Deep Reinforcement Learning for Detection of Inner Ear Abnormal Anatomy in Computed Tomography

Lopez Diez, Paula; Sørensen, Kristine Aavild; Sundgaard, Josefine Vilsbøll; Diab, Khassan ; Margeta, Jan ; Patou, Francois ; Paulsen, Rasmus Reinhold

Published in:
Proceedings of 25th International Conference on Medical Image Computing and Computer Assisted Intervention

Link to article, DOI:
10.1007/978-3-031-16437-8_67

Publication date:
2022

Document Version
Peer reviewed version

Deep Reinforcement Learning for Detection of Inner Ear Abnormal Anatomy in Computed Tomography

Paula López Diez[1], Kristine Sørensen[1], Josefine Vilsbøll Sundgaard[1], Khassan Diab[4], Jan Margeta[3], François Patou[2], and Rasmus 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] Tashkent International Clinic, Tashkent, Uzbekistan
plodi@dtu.dk

Abstract. Detection of abnormalities within the inner ear is a challenging task that, if automated, could provide support for the diagnosis and clinical management of various otological disorders. Inner ear malformations are rare and present great anatomical variation, which challenges the design of deep learning frameworks to automate their detection. We propose a framework for inner ear abnormality detection, based on a deep reinforcement learning model for landmark detection trained in normative data only. We derive two abnormality measurements: the first is based on the variability of the predicted configuration of the landmarks in a subspace formed by the point distribution model of the normative landmarks using Procrustes shape alignment and Principal Component Analysis projection. The second measurement is based on the distribution of the predicted Q-values of the model for the last ten states before the landmarks are located. We demonstrate an outstanding performance for this implementation on both an artificial (0.96 AUC) and a real clinical CT dataset of various malformations of the inner ear (0.87 AUC). Our approach could potentially be used to solve other complex anomaly detection problems.

Keywords: Deep Reinforcement Learning · Anomaly Detection · Inner Ear · Congenital Malformation

1 Introduction

Sensorineural hearing loss (SNHL) in children is a major cause of disability. Generally SNHL is detected early in many parts of the world, which allows the prescription of interventions that mitigate the risk of abnormal social, emotional and communicative development. Such interventions include Cochlear Implant (CI) therapy which is prescribed each year to about 80,000 infants and toddlers. Congenital SNHL is sometimes the consequence of an abnormal embryonic development. Resulting malformations are generally classified according to

two categories: membranous malformations, which are not observable in conventional medical scans, and congenital malformations, which can be detected by Computed Tomography (CT) or Magnetic Resonance Imaging (MRI) [4]. These cases raise surgical challenges during surgical planning of the CI therapy and during the surgery itself, often requiring the surgeon to discover and adapt to the anatomy of the malformation during the operation. Anticipating the presence of such malformations from standard imaging modalities is a complex task even for expert clinicians. Categories for these malformations have been described by Sennaroğ L. *et al.* [16], and heuristics have been proposed to help identify them, such as the ones proposed by Dhanasingh A. *et al.* [7]. These heuristics are however of limited use to inexperienced otologists and ear, nose, and throat (ENT) surgeons, who, given the rarity of some of these conditions, cannot easily learn to detect the associated image patterns reliably. We take a first step towards assisting otologists and ENT surgeons in screening or detecting inner ear malformation by proposing the first automated method to detect these anomalies from clinical CT scans.

