Deep Reinforcement Learning for Detection of Abnormal Anatomies

Paula López Diez *[1], Kristine Aavild Juhl[1], Josefine Vilsbøll Sundgaard[1], Hassan Diab[4], Jan Margeta[3], François Patou[2], and Rasmus R. Paulsen[1]

[1]DTU Compute, Technical University of Denmark, Kongens Lyngby, Denmark
[2]Oticon Medical, Research & Technology group, Smørum, Denmark
[3]KardioMe, Research & Development, Nova Dubnica, Slovakia
[4]The National Medical Research Center for Otorhinolaryngology of the Federal Medico-Biological Agency of Russia, Moscow, Russia

Abstract

Automatic detection of abnormal anatomies or malformations of different structures of the human body is a challenging task that could provide support for clinicians in their daily practice. Compared to normative anatomies, there is a low presence of anatomical abnormalities in patients, and the great variation within malformations make it challenging to design deep learning frameworks for automatic detection. We propose a framework for anatomical abnormality detection, which benefits from using a deep reinforcement learning model for landmark detection trained in normative data. We detect the abnormalities using the variability between the predicted landmarks configurations in a subspace based on a point distribution model of landmarks using Procrustes shape alignment and principal component analysis projection from normative data. We demonstrate the performance of this implementation on clinical CT (Computed Tomography) scans of the inner ear, and show how synthetically created abnormal cochlea anatomy can be detected using the prediction of five landmarks around the cochlea. Our approach shows a Receiver Operating Characteristics (ROC) Area Under The Curve (AUC) of 0.97, and 96% accuracy for the detection of abnormal anatomy on synthetic data.

*Corresponding Author: plodi@dtu.dk

1 Introduction

The detection of abnormal anatomies in medical images has a key role in accurate diagnosis making. The detection of different types of malformations in the clinic is usually performed by visual inspection of clinical images, which makes the diagnosis and detection very sensible to the practitioner's experience and subjectivity. State-of-the-art deep learning methods have shown high performance for automatic detection of anomalies as presented in [3] using well-known architectures. Similar deep learning approaches are also use in a clinical context for anatomical anomalies as in [6]. However, training such a model requires large amounts of labeled medical data, which is challenging and expensive to acquire. Additionally, these approaches suffer from limited generalizability because training data rarely faithfully represents all possible pathological appearances, as images representing the less common malformations are particularly complicated to acquire [1]. It has been shown how simpler methods, such as principal component analysis (PCA), can be employed for anomaly detection in e.g. physiological measurements [2] [8]. While for detection of abnormal anatomies in image data, most studies employ more advanced deep learning methods due to the increased complexity of the data [1, 11].

In this paper, we use a Communicative Multiple Agent Reinforcement Learning (C-MARL) model [9] for landmark detection within the structures of the inner ear. Automatic landmark de-

https://doi.org/10.7557/18.6280

tection is a very active field of research, and with the popularization of deep learning, multiple neural networks have been used for landmark localization. From an object search perspective, the problem of locating landmarks in 3D images can result in drops of accuracy and increased processing times due to the unnecessary exhaustive mapping or scanning of the images. Guesu et al. [5] first showed how a deep reinforcement learning approach allows different agents to learn the optimal policy to locate the landmarks' positions using the image information at different scales. This methodology resembles normal human strategy, where the person who is looking for a certain landmark would initially locate the region of the image where the landmark should be, and zoom-in multiple times for a fine-tuned location of the exact point. To benefit from the use of multiple agents, Vlontzos et al. [13] introduced implicit communication between the agents (MARL model), in which the agents share the weights of the convolutional neural network (CNN) layers. Leroy et al. [9] further included explicit communication between the agents, sharing the average weights of the fully connected layers (C-MARL model). This communication scheme allows for a more robust localization of landmarks, especially when they present a spatial correlation across the dataset.

