

Fully Automatic Teeth Segmentation in Adult OPG Images

Nicolás Vila Blanco¹(✉), Timothy F. Cootes², Claudia Lindner²,
Inmaculada Tomás Carmona³, and Maria J. Carreira¹

¹ University of Santiago de Compostela, Centro de Investigación
en Tecnoloxías da Información (CITIUS), Spain

² University of Manchester, Centre for Imaging Sciences, UK

³ University of Santiago de Compostela, Oral Sciences Research Group,
Health Research Institute Foundation of Santiago (FIDIS), Spain

`nicolas.vila@usc.es`

Abstract. This work addresses the problem of segmenting teeth in panoramic dental images. Random forest regression voting constrained local models were applied firstly to locate the mandible and the approximate pose of each tooth, and secondly to locate the full outline of each individual tooth. An automatically computed quality-of-fit measure was proposed to identify missing teeth. The system was evaluated using 346 manually annotated images containing adult-stage mandibular teeth. Encouraging results were achieved for detecting missing teeth. The system achieved state-of-the-art performance in locating the outline of present teeth with a median point-to-curve error of 0.2 mm for each of the teeth.

Keywords: Teeth segmentation, panoramic dental images, random forest regression-voting, machine learning.

1 Introduction

Dental radiographs have been widely used since the discovery of X-rays in a variety of fields: abnormality detection, treatment and/or surgery planning, prostheses design, assessment of children’s dental development, human identification by dental matching, and many more. X-ray images provide additional information to the simple exploration of the oral cavity since they reveal hidden parts of the teeth and other surrounding structures. There are several types of dental X-ray images depending on the captured area. In intraoral images, the sensor is placed inside the mouth and the images cover some specific area (no more than 3 – 4 complete teeth). In contrast, in extraoral images the sensor is placed outside the mouth and the images cover a bigger area. That is the case for panoramic images, which provide a complete coverage of the dentition and other surrounding bones and tissues with a very small dose of ionising radiation. Although their quality is highly dependent on patient positioning and patient movements during acquisition [1, 2] they have been widely used to diagnose periodontal disease, cysts

in the jaw bones, jaw tumours, oral cancer, impacted teeth, temporomandibular joint disorders or sinusitis, among others.

One of the key tasks in automatic dental image processing is teeth segmentation. This has proven to be useful in a variety of areas such as human identification [3–5], caries detection [6], lesion detection [7] or even dental age estimation [8]. The works in this area tackled automatic or semiautomatic teeth segmentation mostly from intraoral images in a variety of ways, comprising thresholding [4,9], combination of morphological operations [10,11], active contours [3], level sets [12], mixture of Gaussians [5] and many more. Although these algorithms can reach great performance in a variety of applications, they present some problems when working with dental images, mainly because they are very sensitive to intensity changes, dental restorations, teeth injuries and overlapping teeth. Thus, there is a need to follow more robust approaches which use domain knowledge to improve the results.

In this regard, methods utilising statistical models have proven to be accurate and robust in medical image segmentation. One of latest contributions on this area are random forest regression-voting constrained local models (RFRV-CLMs) [13], which are an evolution of the original constrained local models [14] and combines a global shape model with individual point appearance models. Over the last years, this approach has been applied to a variety of medical images with high performance [15–17], which encourages us to use it in the teeth segmentation problem.

Our main contribution is the development of a fully automatic procedure to outline mandibular adult-stage teeth in panoramic images, including the identification of any missing teeth.

2 Methods

2.1 RFRV-CLM

RFRV-CLMs combine a linear shape model with a set of local models designed to locate each point. RFRV-CLMs are summarised in the following, and the reader is referred to [13,15]) for full details.

