Novel data mining-based age-at-death estimation model using adult pubic symphysis 3D scans

Zdeněk Buk[1], Michal Štepanovský[1], Anežka Kotěrová[2], Jana Velemínská[2], Jaroslav Brůžek[2], and Pavel Kordík[1]

[1] Faculty of Information Technology, Czech Technical University in Prague,
Thakurova 9, Prague 160 00, Czech Republic
[2] Department of Anthropology and Human Genetics, Faculty of Science, Charles University,
Vinicna 7, Prague 128 43, Czech Republic

Abstract: The paper introduces a novel age-at-death estimation model based on Convolutional Neural Network (CNN). The model uses 3D scan of human pubic symphysis as an input and estimates the age-at-death of the individual as an output. The Mean Absolute Error (MAE) of this model is about 10.6 years for individuals between 18 and 92 years of age-at-death. Moreover, the results of the study indicate that pubic symphysis can be used to estimate the age of individuals across the entire age range. The study involved a sample of 483 bone scans collected from 374 individuals (from which 109 individuals provided both left and right pubic symphysis).

1 Introduction

Estimating the age of death of unknown human skeletal remains represents one of the major tasks of biological anthropologists. Traditionally, the estimation is performed visually by assessing degenerative changes of join surfaces (e.g. [8], [10]), among which the pubic symphysis of the pelvis is widely used (e.g. [4], [11]). Figure 1 illustrates a few examples of human symphyseal surfaces of individuals with the age-at-death of 25, 35, 45, 65 and 85 years, respectively. Visual observation, however, has its limitations, e.g. it is subjective, dependent on observer experience, and last but not least, its applicability suffers from low accuracy and reliability of estimates (e.g. [7], [9]). To achieve both accurate and reliable age estimates, it is recommended to use three broad intervals [1], [5]. Moreover, single-indicator methods do not work equally well throughout the adult period, for example, it has been reported that the pubic symphysis is no longer suitable for age estimation after the age of 40 years [2], [6]. Currently, the research has shifted to imaging technologies and sophisticated data mining methods (e.g. [3], [15]) that could offer a more objective and accurate perspective on age estimation in adults. The Algee-Hewitt – Slice – Stoyanova team ([12], [13] proposed the most prominent approach [14]) with the estimation error (RMSE) ranging between 13.7 and 16.6 years (based on the dataset consisting of 93 samples) [13]. In this paper, we did not follow that approach. Instead, we developed a novel age-at-death estimation model based on CNN. Our model takes an image representing the 3D structure of the bone (see the following sections) and predicts an individual's age-at-death using the CNN for pattern recognition.

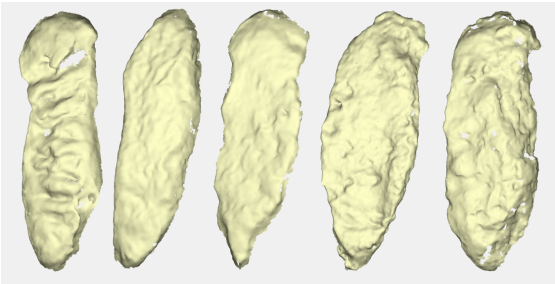

Figure 1: From left to right, examples of symphyseal surface scans of individuals with the age-at-death of 25, 35, 45, 65 and 85 years.

2 Input data

Table 1 shows the age-at-death distribution of our collection. The mean age is 53.7 years, and the standard deviation is 17.1 years. Table 2 shows the structure of the osteological collection.

Table 1: Age distribution of the collection

Age-at-death:	18–29	30–39	40–49	50–59	60–69	70+
Males	24	54	57	56	59	49
Females	10	32	33	39	21	49

Table 2: Structure of the osteological collection

Country:	Portugal	Switzerland	Thailand	Crete
Males	129	45	114	10
Females	91	21	68	5

The input dataset consists of 483 skeletal samples from adult (18–92 years) males and females. All skeletal samples were digitised using the HP 3D Structured Light Scanner PRO S2 or S3 scanner and exported in STL format. The STL format is a file format describing an unstructured triangulated surface using a 3D Cartesian coordinate system. In this format, the surface geometry of a 3D object is represented as a number of small adjacent triangles. Figure 2 shows an example of the 3D scan.