We base our abnormal anatomy detection on the C-MARL model, with the expectation that when a certain structure is not present in the CT scan, the agents of the C-MARL model trained exclusively on normal cases will not converge to a landmark located in the missing or malformed anatomy. We thus exploit the knowledge about the normal anatomies, as this carries implicit information about the abnormal anatomies. The agents will show a lower degree of agreement regarding the final position of the landmark, when presented with an abnormal anatomy. To show this variation, we project the set of estimated landmarks in a space that is configured according to the normal anatomical variations of the landmarks. In order to do so, we align the landmarks of the training set using the Procrustes analysis [7] and use a PCA to decrease the dimensions of the model. We then examine the variability between the agents from the C-MARL model in the PCA space, to determine whether or not the anatomy is normal. The approach is evaluated on a dataset of clinical CT scans of the inner ear which are labelled with the seven landmarks shown in Figure 1. From these images we artificially remove the cochlea structure, thus artificially simulating a CT scan of a patient with cochlear aplasia. It is known, that this malformation in the inner ear is clinically relevant for the diagnosis of hearing loss [12]. Furthermore, the detection of some type of anomaly in the image can be used not only for setting the special cases apart but also as an initial classification that could potentially be the first block of a sub-categorical classification of the abnormalities.

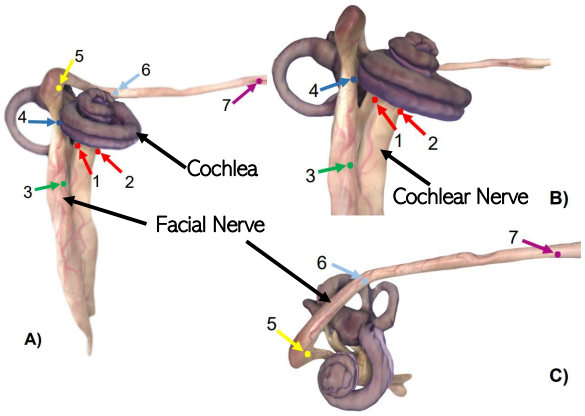

Figure 1: Landmark locations within the facial nerve and cochlear nerve. A) Overview of the seven landmarks B) Close-up of the landmarks 1-4 C) Close-up of landmarks 5-7. Edited from [10]

2 Methods

Our approach is based on using the output variability of a model trained with normative data to detect data anomalies. We employ the C-MARL model [9] for detecting landmarks around the cochlear structure in the inner ear. In this configuration, the deep reinforcement agents navigate through the 3D image (environment) and observe their state, which is defined as a patch of the image centered in the agent location. This patch becomes smaller as the agents get closer to the landmark (multi-scale). Based on the observed state, they take different actions from their action set (move up, down, left, right, forward and backward) and receive a reward, which is a function of the Euclidean distance be-

tween the current position of the agent and the previous one relative to the target point (positive when agent is getting closer and negative otherwise). The expected reward of taking a certain action given a state is known as the Q-value. In deep reinforcement learning the Q-value of a certain state associated with each of the possible actions is estimated by the use of a Deep-Q-Network.

For the C-MARL model the architecture of the Deep-Q-Network resembles a typical image classification architecture, but with a set of fully connected layers for each agent. The model diagram and model architecture is shown in Figure 2. The common CNN weights among all agents provide implicit communication between the agents, meaning the share the same layers responsible of extracting the relevant features for their current state. Meanwhile, the shared average weight of the different fully connected layers allows for implicit communication between agents, sharing information of the layers that are used to map the extracted features from the current state to the Q-value of each action at that point. This setup has been proven specially good when the different landmarks have a consistent spatial correlation as it is the case of the inner ear anatomy. The training configuration employed is the same as presented by López Diez et al. [10], where an overall rate of 2.6% incorrectly located landmarks was reported, with an average error of 1.218 mm.

The agents are randomly initialized within 80% of the image to avoid initialization on an edge. This randomness makes the final estimation of the landmarks a stochastic process. To be able to derive some statistically significant results, we have computed predictions five times for each of the images in the test set.