Each annotated shape is encoded as a vector x with the concatenated coordinates of the n landmark points; $x = (x_0, y_0, x_1, y_1, \dots, x_{n-1}, y_{n-1})^T$. In order to train the model, the shapes are resampled and aligned in a reference frame so a linear model can be built as follows:

$$x = T_\theta(\bar{x} + Pb_i + r_i), \quad (1)$$

where \bar{x} is the mean shape, P are eigenvectors of the covariance matrix, b is a vector of shape parameters, r is a regularisation term which allows small deviations from the model and T_θ is a similarity transformation of parameters θ which maps the shape from the reference frame to the image frame.

In order to locate each individual point, random forest regression-voting is used. The region of interest which encloses all landmark points is resampled into

a standardised reference frame and for each landmark point l in x a set of image patches p_j are sampled at random displacements d_j (i.e. centred at $l + d_j$). Then a set of decision tree regressors are trained from the Haar features [18] of all patches to predict the displacements.

Given a new image and an initial estimation of the pose of the mean shape, the region of interest is resampled into a standardised reference frame and a set of image patches are sampled at random displacements around each initial estimated point. Haar features are extracted from the patches and fed into the random forest regressors. The outputs of all decision trees are accumulated in a voting grid V_l , where the positions of the grid with higher values indicate the most likely position for that landmark point.

The local appearance models and the global shape model are combined as follows:

$$Q(b, \theta, r) = \sum_{l=0}^{n-1} V_l(T_\theta(\bar{x}_l + Pb_l + r_l)), \quad (2)$$

$$\text{s.t. } b^T S_b^{-1} b \leq M_t \quad \text{and} \quad |r_l| < r_t,$$

where M_t and r_t are thresholds on the Mahalanobis distance and the regularisation term, respectively, and S_b is the covariance matrix of the shape model parameters b . This yields the overall quality-of-fit (QoF) measurement Q (2), which represents the total number of votes for a shape defined by parameters $\{b, \theta, r\}$.

The search process is carried out iteratively, so for each search iteration the algorithm gets the set of parameters $\{b, \theta, r\}$ which maximises the overall QoF and updates the landmark points.

2.2 Two-Step Teeth Segmentation

We build separate RFRV-CLMs for each tooth type. Given that the dentition is almost horizontally symmetric, a single model trained from one tooth on one side (left or right) can also be used to segment the corresponding tooth on the opposite side. It is worth mentioning that there are two main problems with teeth segmentation from individual teeth models. First of all, the space occupied by each tooth is very small when compared to the image size, so the search process requires a reasonably good initialisation. Furthermore, teeth of the same type (e.g. single-root and multi-root) are very similar to each other so the search process can easily end up converging to a neighbouring tooth.

To overcome these problems, in addition to individual teeth models, another model was trained from some keypoints in the image. The idea is to identify the most representative points in each tooth and the mandible which give a reasonably good approximation of their poses (see Tab.1). Thus, this model is able to capture the pose variation of each tooth (in terms of position, size and rotation) in relation to neighbouring teeth and the mandible. As the mandible occupies a similar percentage in all panoramic images, a good initialisation of

the search model can be carried out by placing the mean shape in the centre of the image and scaling it to the 75% of the image width.

The search process for a new image is performed fully automatically in two steps. In the first step, the keypoint model looks for the optimal localisation of the teeth and mandible keypoints. Then, the initial pose estimation of each tooth is carried out via (3):

$$\arg \min_{\theta} d(k_t, T_{\theta}(\bar{x}_k)), \quad (3)$$

where k_t is the estimation of the keypoints of tooth t provided by the first model, \bar{x}_k are the keypoints of the mean shape of tooth t and d is the Euclidean distance function. The initial shape estimation for each tooth is, therefore, the result of applying the estimated pose to the mean shape, $T_{\theta}(\bar{x})$.

On completion of the search we estimate the QoF of each model point by computing the magnitude of the mean displacement vector produced by the random forest for the point when evaluated on a patch centred on the point. This should be small for good matches and larger for those points which do not match so well. To obtain a score for the whole tooth we compute the mean, m , and standard deviation, sd , of the values for each point, and construct the final score as $QoF = m + sd$. This has been shown to be a more effective discriminator than just using the mean alone. We treat a tooth as missing if this QoF is above a threshold.

