2016. [PubMed: 26625402]
- [6]. Gao Y and Shen D, "Collaborative regression-based anatomical landmark detection," vol. 60, pp. 9377–9401, 11 2015.
- [7]. Urschler Martin, Ebner T, and tern D, "Integrating geometric configuration and appearance information into a unified framework for anatomical landmark localization," *Medical Image Analysis*, vol. 43, no. Supplement C, pp. 23–36, 2018. [PubMed: 28963961]
- [8]. Ronneberger O, Fischer P, and Brox T, *U-Net: Convolutional Networks for Biomedical Image Segmentation*. Cham: Springer International Publishing, 2015, pp. 234–241.
- [9]. Shen Dinggang, Wu G, and Suk H-I, "Deep learning in medical image analysis," *Annual Review of Biomedical Engineering*, vol. 19, no. 1, pp. 221–248, 2017.



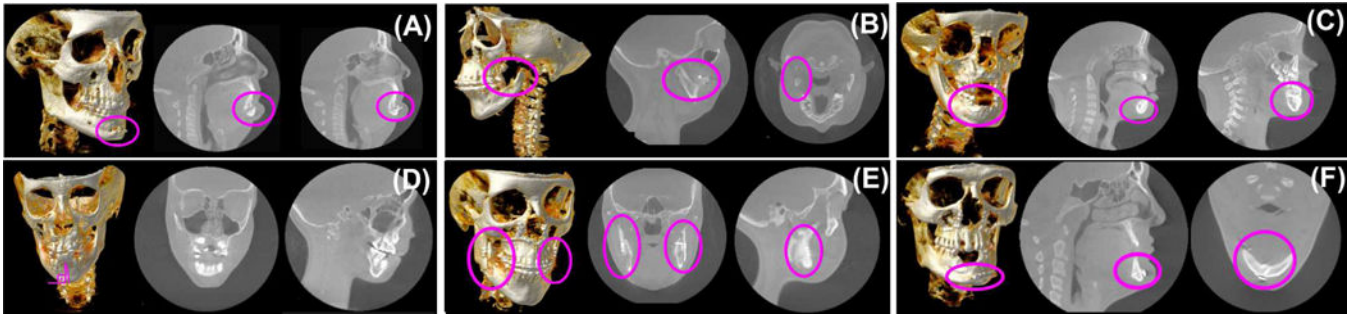
- [10]. Bahrampour S,E, Soltanimehr E, Zamani A, Oshagh M, Moattari M, and Mehdizadeh A, “The accuracy of a designed software for automated localization of craniofacial landmarks on cbct images,” *BMC Medical Imaging*, vol. 14, no. 1, p. 32, 9 2014. [PubMed: 25223399]
- [11]. Li X, Zhang Y, Shi Y, Wu S, Xiao Y, Gu X, Zhen X, and Zhou L, “Comprehensive evaluation of ten deformable image registration algorithms for contour propagation between ct and cone-beam ct images in adaptive head and neck radiotherapy,” *PLOS ONE*, vol. 12, no. 4, pp. 1–17, 4 2017.
- [12]. Gupta A, Kharbanda O, Sardana V, Balachandran R, and Sardana H, “A knowledge-based algorithm for automatic detection of cephalometric landmarks on cbct images,” *International Journal of Computer Assisted Radiology and Surgery*, vol. 10, no. 11, pp. 1737–1752, 11 2015. [PubMed: 25847662]
- [13]. Zhang J, Liu M, and Shen D, “Detecting anatomical landmarks from limited medical imaging data using two-stage task-oriented deep neural networks,” *IEEE Transactions on Image Processing*, vol. 26, no. 10, pp. 4753–4764, 10 2017. [PubMed: 28678706]
- [14]. Shelhamer E, Long J, and Darrell T, “Fully convolutional networks for semantic segmentation,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 39, no. 4, pp. 640–651, 4 2017. [PubMed: 27244717]
- [15]. Mortazi A, Karim R, Rhode K, Burt J, and Bageci U, “Cardiacnet: Segmentation of left atrium and proximal pulmonary veins from mri using multi-view cnn,” in *Medical Image Computing and Computer-Assisted Intervention MICCAI 2017*, Descoteaux M, Maier-Hein L, Franz A, Jannin P, Collins DL, and Duchesne S, Eds. Cham: Springer International Publishing, 2017, pp. 377–385.
- [16]. Huang Gao, Liu Z, and Weinberger KQ, “Densely connected convolutional networks,” *CoRR*, vol. abs/160806993, 2016 [Online]. Available: <http://arxiv.org/abs/1608.06993>
