

Using machine learning on MRI radiomics to diagnose parotid tumours before comparing performance with radiologists: a pilot study.

Samy Ammari^{1,2†}, Arnaud Quillent^{3†}, Víctor Elvira⁴, François Bidault^{1,2}, Gabriel C.T.E. Garcia², Dana M. Hartl⁵, Corinne Balleyguier^{1,2}, Nathalie Lassau^{1,2}, Émilie Chouzenoux^{3*}

¹ *Biomaps, UMR1281 INSERM, CEA, CNRS, Université Paris-Saclay, 94805 Villejuif, France*

² *Department of Imaging, Gustave Roussy Cancer Campus, Université Paris Saclay, 94805 Villejuif, France*

³ *Centre de Vision Numérique, OPIS, CentraleSupélec, Inria, Université Paris-Saclay, 91190 Gif-sur-Yvette, France*

⁴ *School of Mathematics, University of Edinburgh, Edinburgh EH9 3FD, United Kingdom*

⁵ *Department of Otolaryngology Head and Neck Surgery, Gustave Roussy Cancer Campus, Université Paris Saclay, 94805 Villejuif, France*

*Corresponding author: emilie.chouzenoux@inria.fr

† These authors contributed equally to this work

Keywords

Radiomics — Machine learning — Parotid glands — AI benefit analysis

This version of the article has been accepted for publication, after peer review and is subject to Springer Nature's AM terms of use, but is not the Version of Record and does not reflect post-acceptance improvements, or any corrections. The Version of Record is available online at: <http://dx.doi.org/10.1007/s10278-024-01255-y>.

Abstract

The parotid glands are the largest of the major salivary glands. They can harbour both benign and malignant tumours. Preoperative work-up relies on MR images and fine needle aspiration biopsy, but these diagnostic tools have low sensitivity and specificity, often leading to surgery for diagnostic purposes. The aim of this paper is (1) to develop a machine learning algorithm based on MR images characteristics to automatically classify parotid gland tumours and (2) compare its results with the diagnoses of junior and senior radiologists in order to evaluate its utility in routine practice. While automatic algorithms applied to parotid tumours classification have been developed in the past, we believe that our study is one of the first to leverage four different MRI sequences and propose a comparison with clinicians. In this study, we leverage data coming from a cohort of 134 patients treated for benign or malignant parotid tumours. Using radiomics extracted from the MR images of the gland, we train a random forest and a logistic regression to predict the corresponding histopathological subtypes. On the test set, the best results are given by the random forest: we obtain a 0.720 accuracy, a 0.860 specificity and a 0.720 sensitivity over all histopathological subtypes, with an average AUC of 0.838. When considering the discrimination between benign and malignant tumours, the algorithm results in a 0.760 accuracy and a 0.769 AUC, both on test set. Moreover, the clinical experiment shows that our model helps to improve diagnostic abilities of junior radiologists as their sensitivity and accuracy raised by 6 % when using our proposed method. This algorithm may be useful for training of physicians. Radiomics with a machine learning algorithm may help improve discrimination between benign and malignant parotid tumours, decreasing the need for diagnostic surgery. Further studies are warranted to validate our algorithm for routine use.

Introduction

Parotid gland tumours are benign in 70 to 80 % of cases [1, 2]. Preoperative work-up relies on MRI and fine needle aspiration cytology (FNAC), but these diagnostic tools lack sensitivity, often leading to surgery for diagnostic purposes. The sensitivity of FNAC for parotid tumours using the recommended Milan system ranges from 75 to 90 % [3]. Because of the heterogeneity of parotid gland tumours and their possible location in the deep lobe, FNAC is not always sufficient to come up with a correct diagnosis [4, 5]. Moreover, diagnostic parotid surgery carries a high risk of morbidity due to the fact that the facial nerve passes through the gland and dissection of the nerve can lead to physical and aesthetical discomfort, as well as temporary or permanent facial paralysis.

With the development of computer science, new automatic data processing techniques and machine learning can now be leveraged to achieve numerous tasks in the healthcare field [6, 7] such as image classification and segmentation, with some applications to parotid glands [8, 9].

Radiomics are a field of radio-diagnostics which first appeared in the 2010s. It consists in computing several statistics from a segmented area of a medical image such as its shape, volume, texture, intensity, or other high-order moments [10, 11]. Radiomics aim at extracting patterns in the data that would be invisible to the human eye. Combined with machine learning approaches, radiomics have been used by many researchers to automatically analyse MR and CT images [12, 13].

The goal of our study is to (1) introduce a machine learning algorithm for the discrimination of parotid gland tumours into their respective histopathological subtypes as defined by the World Health Organization (WHO) [14]. Previous works were conducted on this matter, highlighting the efficiency of such kind of approach [15–19]. However, they only used at most two MRI sequences, whereas we decide to sample four to observe whether it improves automatic diagnosis performances following the conclusions of [20]. Moreover, (2) we perform a comparative analysis of our proposed algorithm with radiologists in order to evaluate benefits of machine learning algorithms combined to relevant experts' knowledge. To the best of our knowledge, this kind of comparison has never been done before for parotid tumours.

To conduct our research, we create a new dataset coming from a cancer treatment centre. As gathered from an expert hospital, we believe those data encompass very germane information about parotid cancers. On top of that, as all patients coming to the centre have previously been examined in community hospitals, we have at our disposal large collections of images that are not sampled from regular screening, thus resulting in a greater proportion of diagnosed parotid cancer than observed in the whole population.

This research article is organised as follows. [section 1, Materials and methods](#) introduces our data sampling strategy, the subsequent radiomics extraction, the development of the machine learning algorithm and finally the implemented protocol for comparison with radiologists. Our findings are then presented in [section 2, Results](#) and further analysed in [section 3, Discussion](#).

Materials and methods

1.1 Data acquisition

Patient selection The cohort includes patients that underwent an MRI examination for parotid tumours at Gustave Roussy Cancer Campus between 2012 and 2021. All the patients subsequently underwent surgery, with histopathologic examination performed in the same centre. Our cohort includes patients whose diagnosis is definitive, and whose MRI examinations comprise T1-weighted (T1w), contrast-enhanced T1-weighted (T1ce), T2-weighted (T2w) and diffusion-weighted (DWI) sequences (see [Figure 1](#)). These four sequences are standard in MR imaging. Note that we did not use ADC (Apparent Diffusion Coefficient) as it is redundant with DWI, since both sequences are compiled from the diffusion of water molecules in the tissues. We categorise the histopathological subtypes into pleomorphic adenomas (benign), Warthin's tumour (benign) and carcinomas (malignant). Tumours that do not belong to the previous categories are reported as 'Other type'. Descriptive statistics are shown in [Table 1](#). As the patients were hospitalised in an institution specialised in cancer treatment, malignant tumours are over-represented in the studied cohort compared to their frequency in the worldwide population.

This research project was approved by the appropriate institutional board following the General Data Protection Regulation (GDPR) and was declared to the competent national administrative bodies, namely the Health Data Hub and the CNIL. All patients were informed by post and provided their consent to the use of their data within the scope of this study.