3 Experiments and Results

In this work, a set of 346 panoramic images provided by the School of Medicine and Dentistry, University of Santiago de Compostela, Spain, have been used, all of which were collected under ethical approval. To test the proposed segmentation approach, the images where one hemi-arch including all seven left-mandibular teeth (from the first incisor to the second molar) were present have been used as the train set, and the remaining images have been used as the test set. In total, 261 images have been used for training and 85 for testing. In each image the shapes of seven left-mandibular teeth (from 31 to 37) have been manually annotated as well as 7 mandible keypoints (see Fig.1 and Tab.1). In total, each training example consists of a set of 263 landmark points.

The individual tooth models and the keypoint model were built using the RFRV-CLM algorithm. The mean shape of each tooth model is shown in Fig.2. For each model, a coarse-fine approach has been followed, which in this case consists of training a fine model where the reference frame width is approximately the desired object width, and training a coarse model where the frame width is about a quarter of the fine frame width. This gives a rough but more robust shape estimation at first and then refines the shape. In the case of the keypoint model, the search process consists of 3 search iterations with the coarse model and 2 search iterations with the fine model. For the individual teeth models, the iterations of coarse and fine searches have been reduced to 2 and 1, respectively.

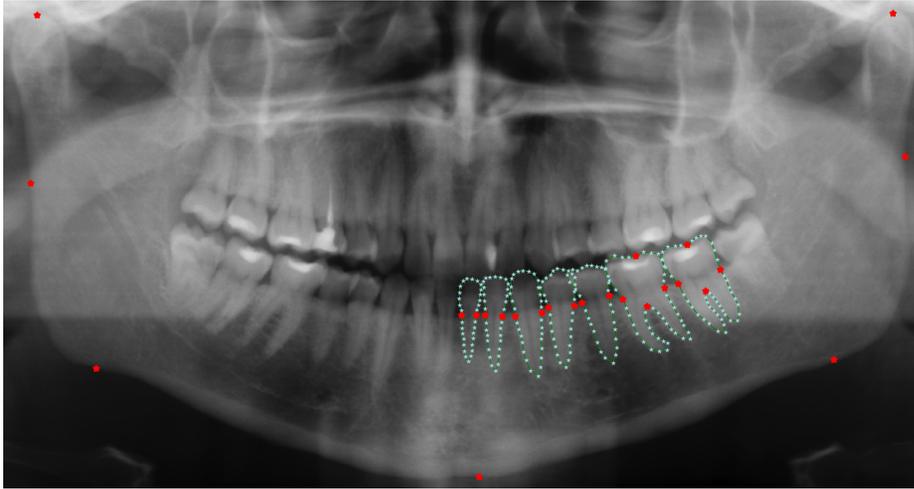


Fig. 1. Annotated points in each single image. In red: teeth and mandible keypoints; in blue: teeth non-keypoints.

The predicted shapes of teeth 31 to 37 have been compared to manually annotated shapes and the performance of the proposed approach has been assessed in three ways. Firstly, the performance of present/missing teeth detection has been measured. Tab. 2 shows the classification results when choosing a threshold to maximise (true positive rate - false positive rate). See Fig. 4 for some examples. Secondly, to assess whether they have been located correctly, the intersection over union (IoU) of annotated and predicted shapes was calculated from the examples

Table 1. Number of points used in each individual model and number of points of each model used in the keypoint model.