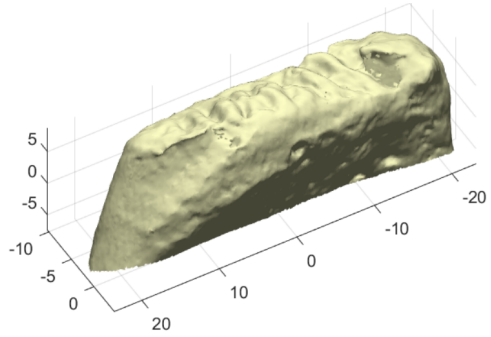

Figure 2: Example of the 3D scan of the symphysis pubica of a 25-year-old female

3 Data preprocessing

The input data coming from the scanner needs to be preprocessed before applying the proposed age estimation method. The STL format is a very convenient input data format, however, it is relatively unsuitable for direct use due to the irregular distribution of all vertices in the 3D space. There are several options how the input data could be represented and regularised. For instance, one can use voxel representation, where each voxel encodes one bit of information – the presence of the bone. However, for a typical scan dimension $50\times15\times15$ mm (see Figure 2) with the resolution of 0.1 mm, this results in over 11 million of voxels per one 3D scan, i.e. in about 11.3 Mbit of information. On the other hand, one can use only the top view of the symphysis surface (see Figure 3), since this area has the highest age prediction capabilities [4], [11]. The surface height could be encoded in colour or in grayscale. Using 8-bit grayscale gives even better resolution for the surface height, compared to the previously described voxel representation. This approach reduces the overall size to 0.6 Mbit, while keeping the same resolution for the other two dimensions. However, this "top view" representation ignores the side walls of the scan, and therefore, eliminates potentially additional age-related information.

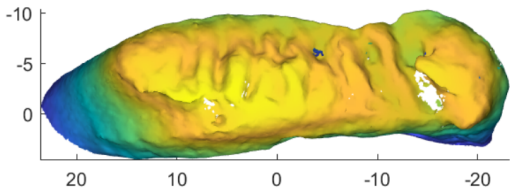

Figure 3: Top view on the symphysis surface from Figure 2. The high of the surface is encoded in colour.

Therefore, we decided to transform the input data from the Cartesian coordinate system to a new coordinate system in such a way that the side walls of the scan could be examined in a similar way to the "top view" representation. First, the position, size, and orientation of all scans needs to be standardised. Then, every point (x, y, z) in the Cartesian coordinate system is transformed to a point (λ, μ, ν) in the new coordinate system using the following equations

$$\begin{aligned} \lambda &= \varphi cos(\theta), \\ \mu &= \varphi sin(\theta), \\ \nu &= z, \end{aligned} \tag{1}$$

where

$$\begin{aligned} \theta &= atan(y/x), \\ \varphi &= acos(z/\rho), \\ \rho &= \sqrt{x^2+y^2+z^2}. \end{aligned} \tag{2}$$

Here, we should note that $\theta = atan(y/x)$ is the four quadrant arctangent of the elements of x and y such that $-\pi \leq atan(y/x) \leq \pi$. The concept of this coordinate system transformation is illustrated in Figure 4. Figure 4 shows some point P on the surface S, where the surface S represents the surface of the symphysis. Initially, the point P is transformed into the spherical coordinates using the equation (2). Consecutively, this point is further transformed into the point P' using the equation (1). If all points obtained from the scan are transformed in such a way, then the surface S is transformed to the new surface S'.

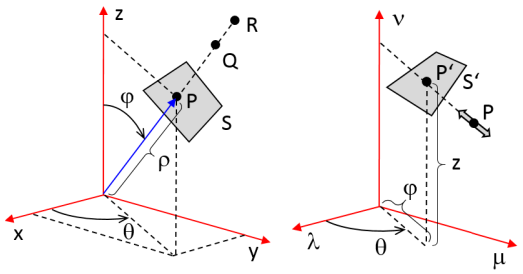

(a) First step (b) Second step

Figure 4: Transformation from the Cartesian coordinate system

The above-described transformation has several interesting properties and offers multiple advantages. Figure 5 helps to understand these properties. Figure 5 (a) shows a cube placed at the centre of the Cartesian coordinate system. Figure 5 (b) shows the same cube but in the new coordinate system. Since the transformation preserves $\nu = z$, all points have the same height above the x-y plane, or λ-μ plane, respectively. The cube is virtually stretched out from the bottom side of the cube in that way that the entire cube can be described as a function of the two variables λ and μ, i.e., for each point (λ, μ) in a portion of the λ-μ plane (the domain of the function) we can assign a unique number $f(\lambda, \mu)$. This is very advantageous, since the complicated 3D shape can be transformed to a 2D image practically without the loss of information.