For normal anatomies, it is assumed that the found landmarks are placed in a certain spatial configuration and that for abnormal cases, the found landmarks will deviate significantly from this configuration. In order to test if a case is within the normal configuration, a point distribution model (PDM) is constructed following the approach in [4]. In the following, the set of found or annotated landmarks from a single scan is called a *shape* and there is point-correspondence over all shapes in the training samples. The shapes are initially aligned using a generalized Procrustes analysis [7]. A similarity transform is used for the alignment and therefore

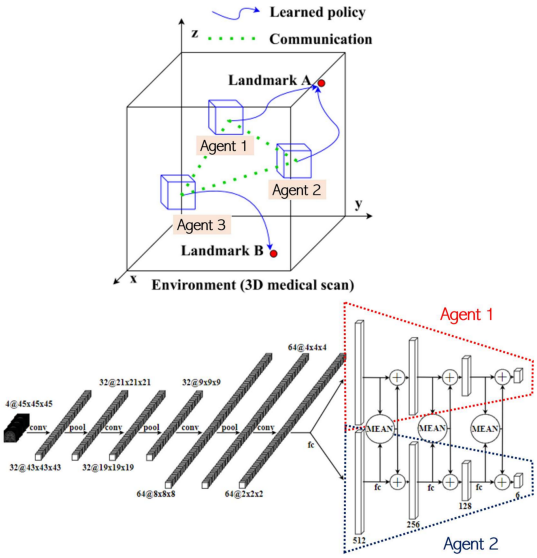

Figure 2: Diagram and architecture of the C-MARL model. Modified from [10]

the PDM will describe the shape variation only and not include the size variation.

Following the Procrustes analysis, a mean shape, $\overline{\mathbf{x}}$ is estimated and the aligned shapes can be used in a principal component analysis (PCA) of shape variation following [4]. The result of the PCA analysis is a set of principal components, concatenated into a matrix Φ, describing the modes of shape variation. A new shape can be synthesized by: $\mathbf{x}_{\text{new}} = \overline{\mathbf{x}} + \Phi\mathbf{b}$. Here $\mathbf{b}$ is a vector of weights controlling the modes of shape variation and Φ contains the first t principal components. A given $\mathbf{x}'$ shape can be aligned to the Procrustes mean and be approximated by the PDM model by projecting the residuals from the average shape into principal component space: $\mathbf{b} = \Phi^T(\mathbf{x}' - \overline{\mathbf{x}})$. The resulting $\mathbf{b}$ vector describes the shape in terms of coordinates in PCA space and is used in the further analysis.

The C-MARL model is trained on images of normal anatomies and their annotated landmarks. These landmarks are also used to build the PDM of the normal anatomy shape configuration. At test time we use three agents per landmark and randomly combine the three predictions into three different shapes of the full anatomical structure. By approximating the shapes with the PDM, we are able to reason about how likely the landmark con-

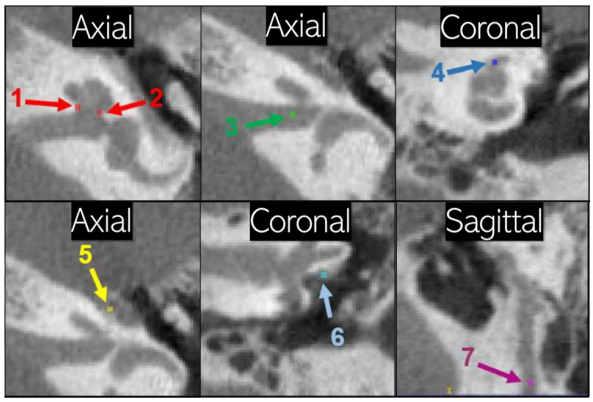

Figure 3: Landmarks' location in a sample CT scan from [10]