Model	#points	#points used in keypoint model
Central incisor (31)	32	2 (one on each neck side)
Lateral incisor (32)	32	2 (one on each neck side)
Canine (33)	32	2 (one on each neck side)
First premolar (34)	32	2 (one on each neck side)
Second premolar (35)	32	2 (one on each neck side)
First molar (36)	48	4 (one on each neck side, one at the top and one at the root bifurcation)
Second molar (37)	48	4 (one on each neck side, one at the top and one at the root bifurcation)
Mandible	7	7 (nearly equidistant points between the top of the condyles)

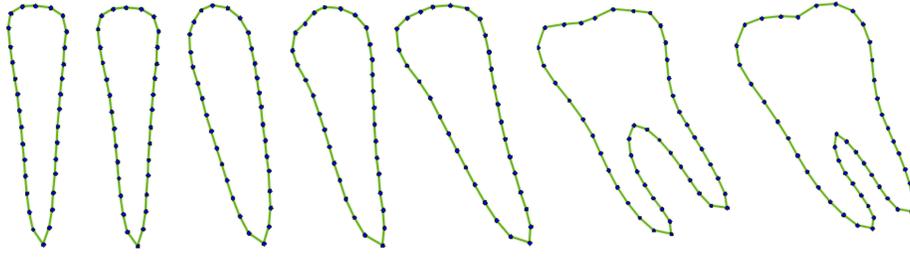


Fig. 2. Mean shape of each individual tooth model.

where both teeth are present and are correctly detected as present. Tab. 3 shows that the detection of multiroot teeth (36 and 37) is slightly more successful than the detection of single root teeth. This is likely to be because the anterior teeth are closer to each other so the model might match a neighbouring tooth. Assuming that an overlap greater than 50% between the prediction and the ground truth indicates that the predicted shape is very likely to match the real tooth, the examples with a IoU value over 0.5 have been treated as correctly located. In general, the proportion of well-located teeth is over 90% among all teeth types. Thirdly, the accuracy of the tooth shape matching has been evaluated on the correctly located teeth (where the overlap between model and true tooth is greater than 50%) with the point-to-curve error, which represents the shortest distance from each estimated point to the curve through the ground truth landmark points (Tab. 4). The median of the errors is less than 0.23 mm for all types of teeth. The 99%-ile is 1.31 mm in the worst case, which demonstrates the robustness of the proposed segmentation approach. Note that all performance measurements have been obtained on the left mandibular teeth only as we did not have manual ground truth annotations for the right side. However, the right

Table 2. Confusion matrix of the missing teeth detection problem and binary classification metrics. In order to obtain these metrics, the “present” class has been considered as the positive class.

		Prediction		Total	Accuracy	95.46%
		Missing	Present			
Real	Missing	31 (TN)	6 (FP)	37	Precision	98.9%
	Present	21 (FN)	537 (TP)	558	Sensitivity	96.24%
Total		53	542	595	Specificity	83.78%

Table 3. Intersection over the union (IoU) statistics for each individual tooth predictions: mean, standard deviation, median. In the last column, the percentage of the examples with an IoU over 0.5, which are treated as correctly located.

Model	mean	sd	median	% over 0.5
31	0.78	0.24	0.86	89.41%
32	0.80	0.20	0.86	92.94%
33	0.84	0.18	0.90	94.12%
34	0.84	0.22	0.91	91.76%
35	0.85	0.23	0.93	92.21%
36	0.87	0.18	0.92	95.77%
37	0.86	0.21	0.92	94.29%

mandibular teeth can be outlined by applying the left mandibular teeth models to the horizontally reflected images. See Fig. 3 for some examples.

4 Discussion and Conclusions

We have shown that a state-of-the-art performance can be achieved in adult mandibular teeth segmentation by using the RFRV-CLM algorithm in two steps. The first step provides an estimation of some teeth and mandible keypoints, which are used to initialise each individual tooth search. In the second step, the search of each tooth is performed independently. This two-step approach overcomes the problem of automatically initialising each individual tooth model, and the results show that the teeth shapes can be matched very accurately, especially if the tooth is correctly located.