Different state-of-the-art deep learning methods have shown high performance for automatic detection of anomalies as presented in [5]. Deep learning approaches mostly based in convolutional neural networks used for classification have been used in a clinical context for anatomical anomalies [8]. Training such models requires large amounts of labeled medical data that faithfully represent these anomalies, which is challenging and expensive to acquire, especially because datasets are usually imbalanced because pathological cases are generally rare [17]. We propose a method that is trained uniquely on normative data for landmark location, which makes the approach suitable for adaptation to other anatomies. Knowledge of normal anatomical structural shapes and arrangements acquired during landmark location training brings implicit information for detecting anomalies within that region. Our method is based on multiple landmark location in CT scans of the inner ear. Because we aim to detect abnormalities indirectly by evaluating the output of the model, we define the landmark location as an object search problem and choose to use a deep reinforcement learning (DRL) architecture. We use both the communicative multiple agent reinforcement learning (C-MARL[11]) model and the standard multiple agent reinforcement learning (MARL[19]) model to locate a set of landmarks in the inner ear. We extract two pieces of critical information from these models: First, the variability of the predicted location of a certain landmark across different runs/agents which we evaluate in a subspace defined by the normative data landmarks after they are all aligned using Procrustes, and a principal component analysis (PCA) of the shape variation is performed to define the subspace as presented by López Diez *et al.* in [12]. Second, as a measurement of abnormality, we use the distribution of the predicted Q-values for each agent over the last ten states, including the final position where the landmark is placed. We initially test our approach using a small set of landmarks in a tight crop of the CT images centered on the cochlea versus synthetically generated images of a specific type of inner ear malformation called cochlear aplasia. Furthermore, we

tested the approach in real clinical data using a set of twelve landmarks in a bigger crop of the inner ear.

Simpler methods such as PCA can be employed for anomaly detection in physiological measurements [3, 10]. Several groups have used models trained on healthy anatomies to derived the detection of anomalies. While conceptually close to the approach we propose here, the methods have relied on spatial autoencoders or CycleGANs, as described by Baur C. *et al.* [2, 1], or segmentation models such as the Bayesian UNet used by Seebock *et al.* [15]. These approaches lack the spatial highlighting and interpretability that our landmark-based approach provides by using highly relevant points of interest defined according to the anatomical malformations.

2 Data

We use two different datasets to test our approach. Our first dataset consists of 119 clinical CT scanners from diverse imaging equipment. These images consist of a region of interest (ROI) with a size of ($32.1^3 mm^3$) with the cochlea in its center and an average voxel resolution of 0.3 mm. To test our approach, we synthetically generated abnormal inner ear CT scans from the original images by removing the cochlea (simulating cochlear aplasia) from the images, thus generating corresponding pairs of normal and abnormal CT scans with the same surrounding structures. The cochlea was segmented using ITK-SNAP software [20] and then replaced by Gaussian noise with mean and standard deviation estimated from the intensities of the tissue surrounding the segmentation [12]. An example of the transformation process as well as the location of the anatomical landmarks we use are shown in Figure 1. This dataset will be called the **Synthetic Set** from now on.

Our second dataset consists of 300 normal anatomy CT scans from heterogeneous sources and 123 CT scans of inner ears that present diverse congenital malformations. This unique dataset contains full-head CT images of CI patients acquired through different CT scanners. This dataset will be referred as the

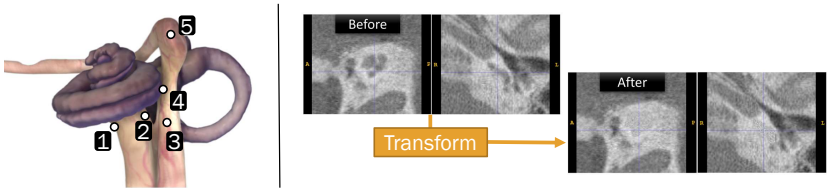

Fig. 1: **Left**: Set of landmarks used in the Synthetic Set. **1, 2** - Opposite sides of bony cochlear nerve canal in axial view. **3** - Facial Nerve (FN) exiting the Internal Acoustic Canal. **4** - Closest point of FN and cochlea. **5** - Geniculate ganglion of the FN. Edited from [18]. **Right**: Example image of CT scan from test set, before and after the synthetic image generation by inpainting.