Another advantage of the proposed transformation is a consequence of preserving $\nu = z$. As already mentioned, all points have the same height above the x-y plane, or λ-μ plane, respectively. This allows us to detect and analyse the disturbances in the symphyseal surface profile quite

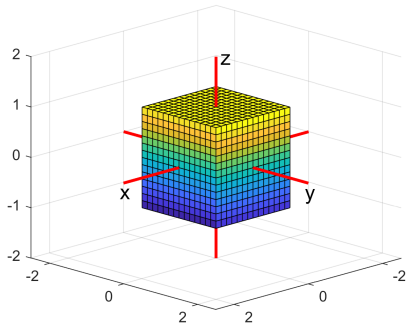

(a) Cube in the Cartesian coordinate system

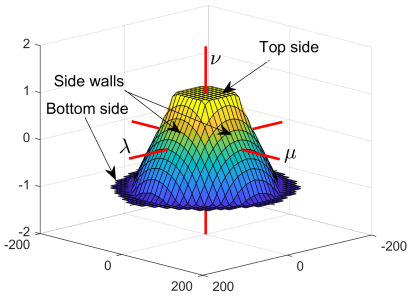

(b) Cube in the proposed coordinate system

Figure 5: Example of the proposed coordinate system transformation

easily. This transformation, however, also has few drawbacks. First, it does not preserve the global shape of the surface as seen from a perpendicular view to that surface. For instance, perfectly square side-walls of a cube (Figure 5 (a)) become increasingly stretched out as φ grows. This can be seen in Figure 5 (b), where the bottom edge of the "square" is much wider compared to the upper edge of that "square". Moreover, the bottom side of the original cube is completely deformed and rather resembles a ring, as seen in Figures 5 (b) in the dark blue areas of the image. However, this disadvantage is of little significance for our purposes, since all pubic scans have no bottom (3D scan captures only the surface of the bone, not internal parts of the bone), and the coordinate system is located in such a way that the most important areas of the scan are deformed only slightly. Figure 6 shows the distribution of a set of points in the λ-μ plane, which were originally uniformly placed on the surface of the cube. This helps to visualise the deformation of the cube shape. The top side of the cube is located around the origin of the λ-μ plane. The edges of the front side of the cube are highlighted in green. As can be seen, the bottom edge of the cube is more stretched compared to the top edge. The second disadvantage of the proposed coordinate system is more fundamental for complicated shapes, as 3D scans can be. Namely, not for all shapes, we can assign a unique number $f(\lambda,\mu)$ in the λ-μ plane. All points with the same value of θ and φ (for instance, points P, Q and R in Figure 4 (a)) are projected into the same (λ, μ) coordinates. In this case, we can select the maximum, minimum, median or the average of all points mapped to the same (λ,μ) point. In this situation, some information from the original shape is lost and cannot be fully recovered anymore. This creates unwanted artefacts in the transformed data. However, we have experimentally observed that it occurs only occasionally for our dataset and affects only small portions of the whole area. In our case, these artefacts are partially suppressed by scaling down all x-coordinates by a factor of 2.5 before applying the above described transformation of the coordinate system.

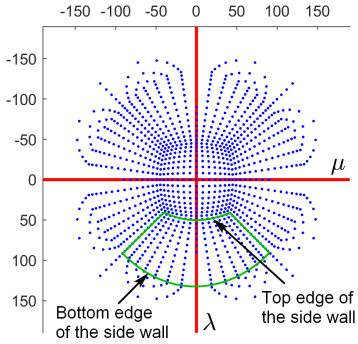

Figure 6: Visualisation of a set of points originally uniformly located on the cube surface in the Cartesian coordinate system after the projection into the λ-μ plane.

Figure 7 shows the symphysis surface from Figure 2 in the new coordinate system. The φ variable is plotted with a resolution of 2° for better visualisation. The actual resolution is set to 0.5°. The surface from Figure 7 is projected onto a regular mesh in the λ-μ plane, where the ν coordinate is encoded in 8-bit (or 16-bit) value, effectively creating a grayscale image as shown in Figure 8. The range of angle φ can be arbitrarily chosen, e.g., if chosen such that $\varphi \in < 0^\circ, 90^\circ >$, then only points above the x-y plane (with a positive z value) are used. The grayscale image can be directly used as input to the age estimation model. Moreover, to increase the variability of the input training dataset and the robustness of our model, we have generated 41 projections (grayscale images) for each 3D scan with a slightly rotated and translated origin of the Cartesian coordinate system.