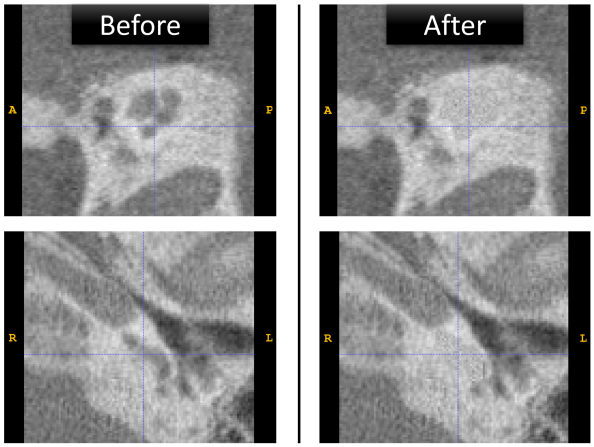

Figure 4: Sample CT scan from test set, before (left) and after (right) the artificial cochlea removal process.

figuration is, as opposed to merely looking at individual distances between single landmarks. Therefore, each shape is projected into the PCA space and described by a vector $\mathbf{b}_i$, where $i = 1, 2, 3$ are the different agents. To quantify the variation between the different agents, we measure the standard deviation between the PCA loadings in the first 6 dimensions (corresponding to 90% of the variation in the PDM).

3 Data

The dataset consists of 120 clinical CT scans of the inner ear from cochlear implant patients with normal inner ear anatomy [10]. From each CT scan, a region of interest is cropped centered in the cochlea with a cubic shape of 32.1x32.1x32.1 mm^3. The average voxel side length of the scans is 0.3 mm in the range of [0.13, 0.45] mm. Each CT scan is annotated with seven landmarks along the facial and the cochlear nerve in the nearby region of the cochlea, as shown in Figure 1 and 3.

To describe the cochlear image with the landmarks, we use 5 of the 7 landmarks defined in [10]. Landmark one and two are placed on opposite sides of the cochlear nerve in the axial view, number three identifies the point where the facial nerve exits the internal acoustic canal, number four shows the closest point of the facial nerve to the cochlea structure, number five shows the high curvature point of the facial nerve, and six and seven characterize the more elongated part of this nerve in a region more distant from the cochlea. Neither of these landmarks are placed within the cochlea.

For the further analysis for abnormality detection, only the five first landmarks are employed. The first five landmarks are all positioned around the cochlea, while landmarks six and seven are further away from the anatomy in question for this work. The last two landmarks, and thus the six corresponding agents in the C-MARL model, are therefore ignored in the further analysis. The dataset is split into a training set containing 97 CT scans, and a test set containing 23 CT scans. The training data is used to train the C-MARL model for landmark detection, while the test data is used for the rest of the analysis of abnormality detection.

Abnormal CT scans are artificially generated by removing the cochlea structure from the images in the test set, thus generating corresponding pairs of normal and abnormal CT scans with the same surrounding structures of the inner ear. The cochlea is delineated using the ITK-SNAP software [14] to generate a rough segmentation of the cochlear structure. This section of the image is then replaced by Gaussian noise, with mean and standard deviation estimated from the intensities of the surrounding region of the segmentation. An example of the transformation process is shown in Figure 4.

4 Results

For evaluation of the detection of abnormal anatomies, the variation between the PCA loadings for the 15 agents trained for the five landmarks is evaluated. The plot in Figure 6a shows the standard deviation for each corresponding pair of normal and abnormal anatomy, and the graphs show how the standard deviation is increased for the abnormal cases compared to the normal ones. In Figure 6b it can be observed that the ROC curves show a high ROC AUC and maximum accuracy for all five runs, but the averaged method, where all five runs are used to compute the overall standard deviation, shows the highest ROC AUC of 0.97 and a 96% accuracy for detection of abnormal anatomies.