A limitation of this study is that we have not taken into account the third molar (also known as the *wisdom tooth*). This is because this tooth is often extracted or missing in some patients so we had very few examples. Moreover, although the QoF statistics are a good starting point for missing teeth detection,

Table 4. Point to curve statistics in each individual tooth model (in mm): mean, standard deviation, median and 90, 95 and 99 centiles. These results have been obtained on the examples where the teeth have been correctly located.

Model	mean	sd	med	90%	95%	99%
31	0.23	0.12	0.21	0.38	0.47	0.61
32	0.30	0.23	0.22	0.58	0.65	1.29
33	0.30	0.24	0.23	0.56	0.76	1.31
34	0.24	0.20	0.19	0.36	0.56	0.90
35	0.21	0.11	0.18	0.30	0.48	0.58
36	0.25	0.16	0.20	0.35	0.50	0.89
37	0.27	0.14	0.23	0.41	0.48	0.82

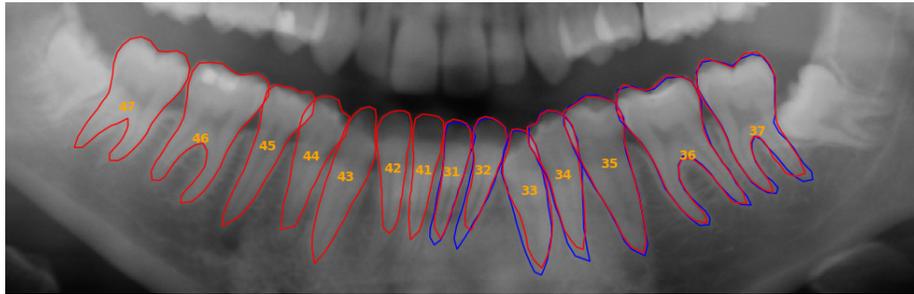
this task could be improved by using other metrics or algorithms developed specifically for that purpose.

Nonetheless, the presented results are promising and are a big step towards a fully automatic dental assessment tool with a variety of applications. Two direct uses of the proposed system are (i) automatic teeth measurements with a view to plan surgical treatments; and (ii) automatic radiograph matching with the aim of identifying people (e.g. in forensics). Other clinical tasks could also be carried out with this system and few functionality additions. For example, the detection of caries, impacted tooth and other abnormalities.

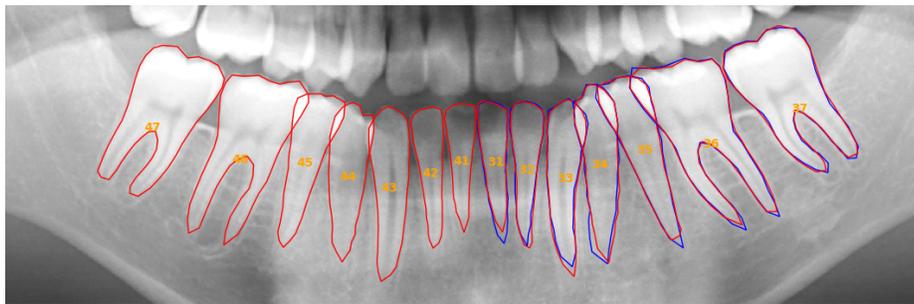
Acknowledgements. This work has received financial support from the Consellería de Cultura, Educación e Ordenación Universitaria (accreditation 2016-2019, ED431G/08, growth potential group 2017-2020 ED431B 2017/029, reference competitive group 2017-2020, ED431C 2017/69, and N. Vila Blanco support ED481A-2017) and the European Regional Development Fund (ERDF). C. Lindner is funded by the Medical Research Council, UK (MR/S00405X/1).