- [17]. Jégou Simon, Drozdal M, Vazquez D, Romero A, and Bengio Y, “The one hundred layers tiramisu: Fully convolutional densenets for semantic segmentation,” *CoRR*, vol. abs/161109326, 2016 [Online]. Available: <http://arxiv.org/abs/1611.09326>
- [18]. Breu H, Gil J, Kirkpatrick D, and Werman M, “Linear time euclidean distance transform algorithms,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 17, no. 5, pp. 529–533, 5 1995.
- [19]. Fabbri R, Costa LDF, Torelli JC, and Bruno OM, “2d euclidean distance transform algorithms: A comparative survey,” *ACM Comput. Surv.*, vol. 40, no. 1, pp. 2:1–2:44, 2 2008 [Online]. Available: <http://doi.acm.org/10.1145/1322432.1322434>
- [20]. Bose P, Maheshwari A, Shu C, and Wuhler S, “A survey of geodesic paths on 3d surfaces,” *Comput. Geom. Theory Appl*, vol. 44, no. 9, pp. 486–498, 11 2011.
- [21]. Datar M, Lyu I, Kim S, Cates J, Styner MA, and Whitaker R, *Geodesic Distances to Landmarks for Dense Correspondence on Ensembles of Complex Shapes*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2013, pp. 19–26.
- [22]. Hinton Geoff, RMSprop-Optimization algorithms, <https://www.coursera.org/learn/deep-neural-network/lecture/BhJlm/rmsprop>.
- [23]. Hochreiter S and Schmidhuber J, “Long short-term memory,” *Neural Comput*, vol. 9, no. 8, pp. 1735–1780, 11 1997. [PubMed: 9377276]
- [24]. Glasmachers T, “Limits of End-to-End Learning,” *ArXiv e-prints*, 4 2017.
- [25]. Shalev-Shwartz S, Shamir O, and Shammah S, “Failures of Gradient-Based Deep Learning,” *ArXiv e-prints*, 3 2017.
- [26]. Raudaschl PF, Zaffino P, Sharp GC, Spadea MF, Chen A, Dawant BM, Albrecht T, Gass T, Langguth C, Lthi M, Jung F, Knapp O, Wesarg S, Mannion-Haworth R, Bowes M, Ashman A, Guillard G, Brett A, Vincent G, Orbes-Arteaga M, Crdenas-Pea D, Castellanos-Dominguez G, Aghdasi N, Li Y, Berens A, Moe K, Hannaford B, Schubert R, and Fritscher KD, “Evaluation of segmentation methods on head and neck ct: Auto-segmentation challenge 2015,” *Medical Physics*, vol. 44, no. 5, pp. 2020–2036. [Online]. Available: <https://aapm.onlinelibrary.wiley.com/doi/abs/10.1002/mp.12197>
- [27]. Ang KK, Zhang Q, Rosenthal DI, Nguyen-Tan PF, Sherman EJ, Weber RS, Galvin JM, Bonner JA, Harris J, El-Naggar AK, Gillison ML, Jordan RC, Konski AA, Thorstad WL, Trotti A, Beitler JJ, Garden AS, Spanos WJ, Yom SS, and Axelrod RS, “Randomized phase iii trial of concurrent accelerated radiation plus cisplatin with or without cetuximab for stage iii to iv head and neck