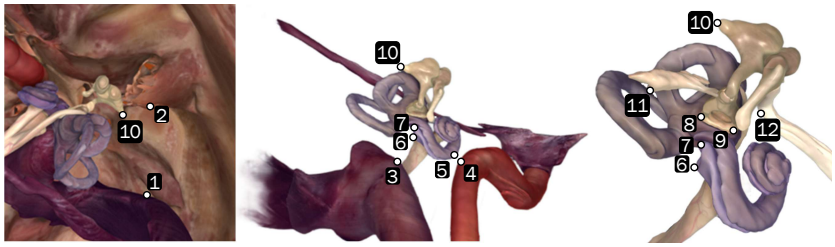

Fig. 2: Set of landmarks annotated in normal anatomy CT images in the Real Abnormality Set. **1** - Sigmoid Sinus (SS) (closest point to EAC). **2** - External Acoustic Canal (EAC) (closest point to SS). **3** - Jugular Bulb (closest point to Round Window (RW)). **4** - Carotid Artery(CA) (closest point to Basal turn of the cochlea). **5** - Basal Turn (closest point to JB). **6,7** - Anterior and posterior edges of RW. **8,9** - Anterior and posterior crus of staples. **10** - Short Process of Incus. **11** - Pyramidal Process **12**- Cochleariform Process. Edited from [18].

Real Abnormality Set further on. Out of the 300 normal ears, 175 were manually annotated by an expert with 12 landmarks that define key points of this anatomy. To optimally characterize certain points of interest, these landmarks were designed in close collaboration with our clinical partner, an ENT surgeon specialized in CI therapy in abnormal anatomies. These landmarks are presented in Figure 2. Simultaneously, the same ROI of $80^3 mm^3$ was extracted from the full-head CTs by using the location of the mandible joint and the beginning of the internal acoustic canal for both normal and abnormal anatomies. All images were re-sampled to a 0.5 mm isotropic resolution.

3 Methods

DRL for landmark location. Deep-Q-Networks [14] are used to find the optimal strategy for agents to reach their goal. These agents navigate through the 3D image (environment) and observe their state, which is defined as a patch of the image centered on the agent location. This patch becomes smaller as the agent gets closer to the landmark (multi-scale). Based on the observed state, the agent performs one action from the action set (move up, down, left, right, forward, and backward) and receives a reward, which is a function of the Euclidean distance between the current position of the agent and the previous position 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. The architecture of the Deep-Q-Network used for landmark location resembles a typical image classification architecture, but with a set of fully connected layers for each agent. The architecture of the model is shown in Figure 3. The common convolutional neural network weights among all agents provide implicit commu-

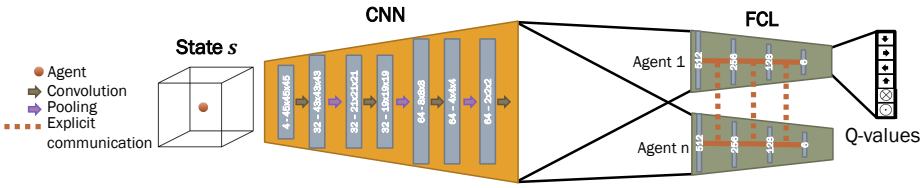

Fig. 3: Diagram of the DRL model used. The explicit communication connections are used in the C-MARL model, but not in the MARL model.

nication between the agents, meaning they share the same layers responsible for 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 predicted Q-value of each agent. This setup has been proven especially effective when the different landmarks present a consistent spatial correlation as it is the case with inner ear anatomy[13]. We trained a C-MARL in the normal anatomies of the Synthetic Set and the Real Abnormality Set. Finally we also trained the MARL model on the normal samples of the Real Abnormality Set as we expect that the explicit communication between agents might influence the variability of the model output when facing an abnormal anatomy. The training configuration employed is the same as presented in [13].

PCA shape distance method. The defined landmarks are placed in a spatial configuration which reveals consistency between patients with normal anatomy. We expect that for abnormal cases the landmarks predicted by the model will deviate significantly from this configuration and from one another. In order to test if a case is within the normal configuration, a point distribution model (PDM) is constructed following the approach presented in [6]. We will refer to a full set of landmarks in an image as *shape* where we know there is a point correspondence across all shapes in the training data. The alignment between all the annotated landmarks in normal anatomy is derived using Procrustes analysis [9]. Using this transform, we obtain a PDM that describes the shape variation only and that is invariant to size variation. Once the training shapes are all aligned, a mean shape is computed $\overline{\mathbf{x}}$, followed by a PCA of the shape variation [6]. The outcome of the PCA analysis is a set of principal components concatenated into a matrix Φ, which describes the modes of shape variation. A new shape $\mathbf{x}_{\text{new}}$ can be then defined as: $\mathbf{x}_{\text{new}} = \overline{\mathbf{x}} + \Phi\mathbf{b}$. The vector $\mathbf{b}$ denotes weights controlling the modes of shape variation and Φ contains the first t principal components. We chose to use a t value such that 90% of the shape variability is contained in the Φ matrix. We found $t = 11$ for the 36-dimension space defined by the twelve 3d-landmarks from Figure 2 over the 175 annotated normal anatomy images of the Real Abnormality Set and $t = 6$ for the 92 normal anatomy images annotated with five landmarks described in Figure 1 from the Synthetic Set. 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 vector $\mathbf{b}$ describes the shape in terms of coordinates in the PCA space. In this space we evaluate the distance between the different shapes predicted by the model. We then define the distance $d_{ji} = ||b_i - b_j||_2$ which measures the variation of all the different shapes predicted for a certain image. Finally we compute the standard deviation of this distribution of distance values for a certain image, D_{image}, which measures the level of agreement among the multiple predictions computed in the PCA space defined by normative shapes. A sketch of this approach is shown in Figure 4a.

Q-value history distribution method. Using deep Q-learning means that the network is trained to estimate the Q-value, or estimated reward, of taking a certain action given a certain state. Our hypothesis is that if the current state of the agent that is looking for a certain landmark resembles the normal anatomical configuration of such a region, the Q-values will present a uniform distribution, as the agent should not expect a high reward for moving in a certain direction. On the other hand, when the anatomy of the state does not look like what the agent would expect, the Q-values should be less uniformly distributed, pushing the agent to move in a certain direction. To test this hypothesis, we have computed a measurement of the variability within the distribution of the predicted Q-values of the action set. To compute this abnormality measurement we collect the buffer with the predicted Q-values of the last 10 states of the agent, which have empirically be found sufficient to define the later states of the landmark search procedure. These Q-value vectors are then normalized for each agent and merged together with the Q-values of the different runs for each specific landmark. Then, the standard deviation of the Q-values distribution is computed for each landmark u_n. These uncertainty measurements are then joined together into a single value per image $U_{\text{image}} = \sqrt{\sum_n u_n^2}$. An overview of the process is outlined in Figure 4b.

The combination of both methods has been computed to evaluate its joint performance. Due to the different magnitude of the measurements of each method, a weighting factor has been included in the combination so both methods have a more balanced contribution. The weighting factor w is defined as $w = \frac{\text{median}(D_{\text{training}})}{\text{median}(U_{\text{training}})}$. Then the combination of both methods is defined as $C_{\text{image}} = \sqrt{D_{\text{image}}^2 + (wU_{\text{image}})^2}$ to analyze the joint performance.

4 Results

Each of our tests has been evaluated over five runs for a more rigorous analysis. Both methods described in the previous section where initially tested in the Synthetic Set. The C-MARL model was trained on 92 images with three agents per landmark for the five landmarks shown in Figure 1. The average error of landmark location is 0.814mm in the test set. Our method is tested on 27 normal anatomy CT crops and their corresponding artificially created abnormal anatomy scans. The results are shown in Figure 5 a) and d).

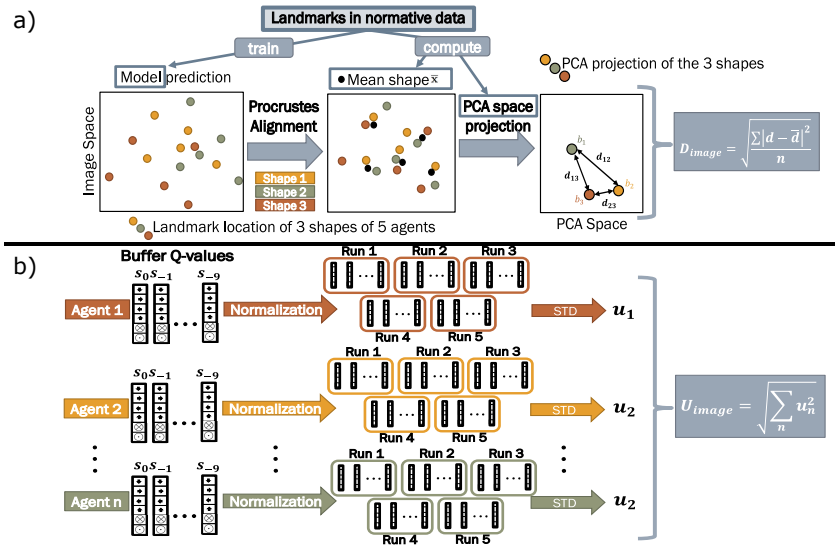

Fig. 4: Diagram of the computation process **a)** D_{image} **b)** U_{image} computation process.

To test the performance in the Real Abnormality Set, we trained a C-MARL model with one agent per landmark for the twelve landmarks described in Figure 2. The model was trained on 150 CT scans of patients with normal inner ear anatomy with an average test error of 1.74 mm for landmark location over the 25 images of the test set. To test our method we used 123 other different CT scans of normal anatomy and 123 of congenital malformations from the Real Abnormality Set, over five runs. The results of our method are shown in Figure 5 b) and e). As was expected, there is a significant drop in performance when comparing the results in the smaller ROI of the Synthetic Set shown in Figure 5 d) and the results with the same architecture (with a different set of landmarks) shown in Figure 5 e). However, we expected our method would benefit from using the MARL model, which does not include the connections that are responsible for the explicit communication between agents. This means that the agents would not share explicit information about their location and search procedure while looking for the landmarks. This makes agents more independent from each other and less tied to the spatial correlation among them. Our hypothesis is that avoiding this communication will derive greater values for the abnormality measurements when facing an anomaly. The MARL model was trained in the exact same configuration as the C-MARL model and obtained an average error of 1.99 mm on landmark location accuracy. The results of applying our method in this configuration can be observed in Figure 5 c) and f). It can be observed that the method does indeed perform better without explicit communication connections in the model.

The combination of both methods C_{image} shows an improved performance for the Synthetic Set as shown in Figure 5 d), and a very close performance to the best-performing method for the Real Abnormality Set see Figure 5 e) and f). We consider that the combination should be used as a more stable measurement which shows an area under the curve (AUC) of 0.96 for the artificial dataset and 0.86 for the large clinical dataset.

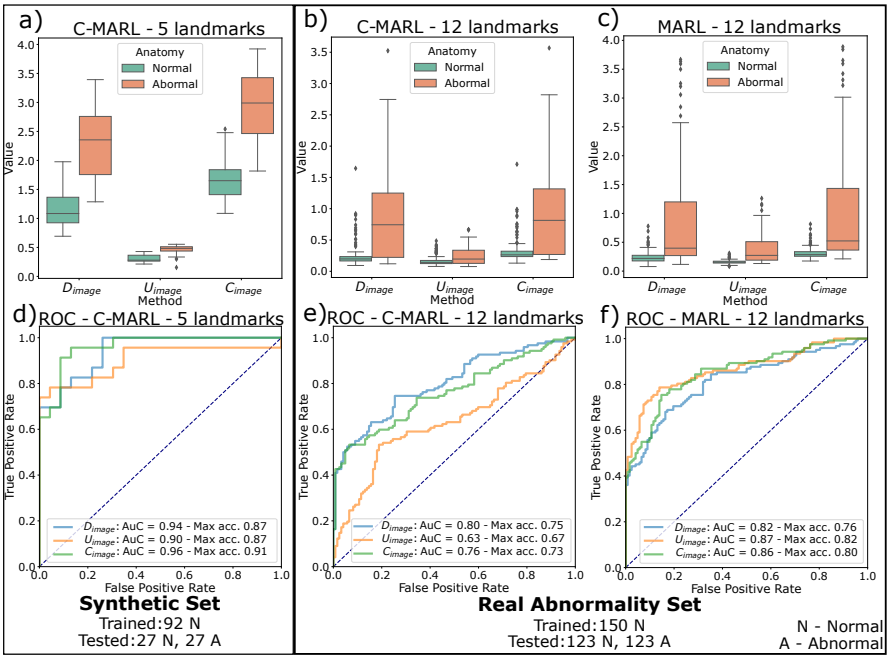

Fig. 5: Evaluation of the different methods over five runs for normal and abnormal anatomies. **Right:** Abnormality measurements distribution. **Left:** Derived ROC curves. **a and b:** 5 landmarks with C-MARL model in artificial data. **c and d:** 12 landmarks with C-MARL model in clinical data. **e and f:** 12 landmarks with MARL model in clinical data.

5 Conclusion

We have demonstrated that we can detect abnormal inner ear anatomies by solely training a DRL model on normative data and evaluating the output variability of certain implicit information. This information looks at the relative position of the predicted landmarks over different runs/agents in a subspace defined by the normative annotations as well as the distribution of the Q-values of the last iterations of the agents as a measurement of the uncertainty of the final

location. Our MARL model achieved the best performance with an AUC of 0.87 in clinical data which is a high score for such a complex classification problem. We showed how uncertainty information can be derived from a trained model to automatically detect abnormal anatomies, meaning no specific classification model needs to be trained, and therefore annotated abnormal data are not required to build the framework. We proved that the stated methods provide a measurement of the abnormality of the model's output which is linked with the presence of malformations. We examined the approach with good results, not only on artificially generated data, but also in a large dataset of real clinical CT scans of patients with diverse inner ear malformations.

References

1. Baur, C., Graf, R., Wiestler, B., Albarqouni, S., Navab, N.: Steganomaly: Inhibiting cyclegan steganography for unsupervised anomaly detection in brain mri. In: Martel, A.L., Abolmaesumi, P., Stoyanov, D., Mateus, D., Zuluaga, M.A., Zhou, S.K., Racoceanu, D., Joskowicz, L. (eds.) Medical Image Computing and Computer Assisted Intervention – MICCAI 2020. pp. 718–727. Springer International Publishing, Cham (2020)
2. Baur, C., Wiestler, B., Muehlau, M., Zimmer, C., Navab, N., Albarqouni, S.: Modeling healthy anatomy with artificial intelligence for unsupervised anomaly detection in brain mri. Radiology: Artificial Intelligence **3**(3), e190169 (2021). https://doi.org/10.1148/ryai.2021190169, https://doi.org/10.1148/ryai.2021190169
3. Ben Amor, L., Lahyani, I., Jmaiel, M.: PCA-based multivariate anomaly detection in mobile healthcare applications. In: Proc. International Symposium on Distributed Simulation and Real Time Applications (DS-RT). pp. 1–8 (2017). https://doi.org/10.1109/DISTRA.2017.8167682
4. Cairo/EG, R.Z.: Congenital inner ear abnormalities:a practical review. EPOS ECR 2019 / C-1911 . https://doi.org/10.26044/ecr2019/C-1911, https://dx.doi.org/10.26044/ecr2019/C-1911
5. Chalapathy, R., Chawla, S.: Deep learning for anomaly detection: A survey (1 2019), http://arxiv.org/abs/1901.03407
6. Cootes, T.F., Taylor, C.J., Cooper, D.H., Graham, J.: Active shape models-their training and application. Computer vision and image understanding **61**(1), 38–59 (1995). https://doi.org/https://doi.org/10.1006/cviu.1995.1004
7. Dhanasingh, A., Erpenbeck, D., Assadi, M.Z., Úna Doyle, Roland, P., Hagr, A., Rompaey, V.V., de Heyning, P.V.: A novel method of identifying inner ear malformation types by pattern recognition in the mid modiolar section. Scientific Reports **11** (12 2021). https://doi.org/10.1038/s41598-021-00330-6
8. Gill, R.S., Hong, S.J., Fadaie, F., Caldairou, B., Bernhardt, B.C., Barba, C., Brandt, A., Coelho, V.C., d'Incerti, L., Lenge, M., Semmelroch, M., Bartolomei, F., Cendes, F., Deleo, F., Guerrini, R., Guye, M., Jackson, G., Schulze-Bonhage, A., Mansi, T., Bernasconi, N., Bernasconi, A.: Deep convolutional networks for automated detection of epileptogenic brain malformations. In: Frangi, A.F., Schnabel, J.A., Davatzikos, C., Alberola-López, C., Fichtinger, G. (eds.) Proc. MICCAI. pp. 490–497. Springer (2018). https://doi.org/10.1007/978-3-030-00931-1_56
9. Gower, J.C.: Generalized procrustes analysis. Psychometrika **40**(1), 33–51 (1975). https://doi.org/10.1007/bf02291478

10. Krenn, V.A., Fornai, C., Webb, N.M., Woodert, M.A., Prosch, H., Haeusler, M.: The morphological consequences of segmentation anomalies in the human sacrum. American Journal of Biological Anthropology (12 2021). https://doi.org/10.1002/ajpa.24466, https://onlinelibrary.wiley.com/doi/10.1002/ajpa.24466
11. Leroy, G., Rueckert, D., Alansary, A.: Communicative reinforcement learning agents for landmark detection in brain images. In: Machine Learning in Clinical Neuroimaging and Radiogenomics in Neuro-oncology. pp. 177–186. Springer (2020). https://doi.org/10.1007/978-3-030-66843-3_18
12. López Diez, P., Juhl, K.A., Sundgaard, J.V., Diab, H., Margeta, J., Patou, F., Paulsen, R.R.: Deep reinforcement learning for detection of abnormal anatomies. Proceedings of the Northern Lights Deep Learning Workshop **3** (3 2022). https://doi.org/10.7557/18.6280
13. López Diez, P., Sundgaard, J.V., Patou, F., Margeta, J., Paulsen, R.R.: Facial and cochlear nerves characterization using deep reinforcement learning for landmark detection. In: Proc. MICCAI. pp. 519–528. Springer (2021). https://doi.org/10.1007/978-3-030-87202-1_50
14. Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A.A., Veness, J., Bellemare, M.G., Graves, A., Riedmiller, M., Fidjeland, A.K., Ostrovski, G., Petersen, S., Beattie, C., Sadik, A., Antonoglou, I., King, H., Kumaran, D., Wierstra, D., Legg, S., Hassabis, D.: Human-level control through deep reinforcement learning. Nature **518**, 529–533 (2 2015)
15. Seeböck, P., Orlando, J.I., Schlegl, T., Waldstein, S.M., Bogunović, H., Klimscha, S., Langs, G., Schmidt-Erfurth, U.: Exploiting epistemic uncertainty of anatomy segmentation for anomaly detection in retinal oct. IEEE Transactions on Medical Imaging **39**(1), 87–98 (2020). https://doi.org/10.1109/TMI.2019.2919951
16. Sennaroğlu, L., Bajin, M.D.: Classification and current management of inner ear malformations. Balkan Medical Journal **34** (2017). https://doi.org/10.4274/balkanmedj.2017.0367
17. Shin, H.C., Tenenholtz, N.A., Rogers, J.K., Schwarz, C.G., Senjem, M.L., Gunter, J.L., Andriole, K.P., Michalski, M.: Medical image synthesis for data augmentation and anonymization using generative adversarial networks. In: Gooya, A., Goksel, O., Oguz, I., Burgos, N. (eds.) Simulation and Synthesis in Medical Imaging. pp. 1–11. Springer International Publishing, Cham (2018)
18. Trier, P., Østergaard Noe, K., Sørensen, M.S., Mosegaard, J.: The visible ear surgery simulator. vol. 132 (2008)
19. Vlontzos, A., Alansary, A., Kamnitsas, K., Rueckert, D., Kainz, B.: Multiple Landmark Detection Using Multi-agent Reinforcement Learning. In: Proc. MICCAI. Springer (2019). https://doi.org/10.1007/978-3-030-32251-9_29
20. Yushkevich, P.A., Piven, J., Cody Hazlett, H., Gimpel Smith, R., Ho, S., Gee, J.C., Gerig, G.: User-guided 3D active contour segmentation of anatomical structures: Significantly improved efficiency and reliability. Neuroimage **31**(3), 1116–1128 (2006). https://doi.org/10.1016/j.neuroimage.2006.01.015