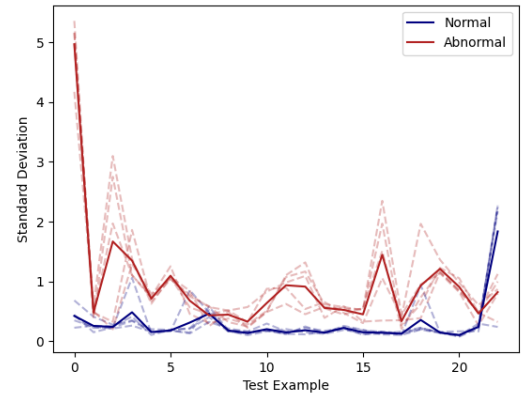

(a)

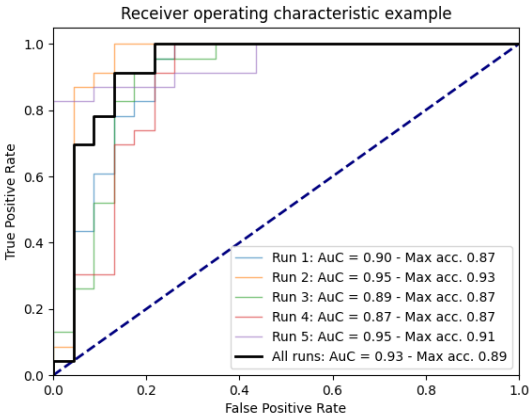

(b)

Figure 5: Results for the direct method on the landmark coordinates. a) Standard deviation between the landmark coordinates for the 15 agents per run (dotted lines) and their average (solid line) for the normal and abnormal cases. b) ROC curves representing the five individual runs (colored lines) and using the average standard deviation across all runs (black line).

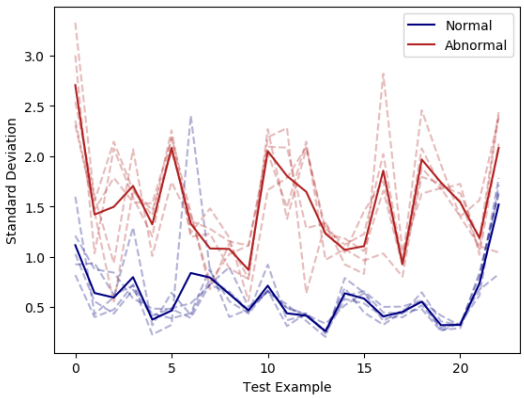

(a)

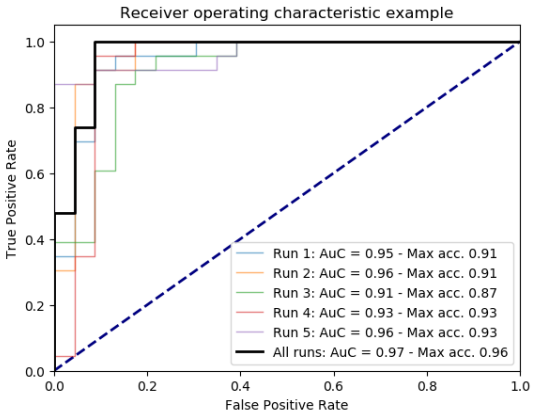

(b)

Figure 6: Results for our proposed PDM and PCA approach. a) Standard deviation between the PCA loadings for the 15 agents per run (dotted lines) and their average (solid line) for the normal and abnormal cases. b) ROC curves representing the five individual runs (colored lines) and using the average standard deviation across all runs (black line).

The results are computed using the test set which contains 23 3D volumes of normative data that has not been seen by the DRL model (trained on 97 images) and the abnormal anatomy version of each image artificially generated as described in Section 3. Overall, the test set contains 46 CT scans: 23 with normal anatomy and 23 with an artificial abnormality. In order to evaluate whether the PCA space of the PDM provides an advantage against analyzing the final configuration of the agents in the original coordinate space, we have also evaluated the variability in the original image space. For each of the three agents per landmark, we calculate the standard deviation of their mutual Euclidean distance. To be able to compare with the representation in the sub-dimensional space, we average all the standard deviations of the five landmarks in each CT scan, so we obtain one value for each scan for each run. Figure 5a shows the representation of these values for the 23 pairs of normal and abnormal CT scans in the test set for each run, together with the overall average among the five runs. It should be noted that the standard deviation in Figures 6 and 5 corresponds to different measurements and the direct comparison should be avoided. The standard deviation seen in Figure 6 is the standard deviation of the distance between each shape projection in the 6-dimensional PCA space of normative data, while in Figure 5 we observe the mean standard deviation of the agents position among the same CT.