- carcinoma: Rtog 0522,” *Journal of Clinical Oncology*, vol. 32, no. 27, pp. 2940–2950, 2014, PMID: [Online]. Available: 10.1200/JCO.2013.53.5633 [PubMed: 25154822]
- [28]. Mannion-Haworth R, Bowes M, Ashman A, Guillard G, Brett A, and Vincent G, “Fully automatic segmentation of head and neck organs using active appearance models,” 01 2016.
- [29]. LaLonde R and Bagci U, “Capsules for Object Segmentation,” *ArXiv e-prints*, 4 2018.
- [30]. Ramachandran P, Zoph B, and Le QV, “Searching for activation functions,” <https://arxiv.org/abs/1710.05941>, 2018.



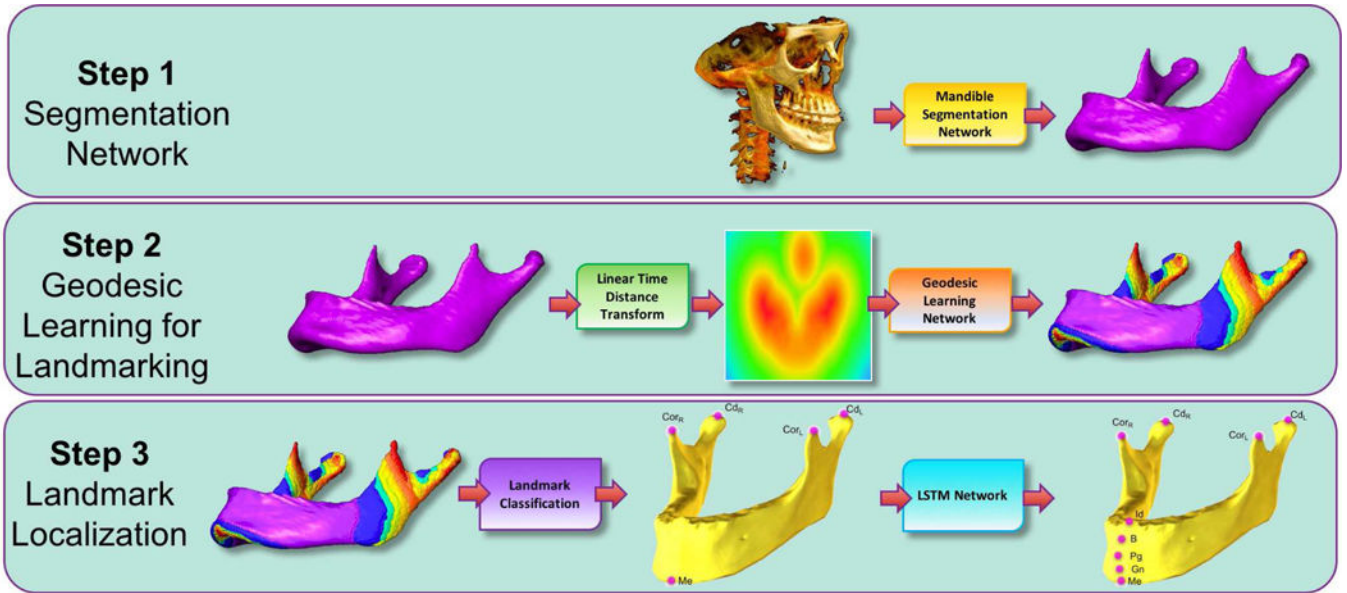
**Fig. 1:**

Anatomical landmarks on the mandible: Menton (*Me*), Gnathion (*Gn*), Pogonion (*Pg*), B Point (*B*), Infradentale (*Id*), Condylar Left (*Cd<sub>L</sub>*), Condylar Right (*Cd<sub>R</sub>*), Coronoid Left (*Cor<sub>L</sub>*), and Coronoid Right (*Cor<sub>R</sub>*). We aim to locate these landmarks automatically.

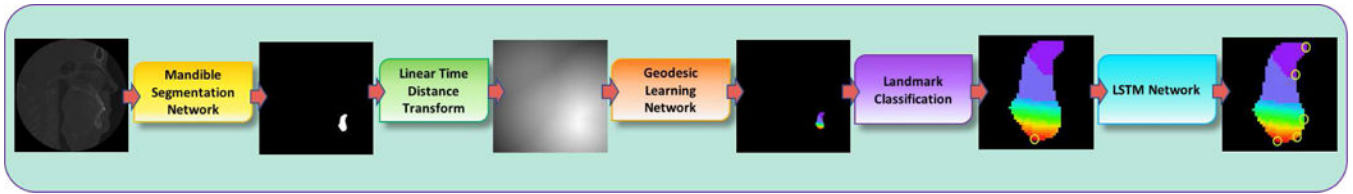


**Fig. 2:**

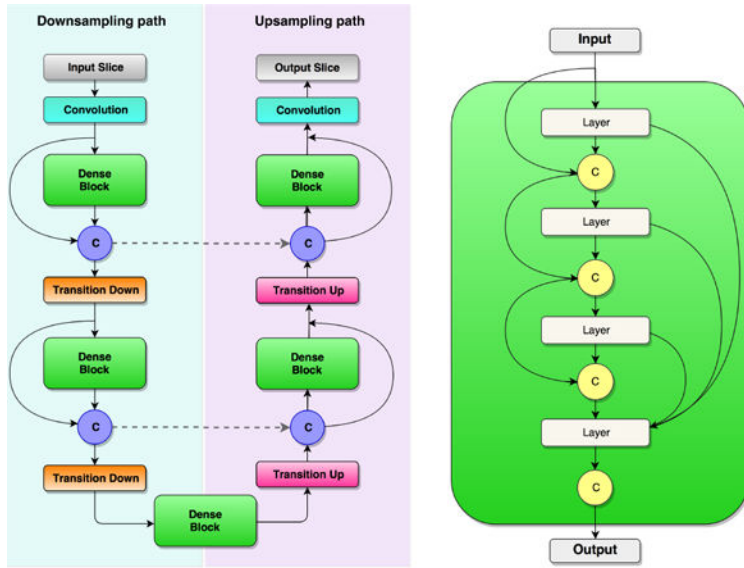
Examples of diverse CMF conditions are illustrated. (A) Surgical treatment, genioplasty with resultant chin advancement and fixation plate (implant) (adult), (B) missing condyle-ramus unit in the mandible in left dominant hemifacial microsomia (adult), (C) unerupted teeth in the anterior mandible with distorted anatomy (pediatric), (D) mid-sagittal plane with respect to lower jaw incisors have a serious degradation from the 90 degrees (pediatric), (E) bilateral bicortical positional screws (implants) in the ascending ramus of the mandible for rigid fixation after a bilateral sagittal split osteotomy (adult), (F) plate and screws (implants) in the anterior mandible for rigid fixation and reduction of an oblique fracture (adult).