MRI protocol Acquisition was performed on three different devices over the study period: a 1.5 T Optima MR450w, a 3 T Discovery MR750w and a 3 T Signa Premier (GE HealthCare, Milwaukee, USA). Corresponding MRI parameters are reported in [Table 2](#). Patients were placed in a supine position, with a head and neck antenna. The contrast agent used was Gadoterate meglumine (Dotarem, Guerbet, Villepinte, France). Resulting MR images were stored in DICOM files and encoded in 24-bit.

Data cleaning To ensure robustness of our machine learning approach, we rely on several data cleaning steps. First, as lymphomas are under-represented, they were removed from the dataset: they were too few for the machine learning algorithm to learn to classify them. Tumours that are described as 'Other type' were too heterogeneous, as expected, to be gathered into one meaningfully consistent class, and too few to be broken down into subgroups. Hence, they were also removed from the dataset. From the 134 patients of the whole cohort, only 111 were kept for the study, after this data cleaning procedure (see [Figure 1](#)).

1.2 Image analysis

Image pre-processing As MR images are very heterogeneous between patients and acquisition devices, a pre-processing step of harmonisation is often performed in order to compensate the high variability of the data. In many radiomics studies, a first normalisation is applied at pixel-level on the images [21, 22]. However, some patients in our cohort were examined with 3 T MRI devices and others with a 1.5 T machine with lower resolution: normalising the images might have resulted in a loss of information. Consequently, no normalisation was performed in our study.

Characteristics	Value
Age	Years
Mean	63.7
Median	62
Minimum	18
Maximum	92
Sex	n (%)
Male	76 (57)
Female	58 (43)
MRI	n (%)
1.5 T	36 (27)
3 T	98 (73)
Tumour status	n (%)
Benign	73 (54)
Malignant	61 (46)
Tumour histopathological subtype	n (%)
Pleomorphic adenoma	36 (27)
Warthin's tumour	27 (20)
Lymphoma	4 (3)
Carcinoma	48 (36)
Other	19 (14)

Table 1. Statistics about the studied cohort.

Tumour segmentation Parotid gland tumours were manually segmented by research engineers using the software Olea Sphere version 3.0.18 (Olea Medical, La Ciotat, France). All patients of the cohort present a single parotid gland tumour except for one who is affected by two tumours which were both used for this study. Each parotid gland tumour was segmented on the four different MRI sequences (see [section 1.1, MRI protocol](#)). Studied MR images are 3-dimensional, but we chose to segment 2D regions of interest (ROI) on the axial slice with the largest tumour surface area. This was done to better reflect the radiologists' behaviour in a clinical setting, as segmenting the whole parotid volumes would be too time-consuming. Within this slice, the entire tumour region is segmented, including its core and eventual necrosis. Figure 2 presents an example of such segmentation. Output ROIs were double-checked by a head and neck radiologist with 15 years of experience.

Features extraction We computed the radiomic features from all four considered sequences using Olea Sphere software with the following settings: 1 mm³ voxel resampling, 64 gray level bins to compute the histogram. A total of 108 radiomics per sequence are extracted, following the IBSI standard [23]: 16 shape-based features, 19 first-order features, 23 gray level

Device	Weighting	Sequence	TR	TE	Slice thickness
Optima MR450w, 1.5 T Installed in 2016, 70 cm tunnel, 32 channels, 40 cm z-axis FOV, gradient 40 mT/m, SR 200 T/m/s	Pre-contrast T1	Fast spin echo	532 ms	Minimal	3 mm
	Post-contrast T1	Fast spin echo	532 ms	Minimal	3 mm
	T2	Fast spin echo	9,430 ms	102 ms	5 mm
	DWI	Spin echo	2,310 ms	Minimal	4 mm
Discovery MR750w, 3 T Installed in 2012, 70 cm tunnel, 32 channels, 50 cm z-axis FOV, gradient 44 mT/m, SR 200 T/m/s	Pre-contrast T1	Fast spin echo	680 ms	10 ms	3 mm
	Post-contrast T1	Fast spin echo	700 ms	11 ms	3 mm
	T2	Fast spin echo	9,700 ms	90 ms	3 mm
	DWI	Spin echo	7,100 ms	73 ms	4 mm
Signa Premier, 3 T Installed in 2021, 70 cm tunnel, 32 channels, 50 cm z-axis FOV, gradient 80 mT/m, SR 200 T/m/s	Pre-contrast T1	Fast spin echo	500 ms	Minimal	3 mm
	Post-contrast T1	Fast spin echo	500 ms	Minimal	3 mm
	T2	PROPELLER	8,831 ms	120 ms	3 mm
	DWI	Spin echo	3,114 ms	60 ms	4 mm

Table 2. MRI devices parameters.

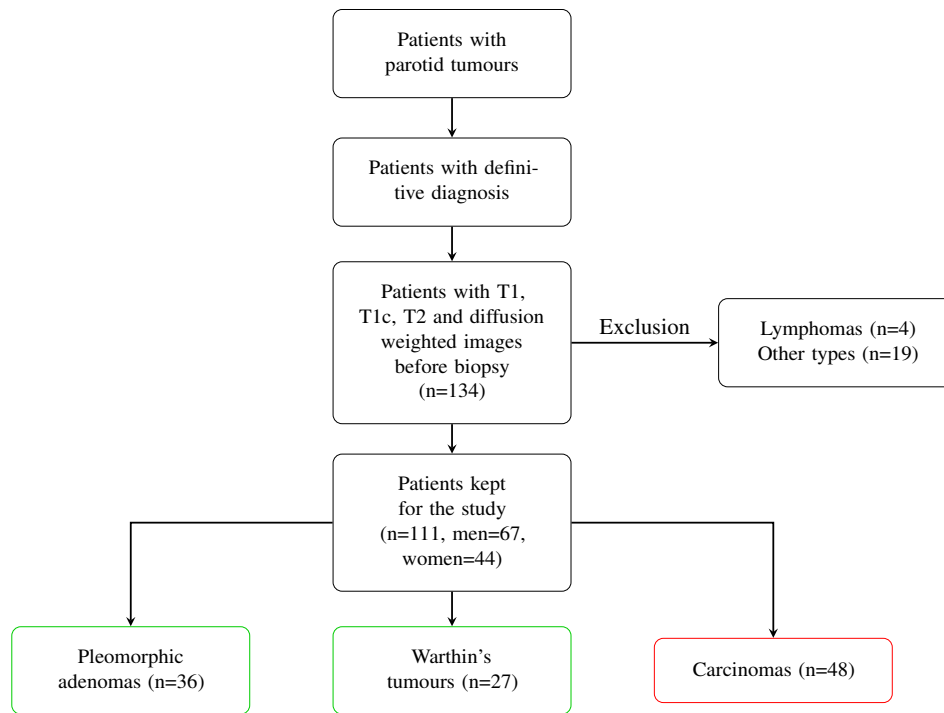


Figure 1. Flowchart showing how relevant data were chosen. The red box indicates malignant tumours, the green ones correspond to benign tumours.

co-occurrence matrix (GLCM) features, 16 gray level run length matrix (GLRLM) features, 15 gray level size zone matrix (GLSZM) features, 5 neighbouring gray tone difference matrix (NGTDM) features and 14 gray level dependence matrix (GLDM) features. Appendix A summarises the computed features. Hence, a single tumour is represented by $4 \times 108 = 432$ real numbers.