Figure 5b shows the ROC curves associated with this approach, as well as the maximum accuracy for each run. In this case, the best performing method is not the average between the five runs, as previously seen. The best results is found for run number two, where the ROC AUC is 0.95, and the maximum accuracy for all runs is 93%.The performance of this method is thus lower than of the proposed method with PDM and PCA, and less stable as the accuracy results of this method have a larger variation between the runs. Figures 6 and 5 cannot be directly compared as in Figure 6 we use the standard deviation of distances in a 6-dimensional PCA space and in 5 we use the standard deviation of distance in mm in the 3-dimensional image space.

5 Discussion

Taking five independent runs of the landmark predictions into account resulted in better results and a more reliable assessment of our method. It can be seen that the accuracy benefits from taking all the runs into account as shown in Figure 6b, where the best performing accuracy of 96% is obtained when taking all five independent runs into consideration. This leads us to believe that the method could benefit from an increased number of runs, given that it appears to provide statistical stability to the method. Potentially, a larger number of agents per landmark could be trained, which may also improve robustness. However, training a greater number of communicating agents would require a significantly longer training and testing time, so a trade-off should be sought between the method's accuracy and the computing cost. We believe our current approach with five runs and three agents per landmark shows a good compromise between accuracy and computational cost.

We have used artificially generated abnormal data to test our approach. This scenario has shown a good performance providing a reliable proof of concept for our method. Further work aims at testing the approach on data with real anomalies and seeing whether or not the results are consistent with the ones presented in this work.

The results show that our proposed approach of measuring the variation of the PCA of the PDM achieves a higher performance compared to measuring the variation of the located landmarks in the image space. The PDM takes into account the full shape of the anatomy, while the inter-agent variance in the image space only takes into account the level of agreement between the multiple agents associated with a certain landmark, but ignores the overall shape. These results show that the method benefits from the shape model, and leads to a more robust prediction of anomalies than just using the analysis in the image space.

Comparison against supervised methods is not possible at this stage due to the lack of labeled data for training, in a further analysis more artificial abnormal data could be generated to train a fully supervised method. However, we consider one of the main advantages of our approach is that it does not require data with abnormalities for the training stages.

6 Conclusion

We have presented an approach for detection of abnormal anatomies in the inner ear based on landmark predictions from a C-MARL deep reinforcement learning model. The method has demonstrated a high performance on the synthetic data for the detection of presence or absence of the cochlea structure. This method manages to both locate the area of the cochlea in the CT scan, and then classify whether or not the structure is present.

References

[1] C. Baur, B. Wiestler, M. Muehlau, C. Zimmer, N. Navab, and S. Albarqouni. Modeling healthy anatomy with artificial intelligence for unsupervised anomaly detection in brain mri. *Radiology: Artificial Intelligence*, 3(3):e190169, 2021. doi: 10.1148/ryai.2021190169. URL `https://doi.org/10.1148/ryai.2021190169`.

[2] L. Ben Amor, I. Lahyani, and M. Jmaiel. PCA-based multivariate anomaly detection in mobile healthcare applications. In *Proc. International Symposium on Distributed Simulation and Real Time Applications (DS-RT)*, pages 1–8, 2017. doi: 10.1109/DISTRA.2017.8167682.

[3] R. Chalapathy and S. Chawla. Deep learning for anomaly detection: A survey. 1 2019. URL `http://arxiv.org/abs/1901.03407`.

[4] T. F. Cootes, C. J. Taylor, D. H. Cooper, and J. Graham. Active shape models-their training and application. *Computer vision and image understanding*, 61(1):38–59, 1995. doi: https://doi.org/10.1006/cviu.1995.1004.

[5] F. C. Ghesu, B. Georgescu, S. Grbic, A. K. Maier, J. Hornegger, and D. Comaniciu. Robust multi-scale anatomical landmark detection in incomplete 3d-CT data. In *Proc. MICCAI*, pages 194–202, 2017. ISBN 978-3-319-66182-7. doi: 10.1007/978-3-319-66182-7_23.

[6] R. S. Gill, S.-J. Hong, F. Fadaie, B. Caldairou, B. C. Bernhardt, C. Barba, A. Brandt, V. C. Coelho, L. d'Incerti, M. Lenge, M. Semmelroch, F. Bartolomei, F. Cendes, F. Deleo, R. Guerrini, M. Guye, G. Jackson, A. Schulze-Bonhage, T. Mansi, N. Bernasconi, and A. Bernasconi. Deep convolutional networks for automated detection of epileptogenic brain malformations. In A. F. Frangi, J. A. Schnabel, C. Davatzikos, C. Alberola-López, and G. Fichtinger, editors, *Proc. MICCAI*, pages 490–497. Springer, 2018. ISBN 978-3-030-00931-1. doi: 10.1007/978-3-030-00931-1_56.

[7] J. C. Gower. Generalized procrustes analysis. *Psychometrika*, 40(1):33–51, 1975. doi: 10.1007/bf02291478.

[8] V. A. Krenn, C. Fornai, N. M. Webb, M. A. Woodert, H. Prosch, and M. Haeusler. The morphological consequences of segmentation anomalies in the human sacrum. *American Journal of Biological Anthropology*, 12 2021. ISSN 2692-7691. doi: 10.1002/ajpa.24466. URL `https://onlinelibrary.wiley.com/doi/10.1002/ajpa.24466`.

[9] G. Leroy, D. Rueckert, and A. Alansary. Communicative reinforcement learning agents for landmark detection in brain images. In *Machine Learning in Clinical Neuroimaging and Radiogenomics in Neuro-oncology*, pages 177–186. Springer, 2020. ISBN 978-3-030-66843-3. doi: 10.1007/978-3-030-66843-3_18.

[10] P. López Diez, J. V. Sundgaard, F. Patou, J. Margeta, and R. R. Paulsen. Facial and cochlear nerves characterization using deep reinforcement learning for landmark detection. In *Proc. MICCAI*, pages 519–528. Springer, 2021. ISBN 978-3-030-87202-1. doi: 10.1007/978-3-030-87202-1_50.

[11] P. Seeböck, J. I. Orlando, T. Schlegl, S. M. Waldstein, H. Bogunović, S. Klimscha, G. Langs, and U. Schmidt-Erfurth. Exploiting epistemic uncertainty of anatomy segmentation for anomaly detection in retinal oct. *IEEE Transactions on Medical Imaging*, 39(1):87–98, 2020. doi: 10.1109/TMI.2019.2919951.

[12] L. Sennaroğlu and M. D. Bajin. Classification and current management of inner ear malformations. *Balkan Medical Journal*, 34, 08 2017. doi: 10.4274/balkanmedj.2017.0367.

[13] A. Vlontzos, A. Alansary, K. Kamnitsas, D. Rueckert, and B. Kainz. Multiple Landmark Detection Using Multi-agent Reinforcement Learning. In *Proc. MICCAI*. Springer, 2019. doi: 10.1007/978-3-030-32251-9_29.

[14] P. A. Yushkevich, J. Piven, H. Cody Hazlett, R. Gimpel Smith, S. Ho, J. C. Gee, and G. Gerig. User-guided 3D active contour segmentation of anatomical structures: Significantly improved efficiency and reliability. *Neuroimage*, 31(3):1116–1128, 2006. doi: 10.1016/j.neuroimage.2006.01.015.