**Fig. 3:** Following the Mandible Segmentation with Fully Convolutional DenseNet, Linear Time Distance Transform (LTDT) of the Mandible Bone is generated. A second U-Net [8] is used to transform LTDTs into a combined Geodesic Map of the mandibular landmarks Menton ( $Me$ ), Condylar Left ( $Cd_L$ ), Condylar Right ( $Cd_R$ ), Coronoid Left ( $Cd_L$ ), and Coronoid Right ( $Cd_R$ ). Finally, an LSTM Network is used to detect Infradentale ( $Id$ ), B point ( $B$ ), Pogonion ( $Pg$ ), and Gnathion ( $Gn$ ) mandibular landmarks according to the detected position of the Menton ( $Me$ ) in previous step. All algorithms in this proposed pipeline run in pseudo-3D (slice-by-slice 2D). To ease understanding of the segmentation results, surface rendered volumes are presented instead of contour based binary images.



**Fig. 4:**  
The example workflow of a single slice in the proposed pipeline (Figure 3). The outputs of the steps (Landmark Classification and LSTM Network) are zoomed in for visual illustration of the process.

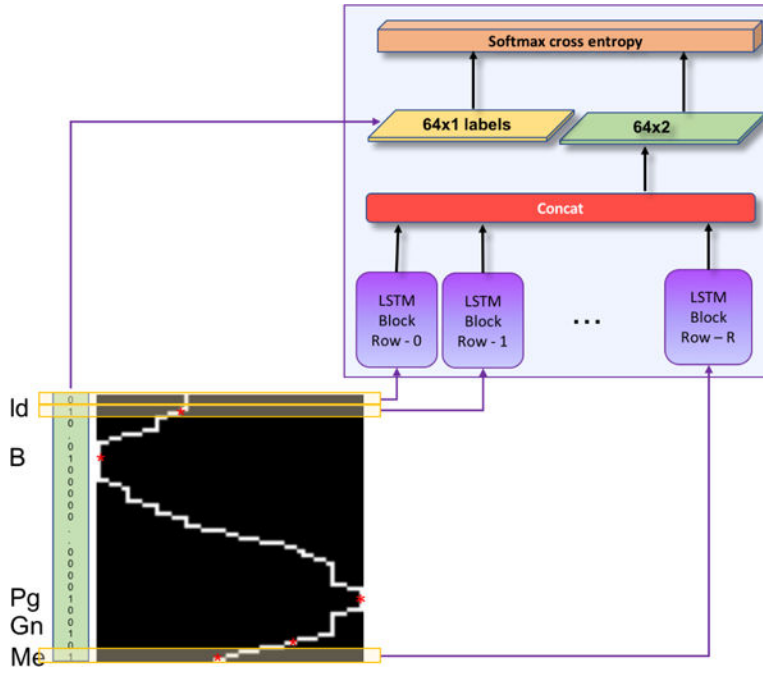


(a) Fully Convolutional DenseNet with 103 layers. (b) Content of a dense block.

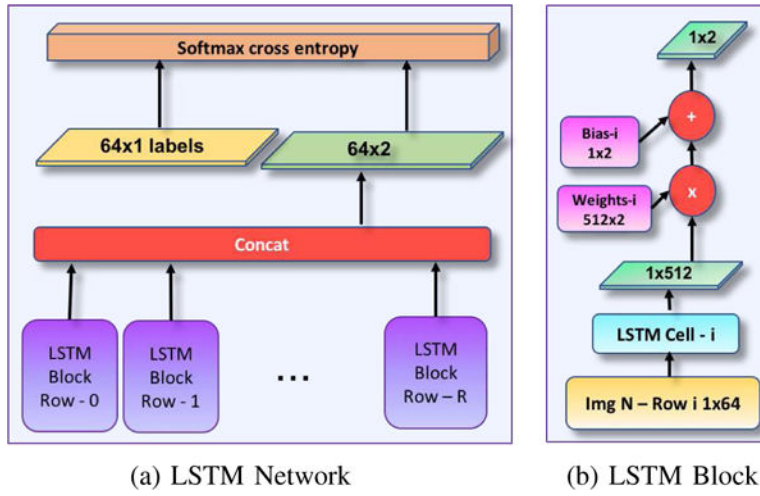
**Fig. 5:**

(a) General architecture of the Tiramisu [17] is illustrated. The architecture is composed of downsampling and upsampling paths including Convolution, Dense Block, Concatenation (C), Skip Connection (dashed lines), Transition Down, and Transition Up layers. Concatenation layer appends the input of the dense block layer to the output of it. Skip connection copies the concatenated feature maps to the upsampling path. (b) A sample dense block with 4 layers is shown to its connections. With a growth rate of  $k$ , each layer in dense block appends  $k$  feature maps to the input. Hence, the output contains  $4 \times k$  features maps.