Radiomics are also subject to high heterogeneity. ComBat normalisation [24,25] can be used to reduce variability across devices, but several safety checks performed on our data suggests no significant improvement when applying this method to our problem. Hence, we decided not to perform any pixel-level normalisation.

From the medical records of the patients, we retrieve two phenotypic variables: age and sex are indeed often linked to different types of tumours [1,2]. Those two variables are then concatenated to the computed radiomic features and used as input to the machine learning system.

1.3 Machine learning task

Our goal is to discriminate parotid gland tumours according to their histopathological subtype. This corresponds to a classification task, where each type of tumour is a different class, and inputs are the radiomic features.

Separation into training and test sets The data are divided into a training group and a test group, composed of respectively 86 and 25 patients. The training group comprises data acquired from both 1.5 T Optima MR450w and 3 T MR750w whereas the test group, which was sampled after the new 3 T Signa Premier was put into operation, includes data from the three devices. This was motivated by [15] which uses different MRI devices for training and validation, but keeps the same magnetic field strength for both. Such a set-up allows us to verify the generalisation to new MRI systems of similar power while avoiding significant bias. Class distribution of the whole dataset is preserved in a stratified manner while creating both training and test sets: the frequency of appearance of each histopathological subtype is preserved.

Model building Random forests are more robust than other machine learning models when dealing with high-dimensional feature sets, and are also more polyvalent [26,27]. Moreover, this family of algorithms is quite often used to classify radiomics [12,28,29]. They are less prone to overfitting [30], which is likely to occur due to the small size of our dataset. Hence, we decide to use this kind of model in our work and compare our results to a baseline logistic regression. In order to mitigate the imbalanced class repartition, classes are given different weights to increase the influence of the least populated ones.




Figure 2. Examples of segmentation on the four acquired sequences: T1w, T1ce, DWI and T2w (clockwise).

Model training A five-fold cross-validation is employed to estimate the model’s performances while keeping an 80:20 ratio between training and test samples. This method prevents overfitting on small datasets [31, 32]. In addition, we use a grid search to identify the best hyperparameters of the random forest and the logistic regression. To further avoid overfitting, Z-score standardisation and parameters tuning are performed inside the cross-validation loop, as recommended in [32]. The random forest parameters we analysed are maximum depth of the trees, splitting criterion and associated splitting method, number of features input to the trees, number of samples in a leaf and number of samples to reach before splitting a node. Regarding the logistic regression, best regularisation strength and solver were searched. In addition, fine-tuned weights were attributed to the classes for both models. We keep track of five metrics to assess the performance throughout our parameters search procedure: balanced accuracy (takes into account class imbalance), accuracy, sensitivity, specificity and area under the ROC curve (AUC). Definitions of those metrics are given in Appendix B, following naming conventions used by Scikit-learn¹.

Feature selection Feature selection is applied inside the cross-validation loop performed on the training set to compensate for the large number of features acquired and the relatively small number of patients selected (“curse of dimensionality”). Radiomics are often strongly correlated due to the high number of tweakable parameters used for their computation. They are very generic and can be used on various problems, but some of these features may be redundant for certain applications and not for others [33]. A selection of a feature subset used to train the algorithm was therefore required to avoid overfitting [34]. An F-test-based filter method was preferred because of its ability to sort out variables that are independent of the target while keeping a low computational cost [34, 35]. As part of our approach, we tested several values for the number of features to retain using the previously mentioned cross-validation procedure. The set of the best features is then used to subsample the test set and evaluate the metrics.

1.4 Comparison to the analysis of radiologists

Once we have the best performing model, we are able to compare the machine learning algorithm to radiologists’ analyses of the same image sets. First, nine patients were randomly sampled from the test set, keeping the class distribution intact: 3 pleomorphic adenomas, 2 Warthin’s tumours and 4 carcinomas. Seven in-training radiologists (‘juniors’) that just attended a lesson on parotid gland tumours and two professionals (‘seniors’ with 5 and 15 years of experience respectively) participated in the study. For each patient whose data were used in the experiment, all nine radiologists were shown the same four MRI sequences as the ones used to extract the radiomic features (*i.e.*, T1w, T2w, T1ce and DWI sequences). Following this procedure ensured that both radiologists and machine learning model were using the same input information to diagnose the patients. The images were displayed to physicians on a same screen, using a single software with fixed settings for every viewing. Directly after making their diagnosis, the radiologists were shown the algorithm’s prediction and were asked whether they would change their mind given the machine’s diagnosis. The true performance of the algorithm was never disclosed during the experiment.

Results

2.1 Machine learning algorithm

Type classification Among all the combinations of parameters tested, the best one is a random forest that reached a multiclass accuracy of 0.720 and a mean OvR (One-vs-Rest, *i.e.*, computed per class) AUC of 0.838 on the test set. The same model

¹<https://scikit-learn.org>

presented an OvR accuracy of 0.840 for both pleomorphic adenomas and Warthin’s tumours. For carcinomas, this score was 0.760. Remark that this latter case also corresponds to a benign vs malignant classification. Metrics results are available in Table 3, confusion matrix in Table 4 and ROC curves in Figure 3. The best model hyperparameters are as follows: 22 as maximum depth of trees, 3 as minimum number of samples in leaves, and 43 estimators. On average, the machine learning classifier has a worse sensitivity than specificity, which means that it is less able to tell a tumour belongs to a specific category than to tell it does not.

Class	Balanced accuracy	Accuracy	Specificity	Sensitivity	F1 score	OvR AUC
Random forest						
All classes	0.702	0.720	0.860	0.720	0.714	0.838
Pleomorphic adenomas	0.724	0.840	0.947	0.500	0.600	0.860
Warthin’s tumours	0.845	0.840	0.833	0.857	0.750	0.889
Carcinomas	0.760	0.760	0.769	0.750	0.750	0.769
Logistic regression						
All classes	0.560	0.600	0.800	0.500	0.592	0.698
Pleomorphic adenomas	0.697	0.800	0.895	0.500	0.545	0.746
Warthin’s tumours	0.687	0.800	0.944	0.429	0.545	0.690
Carcinomas	0.606	0.600	0.462	0.750	0.643	0.660

Table 3. Results of the two models on the test set for several metrics computed across all classes and by class. In this latter case, metrics are computed in an OvR fashion, considering the studied class as the positives. Metrics definitions are available in appendix B. Overall OvR AUC is the mean of each class score. Carcinomas line also corresponds to the benign/malignant classification. The performance of our approach is significantly better.

		Predicted diagnosis		
		Pleomorphic adenomas	Warthin’s tumours	Carcinomas
		Random forest		
Ground truth	Pleomorphic adenomas	3	1	2
	Warthin’s tumours	0	6	1
	Carcinomas	1	2	9
	Logistic regression			
	Pleomorphic adenomas	3	0	3
	Warthin’s tumours	0	3	4
	Carcinomas	2	1	9

Table 4. Confusion matrix of the test set: histopathological subtype classification.