References

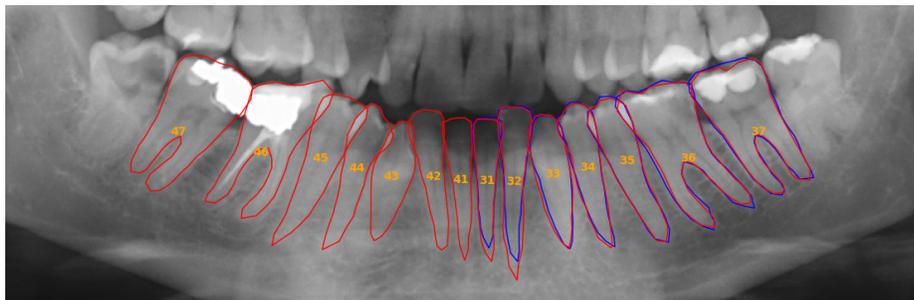
1. Rondon, R., Pereira, Y., do Nascimento, G.: Common positioning errors in panoramic radiography: a review. *Imaging Sci. Dent.* 44(1), 1–6 (2014), <https://doi.org/10.5624/isd.2014.44.1.1>
2. Halperin-Sternfeld, M., Machtei, E., Balkow, C., Horwitz, J.: Patient movement during extraoral radiographic scanning. *Oral Radiol.* 32(1), 40–47 (2016), <https://doi.org/10.1007/s11282-015-0208-6>
3. Chen, H., Jain, A.: Tooth contour extraction for matching dental radiographs. In: Proc. 17th International Conference on Pattern Recognition – ICPR 2004, vol. 3, pp. 522–525. IEEE (2004), <https://doi.org/10.1109/ICPR.2004.1334581>
4. Nomir, O., Abdel-Mottaleb, M.: A system for human identification from X-ray dental radiographs. *Pattern Recognit.* 38(8), 1295–1305 (2005), <https://doi.org/10.1016/j.patcog.2004.12.010>
5. Chen, H., Jain, A.: Dental biometrics: alignment and matching of dental radiographs. In: Proc. 7th IEEE Workshops on Application of Computer Vision – WACV/MOTION 2005, vol. 1, pp. 316–321. IEEE (2005), <https://doi.org/10.1109/ACVMOT.2005.41>
6. Oliveira, J., Proença, H.: Caries detection in panoramic dental X-ray images. In: Tavares, J., Natal Jorge, R. (eds.) *Computational Vision and Medical Image Processing, Computational Methods in Applied Sciences*, vol. 19, pp. 175–190. Springer (2011), https://doi.org/10.1007/978-94-007-0011-6_10
7. Li, S., Fevens, T., Krzyżak, A., Jin, C., Li, S.: Semi-automatic computer aided lesion detection in dental X-rays using variational level set. *Pattern Recognit.* 40(10), 2861–2873 (2007), <https://doi.org/10.1016/j.patcog.2007.01.012>
8. Čular, L., Tomaić, M., Subašić, M., Šarić, T., Sajković, V., Vodanović, M.: Dental age estimation from panoramic X-ray images using statistical models. In: Proc. 10th International Symposium on Image and Signal Processing and Analysis – ISPA 2017, pp. 25–30. IEEE (2017), <https://doi.org/10.1109/ISPA.2017.8073563>



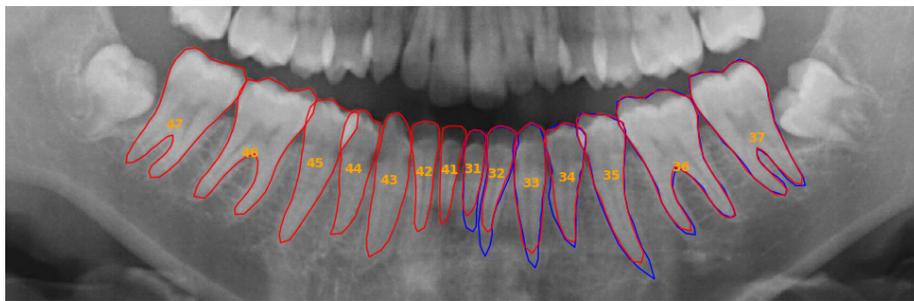
(a)