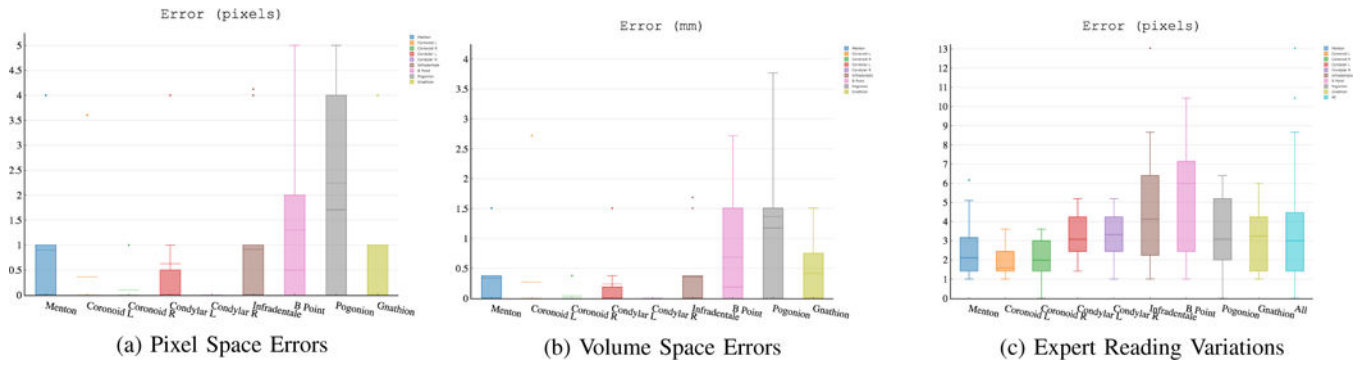




**Fig. 6:** LSTM network input-output. Each row of the scaled sagittal boundary image is input to the corresponding LSTM block, and binary 1D vector of locations annotated as landmark (1), or no-landmark (0) is output.



**Fig. 7:** Details of the network architecture (LSTM) for identifying closely-spaced landmarks. Gnathion ( $Gn$ ), Pogonion ( $Pg$ ), B Point ( $B$ ), and Infradentale ( $Id$ ) are determined once the Menton ( $Me$ ) is detected through U-Net architecture as shown in Step 3 of the Figure 3. Input image resolution is  $R \times K$ , and the LSTM cell is composed of 512 hidden units.



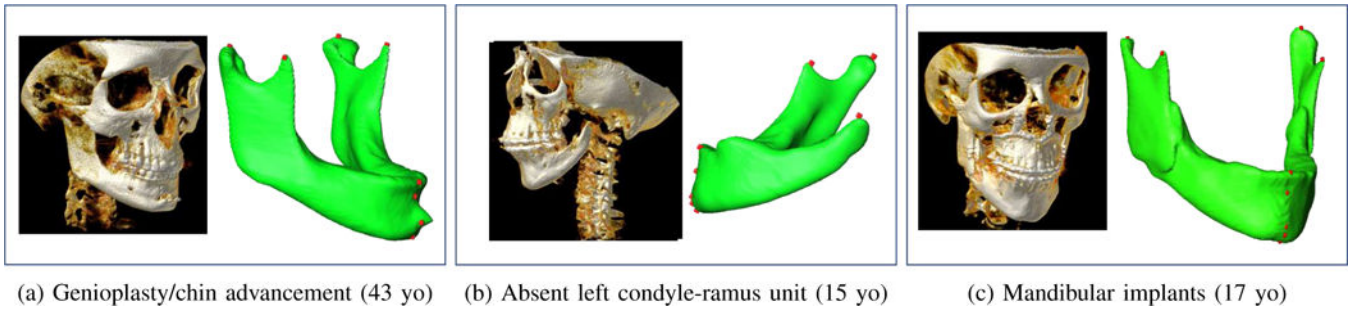
**Fig. 8:**  
 (a) Errors in pixel space, (b) errors in the volume space, (c) inter-observer reading variations in pixel space.