Selected features As previously described in section 1.3, [Feature selection](#), we used an F-test to reduce the dimension of our feature space. This method computes p-values representing the importance of each feature in relationship to the desired output, which is the histopathological type in our case. For both models tested, five relevant features were selected during the training phase, all extracted from the T2w sequence. More information can be found in Table 5. Figure 4 shows the distribution of the five selected features as box plots. For each of these features, we can roughly discriminate the histopathological subtypes directly on the graph. This implies that those features are statistically different as per the associated class: there is a strong relationship between the selected features and the output class found by the algorithm.

2.2 Experiment with radiologists

Radiologists performances To assess the performances of the radiologists, we resort to the same metrics as described in section 1.3, [Model training](#), allowing us to compare their results to the ones of our best machine learning model, *i.e.*, the random forest. As we did not ask radiologists to evaluate their diagnoses with a probability estimate, AUC could not be computed. Table 6 and Figure 3 feature the results of the experiment. It is interesting to observe that senior radiologists seem to have worse performances than both the algorithm and the juniors. This can be explained by the fact that seniors use a lot more information than the four MRI sequences used in this study to make their diagnosis, like patient symptoms for instance. Radiomics cannot be extracted from such data, hence we decided not to use them in our study, since its purpose is to assess efficiency of those

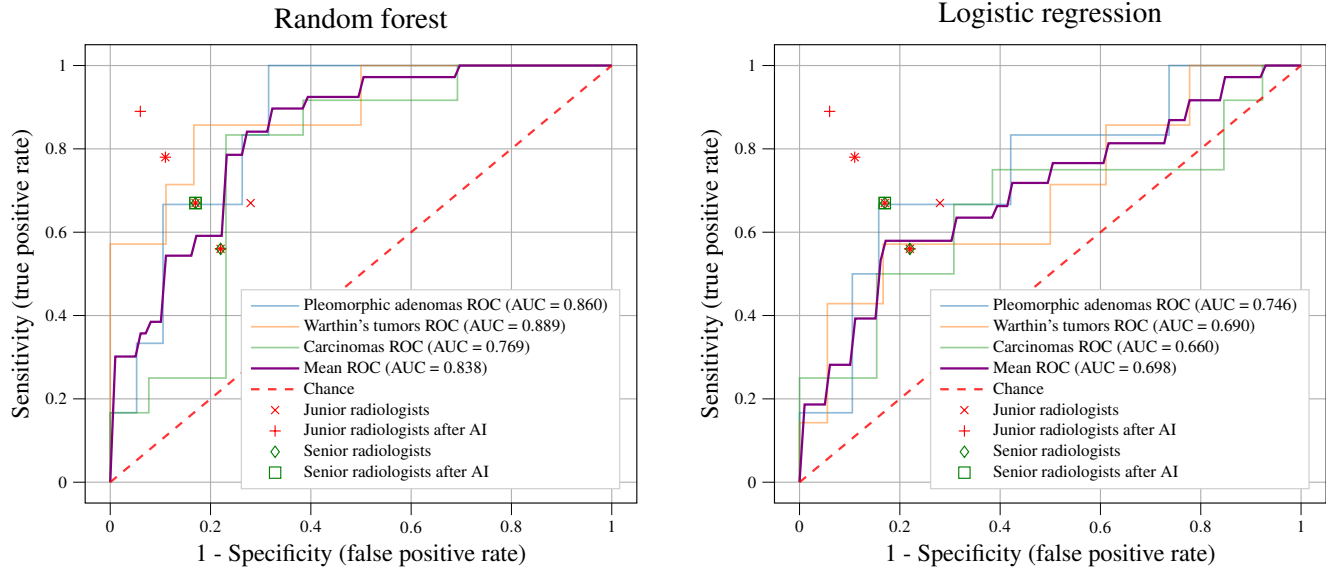


Figure 3. OvR ROC curves calculated from the algorithm predictions for each histopathological subtype class represented in the experiment, along with points corresponding to the metrics of the 9 radiologists. Some points are coinciding, resulting in only 6 of them being distinguishable on the graph (their corresponding overlapping symbols are displayed).

IBSI category	Feature name	F-test p-value
First Order	Skewness	3.88×10^{-5}
First Order	Median	5.83×10^{-5}
Gray Level Run Length Matrix	Long Run High Gray Level Emphasis	7.21×10^{-5}
First Order	Root Mean Squared	12.07×10^{-5}
Gray Level Co-occurrence Matrix	Autocorrelation	12.32×10^{-5}

Table 5. Relevant radiomics features selected during the pre-processing step.




Figure 4. Distribution of selected radiomics features values in the training set.

image statistics for tumour classification. Furthermore, experienced radiologists are used to see a wide variety of different tumours every day and thus have much more doubts than juniors that make use of their knowledge acquired in a lesson given just a few minutes before the experiment. When given access to complete patient files including other information than the four MRI sequences used for this study, radiologists' assessment of tumours is almost perfectly accurate.

Impact of the algorithm on metrics From Table 6, we can see that average performance of both groups improves similarly with the help of the algorithm: 8 % for balanced accuracy, 6 % for accuracy, 3 % for specificity and 6 % for sensitivity.

Candidate	Balanced accuracy	Accuracy	Specificity	Sensitivity
Random forest	0.806	0.778	0.889	0.778
J1	0.833 / 0.917 (+ 0.083)	0.778 / 0.889 (+ 0.11)	0.889 / 0.944 (+ 0.05)	0.778 / 0.889 (+ 0.11)
J2	0.722 / 0.722 (+ 0.000)	0.667 / 0.667 (+ 0.00)	0.833 / 0.833 (+ 0.00)	0.667 / 0.667 (+ 0.00)
J3	0.556 / 0.722 (+ 0.166)	0.667 / 0.778 (+ 0.11)	0.833 / 0.889 (+ 0.06)	0.667 / 0.778 (+ 0.11)
J4	0.389 / 0.556 (+ 0.167)	0.444 / 0.556 (+ 0.12)	0.722 / 0.778 (+ 0.06)	0.444 / 0.556 (+ 0.12)
J5	0.722 / 0.889 (+ 0.167)	0.778 / 0.889 (+ 0.11)	0.889 / 0.944 (+ 0.05)	0.778 / 0.889 (+ 0.11)
J6	0.722 / 0.722 (+ 0.000)	0.667 / 0.667 (+ 0.00)	0.833 / 0.833 (+ 0.00)	0.667 / 0.667 (+ 0.00)
J7	0.556 / 0.556 (+ 0.000)	0.556 / 0.556 (+ 0.00)	0.778 / 0.778 (+ 0.00)	0.556 / 0.556 (+ 0.00)
Mean	0.643 / 0.726 (+ 0.083)	0.651 / 0.714 (+ 0.063)	0.825 / 0.857 (+ 0.032)	0.651 / 0.714 (+ 0.063)
S1	0.556 / 0.722 (0.166)	0.556 / 0.667 (+ 0.111)	0.778 / 0.833 (+ 0.055)	0.556 / 0.667 (+ 0.111)
S2	0.639 / 0.639 (+ 0.000)	0.667 / 0.667 (+ 0.000)	0.833 / 0.833 (+ 0.000)	0.667 / 0.667 (+ 0.000)
Mean	0.597 / 0.681 (+ 0.084)	0.611 / 0.667 (+ 0.056)	0.806 / 0.833 (+ 0.027)	0.611 / 0.667 (+ 0.056)

Table 6. Performance metrics of junior and senior radiologists on 9 sub-sampled patients for the histopathological subtype classification task, first without, then with the help of our machine learning algorithm (difference between brackets). Results are given for each participant to the experiment. J1 to J7 are junior radiologists, S1 and S2 are seniors.

Table 7 shows results of the same experiment but class-wise. In this paragraph, we will consider balanced accuracy as reference metrics. On the subset, our random forest algorithm performs best on Warthin's tumours, which is consistent with the results on the test set. Juniors benefit from the machine learning model for all histopathological subtypes. A sensitivity gain of 21.5 % is obtained on Warthin's tumours diagnosis when using the algorithm, which results in a nearly 12 % increase of corresponding balanced accuracy. Seniors profit from a rise of specificity for both Warthin's tumours and carcinomas (25 % and 10 % respectively) when trusting the model. However, they decided not to follow the machine's predictions concerning pleomorphic adenomas, resulting in a stagnancy of corresponding metrics. The significance of these results has been verified with the help of a Student's t-test (p-value < 0.05) by estimating the variance of each metric thanks to a leave-one-out procedure.

Histopathological subtype	Balanced accuracy	Accuracy	Specificity	Sensitivity
Random forest				
Pleomorphic adenomas	0.834	0.889	1.000	0.667
Warthin's tumours	0.929	0.889	0.857	1.000
Carcinomas	0.775	0.778	0.800	0.750
Juniors				
Pleomorphic adenomas	0.750 / 0.774 (+ 0.024)	0.762 / 0.794 (+ 0.032)	0.786 / 0.833 (+ 0.047)	0.714 / 0.714 (+ 0.000)
Warthin's tumours	0.745 / 0.863 (+ 0.118)	0.841 / 0.905 (+ 0.064)	0.918 / 0.939 (+ 0.021)	0.571 / 0.786 (+ 0.215)
Carcinomas	0.693 / 0.725 (+ 0.032)	0.698 / 0.730 (+ 0.032)	0.743 / 0.771 (+ 0.028)	0.643 / 0.679 (+ 0.036)
Seniors				
Pleomorphic adenomas	0.709 / 0.709 (+ 0.000)	0.722 / 0.722 (+ 0.000)	0.750 / 0.750 (+ 0.000)	0.667 / 0.667 (+ 0.000)
Warthin's tumours	0.750 / 0.875 (+ 0.125)	0.889 / 0.944 (+ 0.055)	1.000 / 1.000 (+ 0.000)	0.500 / 0.750 (+ 0.250)
Carcinomas	0.613 / 0.663 (+ 0.050)	0.611 / 0.667 (+ 0.056)	0.600 / 0.700 (+ 0.100)	0.625 / 0.625 (+ 0.000)

Table 7. Performance metrics of junior and senior radiologists on 9 sub-sampled patients for the histopathological subtype classification task, first without, then with the help of our machine learning algorithm (difference between brackets). Results are given for each considered class.

Limited model influence on radiologists Predictions were followed by juniors in only 8 % of cases, but when they did, the algorithm was correct in 80 % of the cases. Moreover, when juniors kept their diagnosis despite the model estimations, they were mistaken 92 % of the time. As for senior radiologists, they only trusted the algorithm when they were hesitating between two diagnoses, and in 85 % of those cases, the decision was right. When seniors were mistaken in their diagnosis and the model suggested changing their mind, they never did it.

Discussion

3.1 Data selection

Most studies on parotid glands imaged with MR devices rely on T1w and T2w sequences with sometimes alterations (e.g., contrast enhancement or fat saturation) [15, 16, 20, 36]. Nevertheless, the use of DWI is rather rare, but studies such as [37] demonstrated that it could be useful for parotid gland tumours characterisation. Xia et al. in [20] observed that increasing the number of sampled sequences might output better results. In total, four sequences (T1w, T1ce, T2w and DWI) are given as input to our system. However, all the features automatically selected to compute the predictions of the proposed algorithm are extracted from T2w only. This nuances the statements made in those aforementioned works: it seems that the more different sequences we use, the more possible optimisation paths there are, though all sequences might not be useful to an algorithm to make reliable predictions.

Some works on parotid tumours [15, 16, 20, 36, 38] propose to gather all malignant tumours subtypes into a single class. Doing so would have included 23 new images of malignant tissues in our dataset and expanded its size by 20 %. Nonetheless, as medical treatment differs largely between these different types of malignant tumours, it is more helpful for our centre's radiologists to get precise insights about histopathological subtypes rather than broader malignancy assessment.

3.2 Segmentation procedure

Segmentation can be performed on 2D surfaces or 3D volumes. Here, we chose to compute 2D radiomic features on MR image slices with the largest tumour area. This choice, while it might not be as efficient as 3D features [12], spares doctors a very time-consuming task as segmentation would otherwise have to be carried out on each tumour slice. As an example, the segmentation and radiomics computation from the largest 2-dimension surface of a 6,038 mm³ tumour on all 4 considered sequences took us around 7 minutes, as opposed to 12 minutes using a 3D volume. Moreover, further feature analysis and selection can be manually performed as in [15, 16, 39]. Here, we preferred to embed these steps in our algorithm. This makes iterative fine-tuning of the process possible before inputting the features to our machine learning model. Furthermore, radiomics were computed on manually segmented parotid tumours, which can induce a great variability. An automatic segmentation approach (e.g., relying on deep learning) could be beneficial to address this issue, at the cost of a loss of interpretability [8, 9, 20, 36, 38].

3.3 Robustness

We used radiomics which are features whose main role is to sum up multiscale information contained in images. However, the robustness of this method may be dependent on the application, as stated in [33], but is hard to assess [40]. Also, the cohort enrolled in our research is not large enough to acquire a substantial amount of MR images. While we aspire to reduce the number of patients even more as it would mean that fewer persons are suffering from parotid tumours, insufficient data usually results in high variability and low generalisation capacities. Thus, further validation of our experiments in external centres should be carried out in order to guarantee the prospective abilities of our algorithm.

3.4 Efficiency of the method

Authors of [16, 38] observed that Warthin tumours are easier to discriminate than other histopathological subtypes. Table 3 illustrates the same findings, as our machine learning model gets higher accuracy for this kind of affliction. Junior and senior radiologists as well identify Warthin's tumours more accurately than other types (see Table 7) thanks to a high specificity.

Matsuo et al. [36] proposed to gather the information of both T1w and T2w into a single image considering them as two colour channels, thus creating pseudo-coloured pictures. This method allows to quickly compare signal intensities which is certainly a good take since Warthin's tumours are hyper-intense on T1w [41], while pleomorphic adenomas are hypo-intense on T1w and hyper-intense on T2w [42]. However, the whole process needs co-registered MR images, which were not available at the time of writing.

On top of that, it is rather unexpected that all features automatically selected with our approach come from T2w, while radiologists tend to prefer T1ce to make their diagnosis. Gabelloni et al. [16] however got appealing results using only T2w radiomics.

Eventually, we provided an experiment to assess the impact of our algorithm on radiologists' decisions. Nonetheless, it should be noted that since we compare a machine to humans, their respective training datasets are different. Indeed, experience of the physicians can be considered as a form of pre-training, but it conveys a bias as they consequently know about the

frequency of occurrence of each histopathological subtype. Moreover, radiomics given as input to our machine learning model do not encompass the same information as a radiologist's eyes. For instance, the algorithm has no data about the spatial location of the tumour within the head (e.g., deep lobe of the parotid gland, contact with surrounding nerves, etc.).

3.5 Clinicians and AI

Regarding the trust placed in artificial intelligence (AI) by radiologists, our study is in line with previous findings. A survey conducted by the European Society of Radiology (ESR) in 2018 [43] suggested that radiologists tend to trust AI. They see it as a way to save time, allowing a significant improvement of interaction with their patients or other clinicians. As explained in [section 2.2, Experiment with radiologists](#), junior and senior radiologists participating in our experiment followed the predictions of our algorithm in the cases where they were uncertain: Tables 6 and 7 show that seeing the output of our model indeed helps radiologists, as it never mislead them in our experiment. When they follow the predictions of our model, the diagnosis accuracy improves. This may indicate a high degree of confidence in our algorithm from the radiologists who took part in the experiment. Nevertheless, according to the same research work from ESR, around a half of doctors think that patients do not consider AI reliable enough to make a diagnosis without confirmation from a physician. We then believe that junior radiologists would largely benefit from the use of machine learning during their training period. Indeed, it could help them to learn new ways of making their diagnosis, for instance focusing on previously overlooked patterns in MR images, while in the end improving their overall accuracy. This would as well constitute a good way to get used to these new techniques as part of their future job. Moreover, radiologists in general could employ such algorithms as a complement to their daily diagnosis routine.

Nonetheless, as we had limited time access to radiologists to perform our experiment, only nine of them could be enrolled: drawing definitive conclusions or perform meaningful statistical tests remains difficult.

Conclusion

In this research paper, we propose an easy workflow to predict the histopathological subtype of parotid gland tumours using radiomics. Thanks to the analysis of the results gave by the algorithm, we strongly believe that our method could be used as a mean of training junior radiologists. Moreover, the procedure we offer may also help physicians settle doubts when diagnosing parotid tumours.

Based upon the experience acquired while carrying out this research work, we would like to advocate some tactics to conduct future experiments on parotid gland tumours and MR radiomics. First, gathering more images would ensure the robustness of the proposed method, along with a validation on images coming from external centres. Besides, radiomics normalisation techniques applied to machine learning should be further studied. Here, we tried to employ an MRI-adjusted ComBat [24], but it yielded poor standardisation results. This method should actually be adapted to suit machine learning tasks, *i.e.*, modified so that it could be calibrated on a training set before being applied to test data. Moreover, as previously observed in [12], one should consider 3D volumes segmentation as direction of tumour growth is an indication of its histopathological subtype. Correlation is high between radiomics characteristics extracted from those volumes, but redundancy can be diminished thanks to intra-class correlation (ICC) and concordance index computations as proposed in [44]. Regarding the experiment performed with radiologists, new investigations could take advantage of a larger number of physicians involved, as well as the implementation of a method to evaluate the level of confidence into their own diagnosis. This kind of assessment could easily be interpreted as class probabilities in order to assess doubts and compute new metrics. In this work, radiologists had no input on the probability of the algorithm's predictions. Uncertainty quantification would surely help professionals to assess the quality of the AI-based diagnosis. Last but not least, as stated in Section 2.1, all the features selected by the algorithm come from the T2w sequence, which is rather unexpected. Some more experiments should be conducted to further explain this behaviour.

Author contributions

Conceptualisation, S.A., A.Q. and E.C.; methodology, A.Q., V.E. and E.C.; data collection, S.A., F.B., G.G., D.H.; data curation, S.A. and A.Q.; software, A.Q.; supervision, E.C., N.L. and C.B.; formal analysis, A.Q., V.E. and E.C.; writing – original draft, A.Q.; writing – review & editing, A.Q., V.E. and E.C.

Acknowledgements

The authors would like to thank the patients who gave their agreement to the use of their MR images in the scope of this study. E.C. and A.Q. received funding from the European Research Council Starting Grant MAJORIS ERC-2019-STG850925. The work of V.E. is supported by ARL/ARO under grant W911NF-22-1-0235.

References

- [1] Guzzo M, Locati LD, Prott FJ, Gatta G, McGurk M, Licitra L: Major and minor salivary gland tumors. *Critical Reviews in Oncology/Hematology*, 74(2):134–148, 2010
- [2] Spiro RH: Salivary neoplasms: Overview of a 35-year experience with 2,807 patients. *Head & Neck Surgery*, 8(3):177–184, 1986
- [3] Reerds STH, Van Engen–Van Grunsven ACH, Hoogen FJA, Takes RP, Marres HAM, Honings J: Accuracy of parotid gland FNA cytology and reliability of the Milan System for Reporting Salivary Gland Cytopathology in clinical practice. *Cancer Cytopathology*, 129(9):719–728, 2021
- [4] Correia-Sá I, Correia-Sá M, Costa-Ferreira P, Silva Á, Marques M: Fine-needle aspiration cytology (FNAC): Is it useful in preoperative diagnosis of parotid gland lesions? *Acta Chirurgica Belgica*, 117(2):110–114, 2017
- [5] Tartaglione T, Botto A, Sciandra M, Gaudino S, Danieli L, Parrilla C, Paludetti G, Colosimo C: Differential diagnosis of parotid gland tumours: Which magnetic resonance findings should be taken in account? *Acta Otorhinolaryngologica Italica*, 35(5):314–320, 2015
- [6] El-Dahshan ESA, Mohsen HM, Revett K, Salem ABM: Computer-aided diagnosis of human brain tumor through MRI: A survey and a new algorithm. *Expert Systems with Applications*, 41(11):5526–5545, 2014
- [7] Wiens J, Shenoy ES: Machine Learning for Healthcare: On the Verge of a Major Shift in Healthcare Epidemiology. *Clinical Infectious Diseases*, 66(1):149–153, 2018
- [8] Zhang G, Zheng C, He J, Yi S: PCT: Pyramid convolutional transformer for parotid gland tumor segmentation in ultrasound images. *Biomedical Signal Processing and Control*, 81:104498, 2023
- [9] Sunnetci KM, Kaba E, Celiker FB, Alkan A: Deep network-based comprehensive parotid gland tumor detection. *Academic Radiology*, 31(1):157–167, 2024
- [10] Lambin P, Rios-Velazquez E, Leijenaar R, Carvalho S, van Stiphout RG, Granton P, Zegers CM, Gillies R, Boellard R, Dekker A, Aerts HJ: Radiomics: Extracting more information from medical images using advanced feature analysis. *European Journal of Cancer*, 48(4):441–446, 2012
- [11] Kumar V, Gu Y, Basu S, Berglund A, Eschrich SA, Schabath MB, Forster K, Aerts HJ, Dekker A, Fenstermacher D, Goldgof DB, Hall LO, Lambin P, Balagurunathan Y, Gatenby RA, Gillies RJ: Radiomics: The process and the challenges. *Magnetic Resonance Imaging*, 30(9):1234–1248, 2012
- [12] Ortiz-Ramón R, Larroza-Santacruz A, Ruiz-España S, Arana Fernandez De Moya E, Moratal D: Classifying brain metastases by their primary site of origin using a radiomics approach based on texture analysis: a feasibility study, in *European Radiology*, volume 28, pages 4514–4523, Springer-Verlag, 2018
- [13] Zhang J, Jin J, Ai Y, Zhu K, Xiao C, Xie C, Jin X: Differentiating the pathological subtypes of primary lung cancer for patients with brain metastases based on radiomics features from brain CT images. *European Radiology*, 31(2):1022–1028, 2021
- [14] Organisation mondiale de la santé, Centre international de recherche sur le cancer, editors: *WHO Classification of Head and Neck Tumours*, number 9 in World Health Organization Classification of Tumours, International agency for research on cancer, Lyon, 4th ed edition, 2017
- [15] Zheng Ym, Li J, Liu S, Cui Jf, Zhan Jf, Pang J, Zhou Rz, Li Xl, Dong C: MRI-Based radiomics nomogram for differentiation of benign and malignant lesions of the parotid gland. *European Radiology*, 2020
- [16] Gabelloni M, Faggioni L, Attanasio S, Vani V, Goddi A, Colantonio S, Germanese D, Caudai C, Bruschini L, Scarano M, Seccia V, Neri E: Can Magnetic Resonance Radiomics Analysis Discriminate Parotid Gland Tumors? A Pilot Study. *Diagnostics*, 10(11):900, 2020
- [17] Khodabakhshi Z, Motisi L, Bink A, Broglie MA, Rupp NJ, Fleischmann M, von der Grün J, Guckenberger M, Tanadini-Lang S, Balermipas P: MRI-based radiomics for predicting histology in malignant salivary gland tumors: methodology and “proof of principle”. *Scientific Reports*, 14(1):9945, 2024, URL <https://www.nature.com/articles/s41598-024-60200-9>
- [18] Mao K, Wong LM, Zhang R, So TY, Shan Z, Hung KF, Ai QYH: Radiomics analysis in characterization of salivary gland tumors on MRI: A systematic review. *Cancers*, 15(20):4918, 2023, URL <https://www.mdpi.com/2072-6694/15/20/4918>

- [19] Zhang R, Ai QYH, Wong LM, Green C, Qamar S, So TY, Vlantis AC, King AD: Radiomics for discriminating benign and malignant salivary gland tumors; which radiomic feature categories and MRI sequences should be used? *Cancers*, 14(23):5804, 2022, URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC9740105/>
- [20] Xia X, Feng B, Wang J, Hua Q, Yang Y, Sheng L, Mou Y, Hu W: Deep Learning for Differentiating Benign From Malignant Parotid Lesions on MR Images. *Frontiers in Oncology*, 11:632104, 2021
- [21] Shinohara RT, Sweeney EM, Goldsmith J, Shiee N, Mateen FJ, Calabresi PA, Jarso S, Pham DL, Reich DS, Crainiceanu CM: Statistical normalization techniques for magnetic resonance imaging. *NeuroImage: Clinical*, 6:9–19, 2014
- [22] Nyúl LG, Udupa JK: On standardizing the MR image intensity scale. *Magnetic Resonance in Medicine*, 42(6):1072–1081, 1999
- [23] Zwanenburg A, Vallières M, Abdalah MA, Aerts HJWL, Andrearczyk V, Apte A, Ashrafinia S, Bakas S, Beukinga RJ, Boellaard R, Bogowicz M, Boldrini L, Buvat I, Cook GJR, Davatzikos C, Depeursinge A, Desseroit MC, Dinapoli N, Dinh CV, Echegaray S, El Naqa I, Fedorov AY, Gatta R, Gillies RJ, Goh V, Götz M, Guckenberger M, Ha SM, Hatt M, Isensee F, Lambin P, Leger S, Leijenaar RT, Lenkowicz J, Lippert F, Losnegård A, Maier-Hein KH, Morin O, Müller H, Napel S, Nioche C, Orlhac F, Pati S, Pfaehler EA, Rahmim A, Rao AU, Scherer J, Siddique MM, Sijtsema NM, Socarras Fernandez J, Spezi E, Steenbakkers RJ, Tanadini-Lang S, Thorwarth D, Troost EG, Upadhyaya T, Valentini V, van Dijk LV, van Griethuysen J, van Velden FH, Whybra P, Richter C, Löck S: The Image Biomarker Standardization Initiative: Standardized Quantitative Radiomics for High-Throughput Image-based Phenotyping. *Radiology*, 295(2):328–338, 2020
- [24] Orlhac F, Lecler A, Savatovski J, Goya-Outi J, Nioche C, Charbonneau F, Ayache N, Frouin F, Duron L, Buvat I: How can we combat multicenter variability in MR radiomics? Validation of a correction procedure. *European Radiology*, 2020
- [25] Beer JC, Tustison NJ, Cook PA, Davatzikos C, Sheline YI, Shinohara RT, Linn KA: Longitudinal ComBat: A method for harmonizing longitudinal multi-scanner imaging data. *NeuroImage*, 220:117129, 2020
- [26] Caruana R, Karampatziakis N, Yessenalina A: An empirical evaluation of supervised learning in high dimensions, in *Proceedings of the 25th International Conference on Machine Learning*, ICML '08, pages 96–103, Association for Computing Machinery, New York, NY, USA, 2008
- [27] Fernández-Delgado M, Cernadas E, Barro S, Amorim D: Do we need hundreds of classifiers to solve real world classification problems? *The Journal of Machine Learning Research*, 15(1):3133–3181, 2014
- [28] Kniep HC, Madesta F, Schneider T, Hanning U, Schönfeld MH, Schön G, Fiehler J, Gauer T, Werner R, Gellissen S: Radiomics of Brain MRI: Utility in Prediction of Metastatic Tumor Type. *Radiology*, 290(2):479–487, 2018
- [29] Wang H, Zhang J, Bao S, Liu J, Hou F, Huang Y, Chen H, Duan S, Hao D, Liu J: Preoperative MRI-Based Radiomic Machine-Learning Nomogram May Accurately Distinguish Between Benign and Malignant Soft-Tissue Lesions: A Two-Center Study. *Journal of Magnetic Resonance Imaging*, 52(3):873–882, 2020
- [30] Breiman L: Random Forests. *Machine Learning*, 45(1):5–32, 2001
- [31] James G, Witten D, Hastie T, Tibshirani R, Taylor J: *An Introduction to Statistical Learning: with Applications in Python*, Springer Texts in Statistics, Springer International Publishing, 2023
- [32] Hastie T, Tibshirani R, Friedman J: *The Elements of Statistical Learning: Data Mining, Inference, and Prediction, Second Edition*, Springer Series in Statistics, Springer-Verlag, New York, second edition, 2009
- [33] Lambin P, Leijenaar RTH, Deist TM, Peerlings J, de Jong EEC, van Timmeren J, Sanduleanu S, Larue RTHM, Even AJG, Jochems A, van Wijk Y, Woodruff H, van Soest J, Lustberg T, Roelofs E, van Elmpt W, Dekker A, Mottaghy FM, Wildberger JE, Walsh S: Radiomics: The bridge between medical imaging and personalized medicine. *Nature Reviews Clinical Oncology*, 14(12):749–762, 2017
- [34] Kuhn M, Johnson K: *An Introduction to Feature Selection*, pages 487–519, Springer New York, New York, NY, 2013
- [35] Kuhn M, Johnson K: *Feature Engineering and Selection: A Practical Approach for Predictive Models*, CRC Press, 2019
- [36] Matsuo H, Nishio M, Kanda T, Kojita Y, Kono AK, Hori M, Teshima M, Otsuki N, Nibu Ki, Murakami T: Diagnostic accuracy of deep-learning with anomaly detection for a small amount of imbalanced data: Discriminating malignant parotid tumors in MRI. *Scientific Reports*, 10(1):19388, 2020
- [37] Yabuuchi H, Kamitani T, Sagiya K, Yamasaki Y, Hida T, Matsuura Y, Hino T, Murayama Y, Yasumatsu R, Yamamoto H: Characterization of parotid gland tumors: Added value of permeability MR imaging to DWI and DCE-MRI. *European Radiology*, 30(12):6402–6412, 2020
- [38] Chang YJ, Huang TY, Liu YJ, Chung HW, Juan CJ: Classification of parotid gland tumors by using multimodal MRI and deep learning. *NMR in Biomedicine*, 34(1), 2021

- [39] Wang CW, Chu YH, Chiu DY, Shin N, Hsu HH, Lee JC, Juan CJ: JOURNAL CLUB: The Warthin Tumor Score: A Simple and Reliable Method to Distinguish Warthin Tumors From Pleomorphic Adenomas and Carcinomas. *American Journal of Roentgenology*, 210(6):1330–1337, 2018
- [40] Zwanenburg A, Leger S, Agolli L, Pilz K, Troost EGC, Richter C, Löck S: Assessing robustness of radiomic features by image perturbation. *Scientific Reports*, 9(1):614, 2019
- [41] Ikeda M, Motoori K, Hanazawa T, Nagai Y, Yamamoto S, Ueda T, Funatsu H, Ito H: Warthin tumor of the parotid gland: Diagnostic value of MR imaging with histopathologic correlation. *AJNR American journal of neuroradiology*, 25(7):1256–1262, 2004
- [42] Okahara M, Kiyosue H, Hori Y, Matsumoto A, Mori H, Yokoyama S: Parotid tumors: MR imaging with pathological correlation. *European Radiology*, 13(S06):L25–L33, 2003
- [43] European Society of Radiology (ESR): Impact of artificial intelligence on radiology: A EuroAIM survey among members of the European Society of Radiology. *Insights into Imaging*, 10(1):105, 2019
- [44] Cattell R, Chen S, Huang C: Robustness of radiomic features in magnetic resonance imaging: Review and a phantom study. *Visual Computing for Industry, Biomedicine, and Art*, 2(1):19, 2019

Extracted radiomic features

IBSI category	Features	Number of features
Shape	Volume, Surface area, Surface area to volume ratio, Sphericity, Compactness 1, Compactness 2, Spherical disproportion, Maximum 3D diameter, Maximum 2D diameter slice, Maximum 2D diameter column, Maximum 2D diameter row, Major axis, Minor axis, Least axis, Elongation, Flatness	16
First order	Energy, Total energy, Entropy, Minimum, 10th percentile, 90th percentile, Maximum, Mean, Median, Interquartile range, Range, Mean absolute deviation, Robust mean absolute deviation, Root mean squared, Standard deviation, Skewness, Kurtosis, Variance, Uniformity	19
GLCM	Autocorrelation, Joint average, Cluster prominence, Cluster shade, Cluster tendency, Contrast, Correlation, Difference Average, Difference entropy, Difference variance, Joint energy, Joint entropy, Informal measure of correlation 1, Informal measure of correlation 2, Inverse difference moment, Inverse difference moment normalized, Inverse difference, Inverse difference normalized, Inverse variance, Maximum probability, Sum average, Sum entropy, Sum of squares	23
GLRLM	Short run emphasis, Long run emphasis, Gray level non uniformity, Gray level non uniformity normalized, Run length non uniformity, Run length non uniformity normalized, Run percentage, Gray level variance, Run variance, Run entropy, Low gray level run emphasis, High gray level run emphasis, Short run low gray level emphasis, Short run high gray level emphasis, Long run low gray level emphasis, Long run high gray level emphasis	16
GLSZM	Small area emphasis, Large area emphasis, Gray level non uniformity, Gray level non uniformity normalized, Size zone non uniformity, Size zone non uniformity normalized, Zone percentage, Gray level variance, Zone variance, Zone entropy, Low gray level zone emphasis, Small area low gray level emphasis, Small area high gray level emphasis, Large area low gray level emphasis, Large area high gray level emphasis	15
NGTDM	Coarseness, Contrast, Busyness, Complexity, Strength	5
GLDM	Small dependence emphasis, Large dependence emphasis, Gray level non uniformity, Dependence non uniformity, Dependence non uniformity normalized, Gray level variance, Dependence variance, Dependence entropy, Low gray level emphasis, High gray level emphasis, Small dependence low gray level emphasis, Small dependence high gray level emphasis, Large dependence low gray level emphasis, Large dependence high gray level emphasis	14

Metrics definitions

Let TP, FP, TN and FN denote the true/false positives and true/false negatives.

- **Accuracy:** Defined as the ratio of samples that are correctly identified, no matter the class.

$$\frac{TP + TN}{TP + FP + TN + FN}$$

- **Specificity:** Also called true negative rate. Defined as the ratio of negatives that are correctly identified.

$$\frac{TN}{TN + FP}$$

When not used in a binary or OvR context, specificity is computed with overall total TN and FP (*i.e.*, across all classes).

- **Sensitivity:** Also called true positive rate. Defined as the ratio of positives that are correctly identified.

$$\frac{TP}{TP + FN}$$

When not used in a binary or OvR context, sensitivity is computed with overall total TP and FN (*i.e.*, across all classes).

- **Balanced accuracy:** Defined as the mean of sensitivity and specificity.

$$\frac{1}{2} \left(\frac{TP}{TP + FN} + \frac{TN}{TN + FP} \right)$$

When not used in a binary or OvR context, balanced accuracy is the average of sensitivity for each class.

- **Area under ROC curve (AUC):** The receiver operating characteristic (ROC) curve plots the TP rate against the FP rate. AUC corresponds to the area between this curve and the x-axis. An AUC greater than 0.5 means that the algorithm predicts better than random. An AUC of 1 maximizes both specificity and sensitivity.