(b)

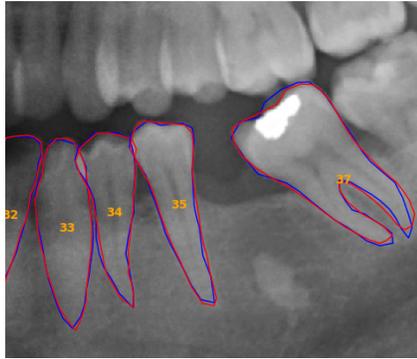


(c)

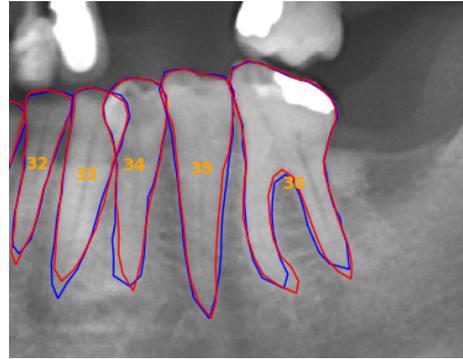


(d)

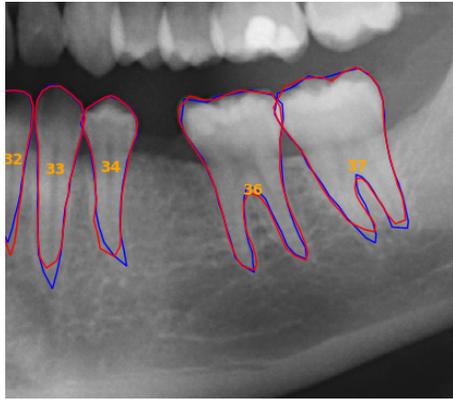
Fig. 3. Results of automatic teeth segmentation process. In red, the predicted shapes. In blue, the left-mandibular manually annotated shapes. The segmentation is robust to some issues such as very bright images (b) or tooth filling (c). It also can manage teeth overlapping. It is worth noting that the most noticeable segmentation errors are observed in the apical regions (around root apices) due to the low contrast of the image in that area.



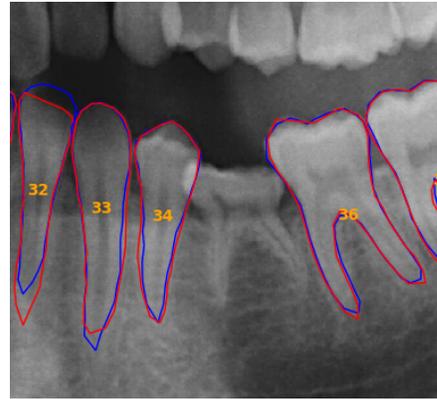
(a) Missing tooth 36 detected correctly.



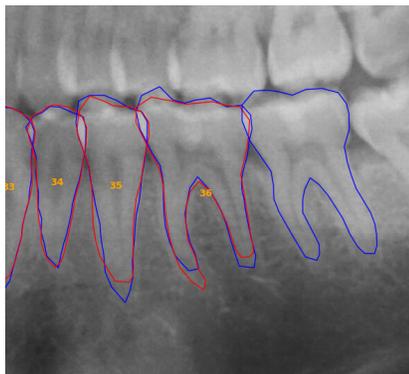
(b) Missing tooth 37 detected correctly.



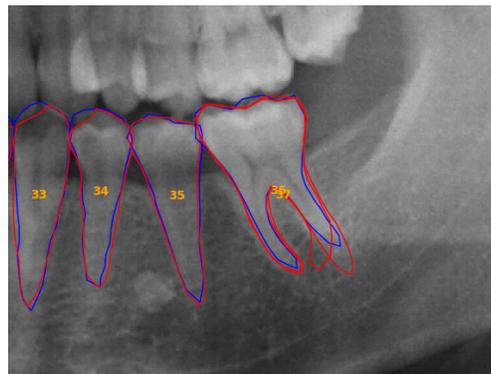
(c) Missing tooth 35 detected correctly.