Author Manuscript

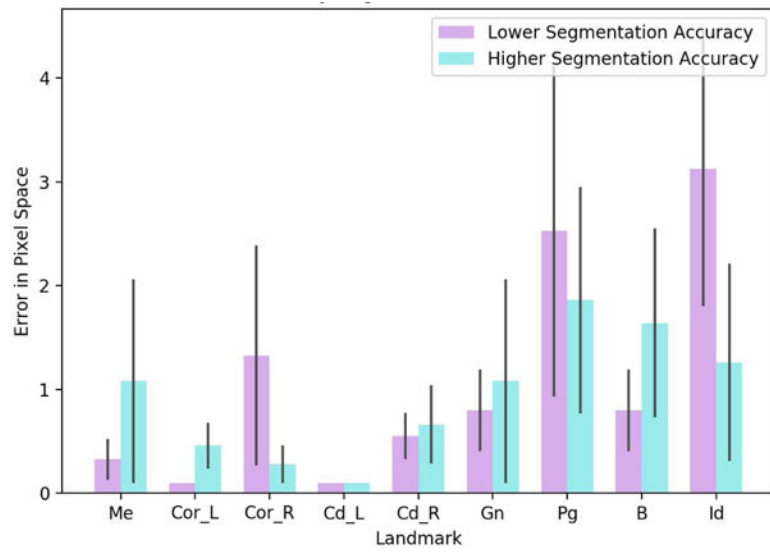
Author Manuscript

Author Manuscript

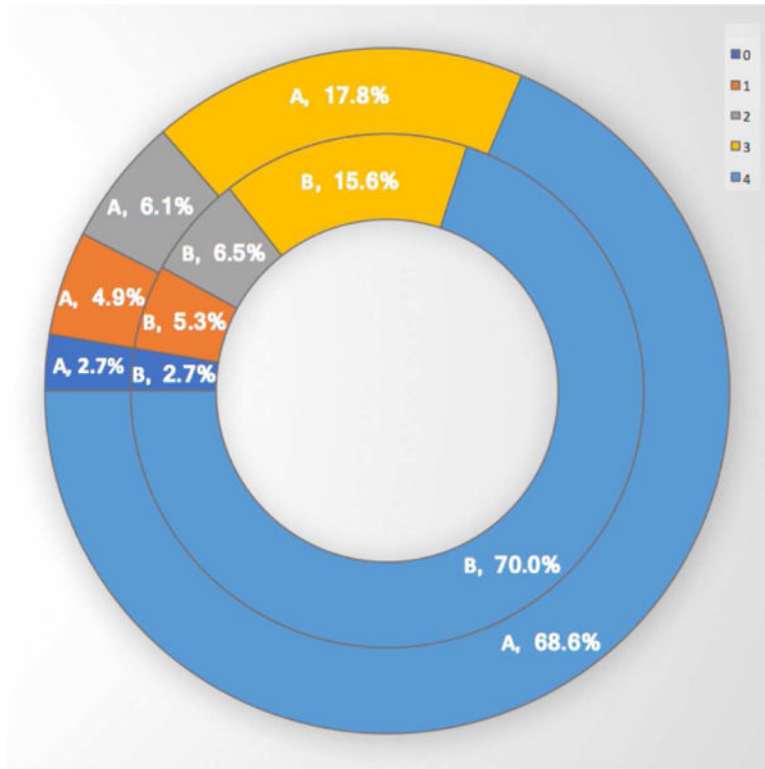
Author Manuscript



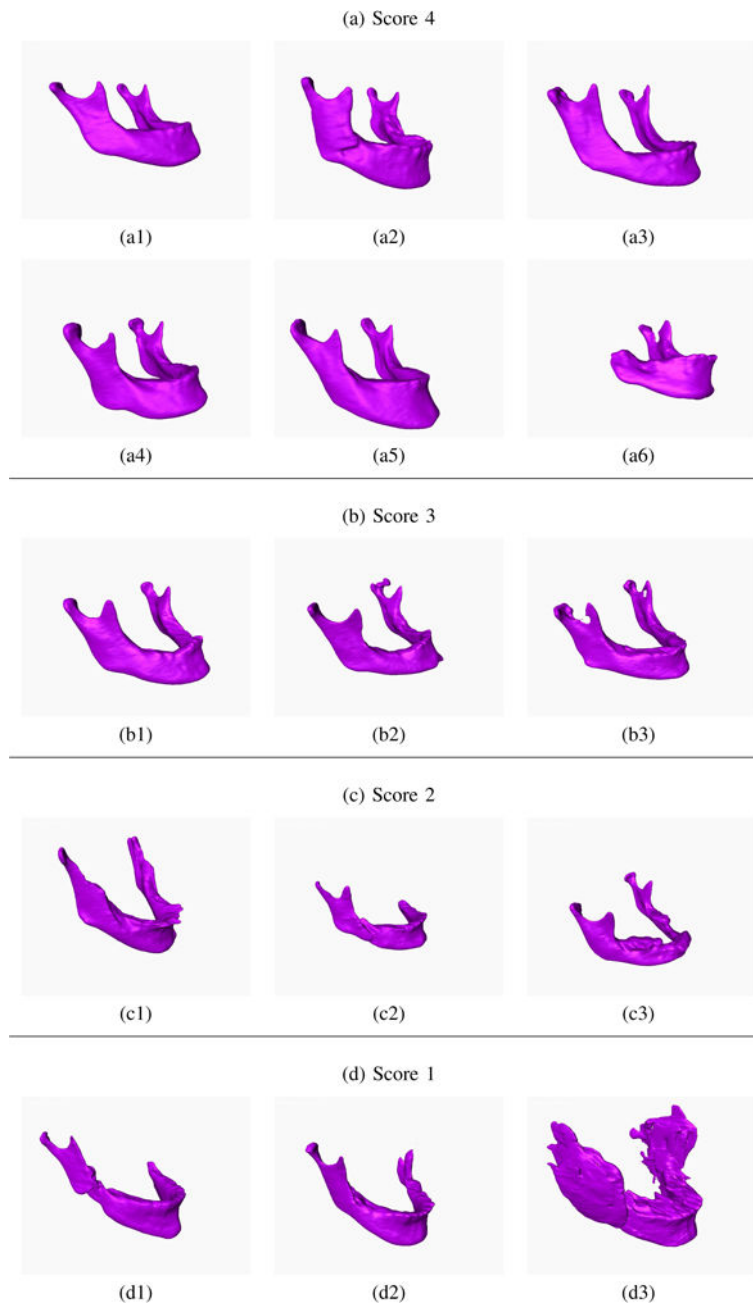
**Fig. 9:**  
Experimental renderings demonstrating segmentation and landmark localization results of patients with high anatomical variability due to deformities and surgical intervention



**Fig. 10:** Impact of segmentation accuracy on the landmark localization process.



**Fig. 11:** Summary of qualitative evaluation of 250 scans from the NIDCR/NIH dataset evaluated by 2 experts, **A** and **B**



**Fig. 12:** Qualitative Evaluation Scores of the Segmentation Results. The experts visually evaluated the performance of the segmentation of the 250 patient scans in the score range 1 to 4, where 1 is inferior. Examples of scans with scores 1–4 are presented.



**TABLE I:**

The network architecture of the Tiramisu segmentation engine.

Layers applied	# of feature maps
Input	1
$3 \times 3$ Convolution	48
Dense Block (4 layers) + Transition Down	112
Dense Block (5 layers) + Transition Down	192
Dense Block (7 layers) + Transition Down	304
Dense Block (10 layers) + Transition Down	464
Dense Block (12 layers) + Transition Down	656
Dense Block (15 layers)	896
Transition Up + Dense Block (12 layers)	1088
Transition Up + Dense Block (10 layers)	816
Transition Up + Dense Block (7 layers)	578
Transition Up + Dense Block (5 layers)	384