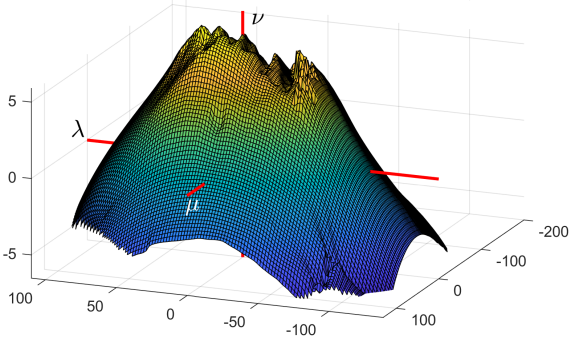

Figure 7: Symphyseal surface fom Figure 2 in the proposed coordinate system with the resolution of 2° of φ

Figure 8: Generated grayscale image

4 Age estimation model

Our age estimation model consists of several identical age predictors. Each predictor is based on convolutional neural network [16]. Figure 9 shows the main idea of the age estimation flow. First, the 3D scan is transformed into several grayscale images (see Figure 8 as an example). These images are consecutively directly used as input for individuals age predictors. Second, an aggregation function is applied in order to combine the results from all predictors, and thus, to provide the final prediction. We have chosen *mean* and *median* as two possible aggregation functions.

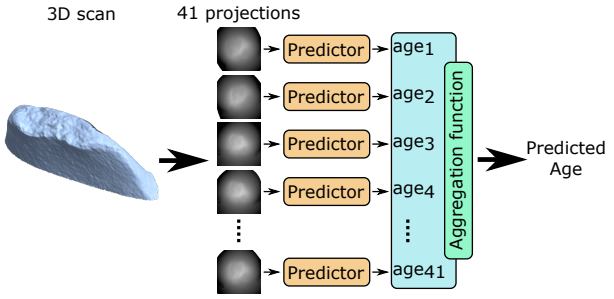

Figure 9: Age prediction for a single individual. There is a single 3D scan for which we build multiple projections (41 in this case). By applying the predictor for each of the projections, we obtain multiple age predictions that are finally aggregated to gain the final predicted age.

4.1 Predictor structure

The application of convolutional neural networks was an obvious choice. We have experimented with various architectures. In the final experiments, we were mainly inspired by the setup used for X-Ray images processing [17]. We also experimented with topologies based on DenseNet [18, 19], which exploits a specific topology that shortens layer connections by connecting each layer to every other layer in a feed-forward fashion. The final predictor structure is shown in Figure 10. The model consists of a total of 20 layers. In the first part, the input image is reduced and transformed into features using convolutional layers in combination with pooling, activation (Ramp) and regularisation layers. The second part represents densely connected feedforward networks that transform 1024 features into a single real value. For better generalisation, we used a dropout layer in the second part.

NetChain

	Input	image array (size: 1×160×160)
Augm	ImageAugmentationLayer	array (size: 1×128×128)
Model	NetChain (20 nodes)	vector (size: 1)
	Output	scalar

Model: NetChain

	Input	array (size: 1×128×128)
1	ConvolutionLayer	array (size: 8×122×122)
2	BatchNormalizationLayer	array (size: 8×122×122)
3	Ramp	array (size: 8×122×122)
4	ConvolutionLayer	array (size: 16×116×116)
5	BatchNormalizationLayer	array (size: 16×116×116)
6	Ramp	array (size: 16×116×116)
7	PoolingLayer	array (size: 16×58×58)
8	ConvolutionLayer	array (size: 16×56×56)
9	BatchNormalizationLayer	array (size: 16×56×56)
10	Ramp	array (size: 16×56×56)
11	PoolingLayer	array (size: 16×28×28)
12	ConvolutionLayer	array (size: 16×26×26)
13	BatchNormalizationLayer	array (size: 16×26×26)
14	Ramp	array (size: 16×26×26)
15	PoolingLayer	array (size: 16×8×8)
16	FlattenLayer	vector (size: 1024)
17	DropoutLayer	vector (size: 1024)
18	LinearLayer	vector (size: 100)
19	Ramp	vector (size: 100)
20	LinearLayer	vector (size: 1)
	Output	vector (size: 1)

Figure 10: Age predictor. The network is wrapped with an Image Augmentation layer, which implements random transformations during training.

4.2 Implementation details

The model was implemented in Wolfram Mathematica using an in-house neural networks package built on the MXNet framework. To improve the robustness of the network training, we incorporated an `ImageAugmentationLayer` to the input layer. This layer takes the input image 160x160 pixels and randomly crops it to 128x128 pixels during the training phase – this allows us to efficiently expand the training dataset without having to implement a custom batch function. During the evaluation phase, the `ImageAugmentationLayer` crops the input image around the centre in a deterministic way, so it does not affect it during evaluation. We have chosen slightly larger kernels (7x7) in the first two layers to better handle larger structures in blurred images. We use the batch normalisation layers which are proposed as a technique to help coordinate the updating of multiple layers in the model [16]. Figures 11 and 12 are used as examples illustrating the extracted patterns for selected layers of the network for 20- and 72-years-old individuals, respectively. These figures show the input image and the output from layers #3, #6, #10 and #19. As it can be seen, the model identifies the vertical structures and the edge of the symphyseal surface reasonably well from the input image. These vertical structures combined with the shape of the symphyseal edge are also used by experts to identify the age of the individual.

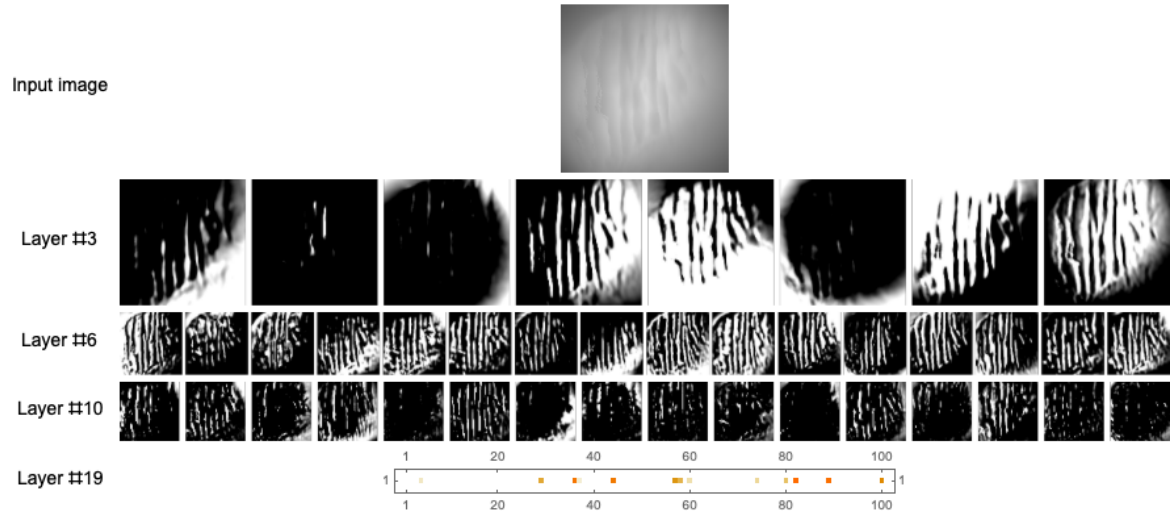

Figure 11: An example of a 20-years-old individual (predicted age = 20.82) evaluation – the surface structure can be clearly recognised.

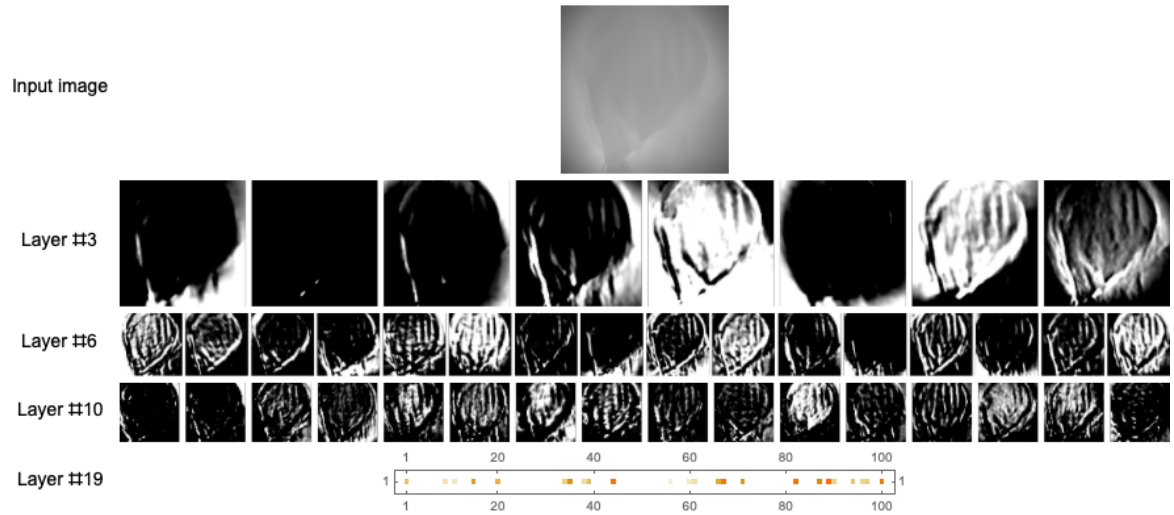

Figure 12: An example of a 72-years-old individual (predicted age = 71.6) evaluation. Even though we can see almost no details in the input image, the model can identify vertical (in this orientation) structures, which seems to be a key for the age identification.

4.3 Training

As already mentioned, we have multiple images for a single individual (41 images per bone). For some individuals, we have at our disposal left and right bones, so there are 82 images for a given individual. All models are designed to process a single image on input, so the training set is an unstructured list of pairs $\{image, age\}$. Since we need proper testing, all images for a single individual must be in the same fold. Cross-validation uses the information about the ID of the individual to split the dataset properly. The folds are then flattened, and we perform standard training for mapping images to real values (age).

We use a standard Adam optimiser [20] with a batch size 16 (experimentally chosen) running on the GPU for training. Each training runs approximately 300 rounds (we have performed many experiments from 100 to many thousands), representing almost 10 million processed records.

4.4 Evaluation

To compare the models and study the behaviour, we have performed two ways of model evaluation. Obviously, the main goal is to predict the age of an (unknown) individual based on the 3D scan of the bone, so we evaluated the model for each individual (41 or 82 images) and computed the *predicted age* using aggregation. Figure 13 shows the actual age vs. the predicted age per individual. Based on the size of the dataset, we have chosen 5-fold validation for all of our experiments.

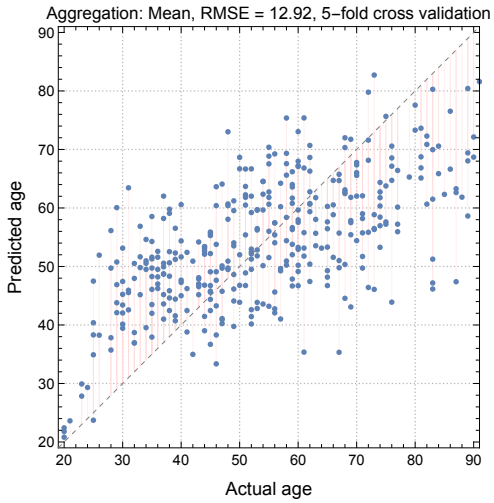

Figure 13: Age prediction with mean aggregation function. Each dot represents a single individual.

5 Results and discussion

Stoyanova et al. presented in their study [13] five age-estimation models with similar age estimation performance and provided an open source software called forAGE (available at http://morphlab.sc.fsu.edu/). The best model (according to their study) is "SAH&VC" (SAH+Outline) and provides MAE of approximately 10.79 for their entire dataset (93 samples). To compare our model with their state-of-the-art model, we used their software and evaluated the MEA for our dataset as well. The results are summarised in Table 3. It should be noted here that for our models, MAE and RMSE are computed using 5-fold cross-validation, whereas for the model from [13] MAE and RMSE are computed directly.

Table 3: Age prediction results. The table compares MAE and RMSE of our models with the model designed by Stoyanova et al. [13] used on our dataset.

Age estimation model	MAE	RMSE
Our model (Median)	10.63	12.94
Our model (Mean)	**10.60**	**12.92**
SAH&VC	19.52	25.04

As the presented results indicate, our models outperform the model developed by Stoyanova et al. [13] in terms of both MAE and RMSE. There is a significant discrepancy between the MAE presented by Stoyanova et al. and the MAE computed on our dataset (i.e., 10.79 vs. 19.52 years). This discrepancy is discussed below in the text.

To determine whether our model contains any systematic error or whether any particular age intervals introduce some anomalies in the prediction, we processed the predicted ages per one-year age intervals. Figure 14 shows the variation in age predictions for each age class. As can be seen, our model generally overestimates younger individuals (under the age of 55 years) and underestimates mature ones (above the age of 55 years). This is a result of the tendency to predict the age towards the mean age of the sample.

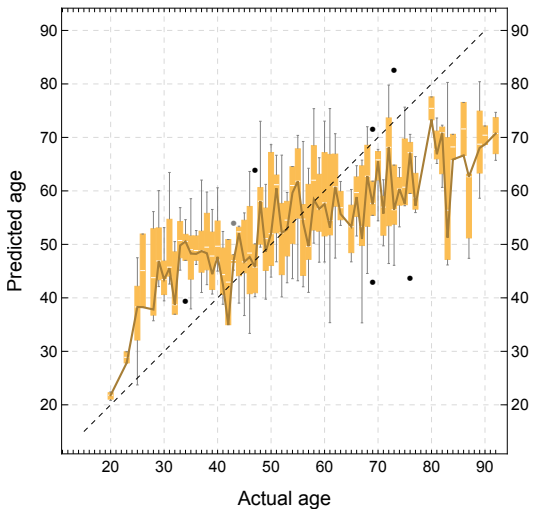

Figure 14: Variation of age predictions for particular age class (all individuals mixed) for our model. The image shows the minimum, maximum, q_1, q_2 (median) and $q3$ quartiles. Black dots represent outliers defined by quartiles and 1.5× interquartile range. The central line connects the median values. Aggregation function: *Mean*

Similarly, we analysed the SAH&VC model from [13]. Figure 15 shows the variation of age predictions for particular age class. As can be seen, the SAH&VC model generally underestimates all individuals above 40-50 years. More specifically, for individuals over 45 years old, the average of all estimations reaches only 37.7 years (for our dataset). We believe that this primarily results from the unbalanced age distribution of the dataset used in [13].

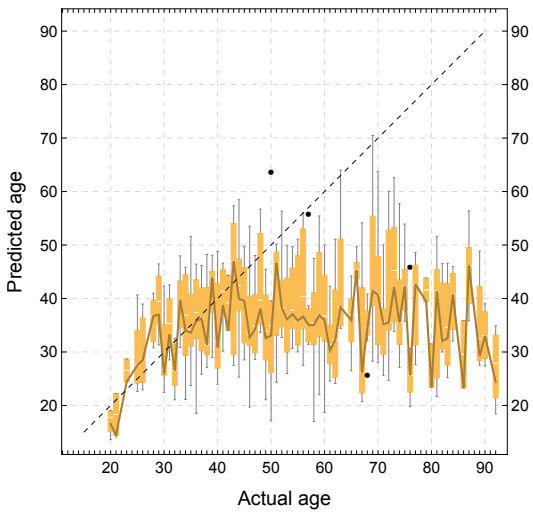

Figure 15: Variation of age predictions for particular age class according to SAH&VC model by Stoyanova [13].

Moreover, we observed that our model can estimate the age-at-death of an individual over the entire age interval (in our case between 19–92 years) – see Figure 14. This contrasts with e.g. [2] and [21], where pubic symphysis is considered appropriate for individuals up to 40 years, or 60 years, respectively. When the maturation process of pubic symphysis is complete, the morphological changes are degenerative and highly variable between individuals [2], [22]. However, we believe that our model can capture even such changes.

6 Conclusion

We have developed a novel age-at-death estimation model based on convolution neural networks. Our model provides a mean absolute error of approximately 10.6 years and is suitable for adult and mature individuals. Our results indicate that the pubic symphysis reflects the age of an individual throughout their entire adult life. In other words, we have observed no limitations in terms of age prediction capabilities of pubic symphysis of adult individuals.

Acknowledgments This research was supported by a research grant awarded by the Technology Agency of the Czech Republic; project number TL03000646.

References

[1] Adserias-Garriga, & Wilson-Taylor, R. (2019). Skeletal age estimation in adults. In J. Adserias-Garriga (Ed.), Age Estimation: A Multidisciplinary Approach (First, pp. 55–73). London: Academic Press.

[2] Baccino, Sinfield, L., Colomb, S., Baum, T. P., & Martrille, L. (2014). Technical note: The two step procedure (TSP) for the determination of age at death of adult human remains in forensic cases. Forensic Science Int., 244, 247–251.

[3] Biwasaka, H., Sato, K., Aoki, Y., et al. (2013). Three dimensional surface analyses of pubic symphyseal faces of contemporary Japanese reconstructed with 3D digitized scanner. Legal Medicine, 15, 264–268.

[4] Brooks, S., & Suchey, J. M. (1990). Skeletal age determination based on the os pubis: a comparison of the Acsádi-Nemeskéri and Suchey-Brooks methods. Human Evolution, 5, 227–238.

[5] Buk, Z., Kordík, P., Brůžek, J., Schmitt, A., & Šnorek, M. (2012). The age at death assessment in a multi-ethnic sample of pelvic bones using nature-inspired data mining methods. Forensic Science International, 220, 294.e1-e9.

[6] Dudzik, B., & Langley, N. R. (2015). Estimating age from the pubic symphysis: A new component-based system. Forensic Science International, 257, 98–105.

[7] Hartnett, K. M. (2010). Analysis of age-at-death estimation using data from a new, modern autopsy sample - Part I: Pubic bone. Journal of Forensic Sciences, 55, 1145–1151.

[8] Langley, & Tersigni-Tarrant, M. A. (2017). Forensic Anthropology: A Comprehensive Introduction (2nd edn.). Boca Raton: CRC Press.

[9] Milner, G. R., Wood, J. W., & Boldsen, J. L. (2018). Paleodemography: problems, progress, and potential. In: Biological Anthropology of the Human Skeleton (Third, pp. 593–633). New York: John Wiley & Sons.

[10] Nikita, E. (2017). Osteoarchaeology: A guide to the macroscopic study of human skeletal remains (First). London: Academic Press.

[11] Schmitt, A., Murail, P., Cunha, E., & Rougé, D. (2002). Variability of the pattern of aging on the human skeleton: Evidence from bone indicators and implication on age at death estimation. J Forensic Sci, 47, 1203–1209.

[12] Slice, D. E., & Algee-Hewitt, B. F. B. (2015). Modeling Bone Surface Morphology: A Fully Quantitative Method for Age-at-Death Estimation Using the Pubic Symphysis. Journal of Forensic Sciences, 60, 835–843.

[13] Stoyanova, Algee-Hewitt, B. F. B., Kim, J., & Slice, D. E. (2017). A Computational Framework for Age-at-Death Estimation from the Skeleton: Surface and Outline Analysis of 3D Laser Scans of the Adult Pubic Symphysis. J Forensic Sci, 62, 1434–1444.

[14] Stoyanova, Algee-Hewitt, B. F. B., & Slice, D. E. (2015). An enhanced computational method for age-at-death estimation based on the pubic symphysis using 3D laser scans and thin plate splines. American Journal of Physical Anthropology, 158, 431–440.

[15] Villa, C., Gaudio, D., Cattaneo, C., et al (2015). Surface Curvature of Pelvic Joints from Three Laser Scanners: Separating Anatomy from Measurement Error. Journal of Forensic Sciences, 60, 374–381.

[16] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. Deep Learning. MIT Press, 2016.

[17] Christie Natashia Archie. Chest X-rays Pneumonia Detection using Convolutional Neural Network, 2020, https://towardsdatascience.com/chest-x-rays-pneumonia-detection-using-convolutional-neural-network-63d6ec2d1dee

[18] Huang, Gao, Zhuang Liu, and Kilian Q. Weinberger. Densely Connected Convolutional Networks. CoRR abs/1608.06993 (2016). http://arxiv.org/abs/1608.06993.

[19] Sultana, Farhana, Abu Sufian, and Paramartha Dutta. Advancements in Image Classification Using Convolutional Neural Network. CoRR abs/1905.03288 (2019). http://arxiv.org/abs/1905.03288.

[20] Diederik P. Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization. In: 3rd International Conference for Learning Representations, San Diego, 2015. Revised 30 Jan 2017.

[21] Teixeira, F., Cunha, E. (2021). Aging the elderly: Does the skull tell us something about age at death? In B. F. B. Algee-Hewitt and J. Kim (Eds.), Remodeling Forensic Skeletal Age (pp. 75–97). Academic Press.

[22] Meindl, R.S., Lovejoy, C.O., Mensforth, R.P., WalkerR.A., A revised method of age determination using the os pubis, with a review and tests of accuracy of other current methods of pubic symphyseal aging, Am. J. Phys. Anthropol. 68 (Sep (1)) (1985) 29–45.