(d) Tooth 35 searcher detects correctly tooth as missing. This is a special clinic case called agenesis, where the tooth (premolar 35 in the image) fails to form, so the related primary tooth is not pushed outward.



(e) Tooth 37 detected incorrectly as missing.



(f) Tooth 37 detected incorrectly as present due to it matches tooth 36.

Fig. 4. Results of present/missing teeth detection. In red, the predicted shapes. In blue, the left-mandibular manually annotated shapes.

9. Razali, M., Ahmad, N., Zaki, Z., Ismail, W., et al.: Region of adaptive threshold segmentation between mean, median and Otsu threshold for dental age assessment. In: Proc. International Conference on Computer, Communications, and Control Technology – I4CT 2014, pp. 353–356. IEEE (2014), <https://doi.org/10.1109/I4CT.2014.6914204>
10. Lira, P., Giraldo, G., Neves, L.: Panoramic dental X-ray image segmentation and feature extraction. In: Proc. V Workshop of Computing Vision (Sao Paulo, Brazil) (2009)
11. Amer, Y., Aqel, M.: An efficient segmentation algorithm for panoramic dental images. *Procedia Comput. Sci.* 65, 718–725 (2015), <https://doi.org/10.1016/j.procs.2015.09.016>
12. Shah, S., Abaza, A., Ross, A., Ammar, H.: Automatic tooth segmentation using active contour without edges. In: Proc. 2006 Biometrics Symposium: Special Session on Research at the Biometric Consortium Conference, pp. 1–6. IEEE (2006), <https://doi.org/10.1109/BCC.2006.4341636>
13. Lindner, C., Bromiley, P., Ionita, M., Cootes, T.: Robust and accurate shape model matching using random forest regression-voting. *IEEE Trans. Pattern Anal. Mach. Intell.* 37(9), 1862–1874 (2015), <https://doi.org/10.1109/TPAMI.2014.2382106>
14. Cristinacce, D., Cootes, T.: Automatic feature localisation with constrained local models. *Pattern Recognit.* 41(10), 3054–3067 (2008), <https://doi.org/10.1016/j.patcog.2008.01.024>
15. Lindner, C., Thiagarajah, S., Wilkinson, J., The arcOGEN Consortium, Wallis, G., Cootes, T.: Fully automatic segmentation of the proximal femur using random forest regression voting. *IEEE Trans. Med. Imaging* 32(8), 1462–1472 (2013), <https://doi.org/10.1109/TMI.2013.2258030>
16. Cootes, T., Ionita, M., Lindner, C., Sauer, P.: Robust and accurate shape model fitting using random forest regression voting. In: Fitzgibbon, A., et al. (eds.) Proc. 12th European Conference on Computer Vision – ECCV 2012, Lect. Notes Comput. Sc., vol. 7578, pp. 278–291. Springer (2012), https://doi.org/10.1007/978-3-642-33786-4_21
17. Bromiley, P., Adams, J., Cootes, T.: Localisation of vertebrae on DXA images using constrained local models with random forest regression voting. In: Yao, J., et al. (eds.) Proc. 2nd MICCAI Workshop on Computational Methods and Clinical Applications for Spine Imaging – CSI 2014, Lect. Notes Comput. Vis. Biomech., vol. 20, pp. 159–171. Springer (2015), https://doi.org/10.1007/978-3-319-14148-0_14
18. Viola, P., Jones, M.: Rapid object detection using a boosted cascade of simple features. In: Proc. 2001 IEEE Conference on Computer Vision and Pattern Recognition – CVPR 2001. IEEE (2001), <https://doi.org/10.1109/CVPR.2001.990517>