Transition Up + Dense Block (4 layers)	256
$1 \times 1$ Convolution	2
Softmax	2

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

**TABLE II:**

The network architecture parameters of the Tiramisu segmentation engine

<b>Hyper-Parameters</b>	<b>Value</b>
Learning-Rate	0.00005
Drop-out	0.2
Network Weight Initialization	Xavier Initializer
Bias Initializer	Zero Initializer
Activation Function	ReLU
Growth Rate	24
Normalization	Batch Normalization

<b>Network Parameters</b>	<b>Value</b>
Pooling	Average
Batch-Size	3
Optimization	Adam

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

**TABLE III:**

The network architecture of the improved Zhang's U-Net for sparsely-spaced landmarks

Layers applied	Slice Size	Number of feature maps
Input	$256 \times 256$	1
$5 \times 5$ Convolution	$256 \times 256$	32
$5 \times 5$ Convolution	$256 \times 256$	32
Max-pooling	$128 \times 128$	32
$5 \times 5$ Convolution	$128 \times 128$	64
$5 \times 5$ Convolution	$128 \times 128$	64
Max-pooling	$64 \times 64$	64
$5 \times 5$ Convolution	$64 \times 64$	128
$5 \times 5$ Deconvolution	$64 \times 64$	64
Upsampling + Copy	$128 \times 128$	128
$5 \times 5$ Deconvolution	$128 \times 128$	64
$5 \times 5$ Deconvolution	$128 \times 128$	32
Upsampling + Copy	$256 \times 256$	64
$5 \times 5$ Deconvolution	$256 \times 256$	32
$5 \times 5$ Deconvolution	$256 \times 256$	32
$5 \times 5$ Deconvolution	$256 \times 256$	21
Softmax	$256 \times 256$	21

**TABLE IV:**

The network architecture parameters of the improved Zhang's U-Net for sparsely-spaced landmarks

<b>Hyper-Parameters</b>	<b>Value</b>
Learning-Rate	1e-3
Decay-Rate	0.995
Drop-out	0.2
Network Weight Initialization	Xavier Initializer
Bias Initializer	Zero Initializer
Normalization	Batch Normalization
Pooling	Maxpool
Batch-Size	3
Optimization	RMSProp

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

**TABLE V:**

Evaluation of the segmentation algorithms. Higher IoU(%) and DSC (%), and lower HD (mm) indicate better segmentation performance. Improved Zhang's U-Net is built on top of Zhang's U-Net implementation [3].

Method	IoU	DSC	HD	Layers	# of params.
Baseline U-Net [8]	100	91.93	5.27	31	≈50M
<b>Improved Zhang's U-Net</b>	100	93.07	5.87	19	≈1M
<b>Proposed (Tiramisu)</b>	100	<b>93.82</b>	<b>5.47</b>	103	≈9M

**TABLE VI:**

Segmentation performances in different datasets, training paradigms (mixed vs. sequential), and post-processing algorithms.

	<b>Post-processing</b>	<b>DSC(%)</b>	<b>Sensitivity(%)</b>	<b>Specificity(%)</b>	<b>HD(mm)</b>
<b>Sequential</b> Tiramisu Segmentation MICCAI 2015	-	92.30	86.43	99.96	5.09
	connected component analysis, 3D fill	93.86	95.23	99.99	4.58
<b>Sequential</b> Tiramisu Segmentation NIH Dataset	-	92.61	93.42	99.97	8.80
	connected component analysis, 3D fill	93.82	93.42	99.97	6.36
<b>Mixed</b> Tiramisu Segmentation → U-Net Landmark Localization NIH Dataset	-	92.09	92.10	99.96	8.30
	connected component analysis, 3D fill	92.28	92.10	99.96	7.11
Mixed Tiramisu Segmentation → Tiramisu Landmark Localization NIH Dataset	-	90.10	90.53	99.97	8.80
	Connected component analysis, 3D fill	90.10	90.52	99.97	6.36

**TABLE VII:**

Landmark localization performances are evaluated for each anatomical landmark and with respect to different pooling functions. Errors (in mm) are given both in average (avg) and median (md) values.

		max pool	avg pool	stoc. pool	max pool + wo drop out
<i>Me</i>	avg	0.33	1.35	0.37	0.03
	md	0	0	0	0
<i>Cor<sub>L</sub></i>	avg	0.27	0.07	0	0
	md	0	0	0	0
<i>Cor<sub>R</sub></i>	avg	0.33	0.3	0.37	0.45
	md	0	0	0	0
<i>Cd<sub>L</sub></i>	avg	1.01	0.037	0.56	0.33
	md	0	0	0	0
<i>Cd<sub>R</sub></i>	avg	0	0.11	0.07	0.07
	md	0	0	0	0
<i>Gn</i>	avg	0.41	1.64	1.35	0.49
	md	0	0	0.18	0
<i>P<sub>g</sub></i>	avg	1.36	2.34	2.4	1.54
	md	1.17	0.75	1.6	0.75
<i>B</i>	avg	0.68	1.47	1.24	0.33
	md	0.18	0	0.56	0
<i>Id</i>	avg	0.35	1.74	0.75	0.52
	md	0	1.131	1.67	0



TABLE VIII:

Comparison of segmentation accuracies with respect to different regularization choices. Drop ratio of 0 denotes “no” use of dropout layer.

Pooling Activation Growth Rate	max pool		avg pool		stoc. pool		avg pool	
	ReLU 16	ReLU 24	ReLU 16	ReLU 24	ReLU 16	ReLU 24	SWISH 16	SWISH 24
drop ratio=0.0	93.09	92.64	93.16	92.16	92.59	90.93	93.14	92.60
drop ratio=0.2	93.10	93.08	93.36	<b>93.82</b>	92.14	92.53	91.19	93.61

**TABLE IX:**

Resulting segmentation DSC accuracies with respect to the drop ratio (avg pooling, ReLU, and growth rate of 24). Note that drop ratio of 0 denotes “no” use of dropout layer.

Drop Ratio	0.5	0.3	0.2	0.1	0.0
DSC(%)	91.21	93.31	<b>93.82</b>	92.90	92.88

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

**TABLE X:**

Effect of different growth rates on segmentation performance using Tiramisu with avg. pooling.

Growth Rate( $k$ )	12	16	24	32
DSC(%)	92.63	93.36	<b>93.82</b>	92.60
HD(mm)	6.44	5.50	5.41	<b>5.02</b>

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript