

Influence of uncertainty estimation techniques on false-positive reduction in liver lesion detection

Ishaan Bhat

Image Sciences Institute, University Medical Center Utrecht, Heidelberglaan 100, 3584 CX Utrecht, The Netherlands

i.r.bhat@umcutrecht.nl

Josien P.W. Pluim

Department of Biomedical Engineering, Eindhoven University of Technology, Groene Loper 3, 5612 AE Eindhoven, The Netherlands

j.pluim@tue.nl

Max A. Viergever

Image Sciences Institute, University Medical Center Utrecht, Heidelberglaan 100, 3584 CX Utrecht, The Netherlands

m.a.viergever@umcutrecht.nl

Hugo J. Kuijf

Image Sciences Institute, University Medical Center Utrecht, Heidelberglaan 100, 3584 CX Utrecht, The Netherlands

h.kuijf@umcutrecht.nl

Abstract

Deep learning techniques show success in detecting objects in medical images, but still suffer from false-positive predictions that may hinder accurate diagnosis. The estimated uncertainty of the neural network output has been used to flag incorrect predictions. We study the role played by features computed from neural network uncertainty estimates and shape-based features computed from binary predictions in reducing false positives in liver lesion detection by developing a classification-based post-processing step for different uncertainty estimation methods. We demonstrate an improvement in the lesion detection performance of the neural network (with respect to F1-score) for all uncertainty estimation methods on two datasets, comprising abdominal MR and CT images, respectively. We show that features computed from neural network uncertainty estimates tend not to contribute much toward reducing false positives. Our results show that factors like class imbalance (true over false positive ratio) and shape-based features extracted from uncertainty maps play an important role in distinguishing false positive from true positive predictions. Our code can be found at <https://github.com/ishaanb92/FPCPipeline>

Keywords: Deep learning, Uncertainty estimation, False-positive reduction

1. Introduction

Deep learning systems have become the go-to approach for various medical image analysis tasks (Zhou et al., 2021; Litjens et al., 2017). However, such systems may make erroneous predictions. These arise due to a variety of reasons, for example, the model overfitting to the training data, presence of noise/artefacts in the image, or mismatch between the training and clinical data distributions (Castro et al., 2020). False-positive prediction is one such type of error and may hinder accurate diagnosis.

Lesion detection in medical imaging is usually formulated as an image segmentation problem, where labels are assigned to each voxel in the image. The task suffers from class imbalance, because lesions occupy a small volume in the image. The imbalance is usually corrected by re-weighting components of the loss function and/or selectively sampling data. While this may lead to more voxels being correctly classified as lesions, the predic-

tion algorithm may become biased toward detecting more lesions than are actually present. False-positive reduction is accomplished via negative sampling i.e. sampling background to reduce algorithmic bias toward finding lesions (Ciga et al., 2021), architectural modifications (Tao et al., 2019), or by using a classifier as a post-processing step (Chlebus et al., 2018; Bhat et al., 2021).

In this paper, we extend initial work performed in Bhat et al. (2021) by developing a system to reduce false-positive predictions using uncertainty estimations made by the neural network. We study the effect uncertainty estimation methods have on the detection of false positive lesions. To do so, we develop a model-agnostic post-processing pipeline that can work with predictions made by any probabilistic classifier.

We summarize related work in Section 2. Section 3 describes the datasets used in this paper. In Section 4 we provide background about the different uncertainty methods we studied, and describe the false-positive classification pipeline. In Section 5 we describe the neural network architecture and training (Section 5.1), and the different evaluation methods used to analyse the performance of uncertainty estimation methods to reduce false positives (Section 5.2). Results are shown in Section 6 and their implications are discussed in Section 7.

2. Related Works

The estimation of prediction uncertainty has been useful in domains such as detecting adversarial (Smith and Gal, 2018) and out-of-distribution (Ovadia et al., 2019) examples, robotics (Loquercio et al., 2020), autonomous driving (Hoel et al., 2020), and reinforcement learning (Depeweg et al., 2018). Uncertainty estimation has also found numerous promising applications in the field of medical image analysis (Ching et al., 2018).

Uncertainty estimates have been used as a proxy to determining the quality of the predictions, with poor quality predictions being referred to an expert. Leibig et al. (2017) showed that the detection of diabetic retinopathy improved when predictions with estimated uncertainty above a certain threshold were referred to an expert. Seeböck et al. (2020) used the uncertainty map to detect anomalous regions in retinal OCT images. For the task of cardiac MR segmentation, Sander et al. (2019) referred voxels with high uncertainty to an expert to produce segmentations with more accurate boundaries. Sedai et al. (2018) showed the benefits of using the uncertainty map as a visual aid to clinicians for the task of retinal layer segmentation in OCT images. Camarasa et al. (2021) show metrics derived from the uncertainty maps can be used to evaluate the quality of multi-class segmentations of the carotid artery. Karimi et al. (2019) used uncertainty estimation to flag poor segmentations and used a statistical shape model to improve the result for the task of prostate segmentation from ultrasound images. In addition to determining quality of predictions, uncertainty estimates have also been shown to improve performance of downstream tasks within a cascade (Mehta et al., 2019).

When it comes to segmenting structures of interest within images, such as lesions, there is a need to quantify uncertainty at the structure level, by aggregating per-voxel uncertainty estimates, because it is at the structure level that diagnosis is performed. Nair et al. (2020) showed that the log-sum of voxel-wise uncertainties computed over the detected structures can be used to filter small false positives. Similarly, Graham et al. (2019) used the mean

uncertainty over segmented glandular structures to disregard poor quality predictions made on colon histopathological images. Mehrtaash et al. (2020) showed that for prostate, brain lesion and cardiac segmentation, the mean uncertainty over the segmented structure had a negative correlation with the Dice overlap and therefore the mean uncertainty could be used as a proxy for segmentation quality. Similarly, Ng et al. (2020) and Roy et al. (2018) developed metrics to check variations in the prediction at the structure level and showed a negative correlation between this metric and the Dice coefficient for cardiac and brain lesion segmentation, respectively. Eaton-Rosen et al. (2018) used uncertainty estimation to compute well-calibrated error bars for the estimated brain lesion volumes for clinical use.

There is also a class of methods that use uncertainty estimates as an input to a second stage that is used to predict segmentation quality or detect segmentation failures. Jungo et al. (2020) used the task of segmenting different types of lesions from brain MR images to show that aggregating features computed from the per-voxel uncertainty maps can predict segmentation quality. These features correlated well with the Dice score of the prediction and therefore could be used to flag a poor segmentation. In a similar vein, DeVries and Taylor (2018) used a second neural network as a regressor to predict the Jaccard index using the per-voxel uncertainty map and binary segmentation as inputs for their skin lesion segmentation task. Sander et al. (2020) used a second neural network to detect local segmentation errors in a cardiac segmentation task using the uncertainty map, the prediction, and image patches as inputs.

Most approaches use voxel-wise uncertainty estimates to detect segmentation failures, or use them as a proxy for segmentation quality. Nair et al. (2020) showed that aggregating voxel-wise uncertainty estimates over the predicted lesion volume using the log-sum operation can be used to filter small false positives. Jungo et al. (2020) compared different ways to perform spatial aggregation at the subject level and concluded using radiomics features to construct subject-level feature vectors performed best at detecting failures while segmenting tumours in brain MR images. Since this is a more general form of spatial aggregation, we construct radiomics feature vectors corresponding to predicted lesion volumes from the uncertainty map estimated by the neural network to investigate the role of uncertainty estimates in reducing false-positive detections.

3. Data

3.1 UMC MR dataset

We used abdominal dynamic contrast enhanced (DCE) and diffusion weighted (DW) MRI of 71 patients with liver metastases from the University Medical Center Utrecht, the Netherlands. This data was acquired between February 2015 and February 2018. The DCE-MR series was acquired in six breath holds, resulting in a total of 16 3-D contrast phases per patient. Voxel spacing for these images is $1.5 \text{ mm} \times 1.5 \text{ mm} \times 2 \text{ mm}$. The liver and the metastases within the liver were manually segmented on the DCE-MRI by a radiologist in training and verified by a radiologist with more than 10 years of experience. The dataset mainly included colorectal and neuroendocrine metastases, with few other types in addition (i.e., other gastrointestinal metastases and breast metastases). Motion correction between the different contrast phases for each DCE MR image was performed using the techniques described in Jansen et al. (2017).

The DW-MR images were acquired with three b-values: 10, 150, and 1000 s/mm², using a protocol with the following parameters: TE: 70 ms; TR: 1.660 ms; flip angle: 90 degrees. For each patient, the DW MR image was non-linearly registered to the DCE MR image using the elastix (Klein et al., 2010) toolbox.

We applied the manually created liver masks to the abdominal DCE and DW MR images. All images were resized to in-place dimensions of 256×256 , followed by z-score normalization of the intensities. The 16 contrast phases of the DCE MR images were divided into groups and averaged to generate a 6-channel image (Jansen et al., 2019). The contrast phases were grouped as follows, the pre-contrast image, the early arterial phase, the late arterial phase, the hepatic/portal-venous phase, the late portal-venous/equilibrium phase, and the late equilibrium. The DW images were concatenated with the DCE images along the channel axis to create a 9-channel input image for the neural network.

The data was split into 50 training patients, 5 validation patients and 16 test patients.

3.2 Liver lesion Segmentation Challenge (LiTS)

The LiTS dataset (Bilic et al., 2019)¹ contains 131 abdominal CT images with reference liver and lesion segmentations from seven hospitals and research institutes. Out of these 131 images, we excluded 13 images that had no lesions. All images in the dataset had in-plane dimensions of 512 x 512 with variable voxel spacings. The median voxel spacing for the dataset was 0.76 x 0.76 x 1.0 mm.

Pre-processing on the images was performed by first clipping the intensities to the range $[-100, 200]$ and then re-scaling the clipped intensities to the range $[0, 1]$. Non-liver regions in the image were masked using the reference liver segmentation.

We divided our reduced dataset of 118 CT images into non-overlapping sets with 94 training, 6 validation and 18 test patients.

4. Materials and Methods

We have developed a classification pipeline to study the effect of uncertainty estimation techniques on false-positive reduction. Section 4.1 provides a general framework to describe different uncertainty estimation methods. Section 4.2 explains our graph-based lesion counting method to account for many-to-one and one-to-many correspondences between the predicted and ground truth segmentations while evaluating lesion detection performance. Finally, Section 4.3 describes the classification pipeline used to reduce false-positive lesion segmentations.

4.1 Theory

Standard neural network training and inference provide point estimates for their prediction. There has been early work on uncertainty estimation utilizing a Bayesian framework to train neural networks that can produce a distribution instead of a point estimate (MacKay, 1992; Neal, 1996). However, neural network training and inference with these approaches was computationally expensive. Recent approaches build upon these earlier approaches to make them computationally efficient using variational inference techniques (Blundell et al.,

1. <https://competitions.codalab.org/competitions/17094>

2015; Gal and Ghahramani, 2016; Kingma et al., 2015), the most widely adopted amongst these being *MC-Dropout* (Gal and Ghahramani, 2016). In addition to these methods, techniques such as using model ensembles (Lakshminarayanan et al., 2017) and test-time augmentation (Wang et al., 2019) are widely used and have been described later in this subsection.

We first present a general method to compute the output distribution of a neural network, and then show how each of the methods mentioned before can be derived from this general formulation. In a Bayesian framework, the output distribution of a neural network, $p(y^*|x^*, X, Y)$, is obtained by marginalizing over the posterior distribution of the model parameters and image transformations:

$$p(y^*|x^*, X, Y) = \int_{\omega \in \Omega} \int_{\tau \in \Lambda} p(y^*|x^*, \omega, \tau) p(\omega|X, Y) p(\tau) d\tau d\omega \quad (1)$$

Here x^* , y^* are the input and output values respectively, X and Y are the data and labels used for training respectively, and ω , τ are the set of neural network parameters and input transformations respectively. The integral in Equation 1 is approximated using a Monte-Carlo simulation:

$$p_{\text{approx}}(y^*|x^*, X, Y) = \frac{1}{NM} \sum_{i=1}^N \sum_{j=1}^M p(y^*|x^*, \omega_i, \tau_j) \quad (2)$$

4.1.1 BASELINE

As a baseline, we included a convolutional neural network with a softmax output producing per-voxel point probability estimates, $p(y^*|x^*, \omega^*)$, where y^* , x^* are the output and input respectively, and ω^* is the set of neural network parameters obtained via optimization. Therefore, the distribution over the model parameters, $p(\omega|X, Y)$ becomes a Dirac delta function at ω^* , and the distribution over input transformations, $p(\tau)$ a Dirac delta function at the identity transform.

4.1.2 MC-DROPOUT

Dropout (Srivastava et al., 2014) is a regularization technique where different random subsets of the neural network weights are set to zero during training, based on sampling masks from a Bernoulli distribution with a fixed probability or *dropout rate*. Gal and Ghahramani (2016) show that weights retained during dropout can be used as samples of the posterior distribution over the weights, $p(\omega|X, Y)$.

Therefore, for MC-Dropout, Equation 2 reduces to:

$$p_{\text{mcd}}(y^*|x^*) = \frac{1}{N_{\text{mcd}}} \sum_{i=1}^{N_{\text{mcd}}} p(y^*|x^*, \omega_i) \quad (3)$$

Here ω_i are the model parameters retained during the i^{th} pass during inference.

4.1.3 MODEL ENSEMBLES

Model ensembles (Lakshminarayanan et al., 2017) have been shown to have superior performance and yield high quality uncertainty estimates as compared with a single neural

network. Models within an ensemble are trained independently of each other, and the final prediction is obtained by averaging the outputs of the different models. From a Bayesian viewpoint, this can be seen as another way to draw samples from Ω . Therefore, for model ensembles, Equation 2 becomes:

$$p_{\text{ensemble}}(y^*|x^*, X, Y) = \frac{1}{N_{\text{ensemble}}} \sum_{i=1}^{N_{\text{ensemble}}} p(y^*|x^*, \omega_i) \quad (4)$$

Here ω_i is the set of weights belonging to the i^{th} model in the ensemble.

4.1.4 TEST-TIME AUGMENTATION

Test-time augmentation (TTA) (Wang et al., 2019) computes the output probability distribution by marginalizing over a distribution of image transformations, $p(\tau)$, used to transform the input image.

$$p_{\text{tta}}(y^*|x^*, \omega^*) = \frac{1}{M_{\text{tta}}} \sum_{j=1}^{M_{\text{tta}}} p(y^*|x^*, \omega^*, \tau_j) \quad (5)$$

Here τ_j is the j^{th} transform sampled from $p(\tau)$ and ω^* is the point estimate of the model parameters obtained via optimization.

4.1.5 COMBINING ANY METHOD WITH TTA

MC-Dropout or model ensembles can be combined with test-time augmentation, as shown in Equation 2 where ω_i could either be the weights retained during i^{th} dropout pass or weights from the i^{th} model in the ensemble

4.1.6 UNCERTAINTY ESTIMATION AND DECOUPLING

The predictive uncertainty for classification problems, i.e. with a discrete set of C labels as possible output values, is quantified by the entropy of the output distribution. The entropy for such a distribution p is given by:

$$u_{\text{predictive}} = H[p] = - \sum_{c=1}^C p(y^* = c|x^*) \log(p(y^* = c|x^*)) \quad (6)$$

Entropy is low when all the probability mass is concentrated on a single class, and high when it is distributed more evenly over classes.

The predictive uncertainty can be decomposed into uncertainty originating from the model (epistemic) and uncertainty originating from the data (aleatoric) (Kiureghian and Ditlevsen, 2009). This decomposition is possible if there is variation in the model parameters used to generate the prediction, as is the case with MC-Dropout or model ensembles (Smith and Gal, 2018).

Aleatoric uncertainty reflects the true noise in the data and cannot be reduced by using more data to train the model. Inherent noise in the data may be reflected in a prediction

with high entropy. Thus, the aleatoric uncertainty is the mean entropy computed over predictions with fixed parameters:

$$u_{\text{aleatoric}} = \frac{1}{T} \sum_{i=1}^T H[p(y^*|x^*, \omega_i)] \quad (7)$$

The epistemic uncertainty in the prediction is the mutual information between the model parameters and the data (Smith and Gal, 2018). This is equivalent to the epistemic uncertainty being the difference between the predictive uncertainty (Equation 6) and the aleatoric uncertainty (Equation 7):

$$u_{\text{epistemic}} = u_{\text{predictive}} - u_{\text{aleatoric}} \quad (8)$$

For models that use a single fixed set of parameters (eg: TTA), Equation 7 shows that the aleatoric uncertainty is equal to the predictive uncertainty, i.e. the predictive uncertainty is estimated from the data only. For predictions from either MC-Dropout or model ensembles, all three uncertainty measures may be computed.

Image segmentation can be viewed as a per-voxel classification problem, and per-voxel uncertainty estimates form an uncertainty map. As an example, we show the original image (with the binary prediction overlaid) and uncertainty maps for a single CT image slice from the LiTS dataset in Figure 1. The epistemic uncertainty maps for MC-Dropout and MC-Dropout+TTA are almost completely empty. This signifies that there is low variation in the MC-Dropout and MC-Dropout+TTA outputs, produced by different passes through the model for a given input (See Equations 6, 7, 8). On the other hand, in Figure 1, the rows corresponding to ensemble and ensemble+TTA, the epistemic uncertainty maps show signal outlining the lesion boundaries, comparable to the predictive and aleatoric uncertainty maps. Following a similar reasoning, this indicates that there is high variation in the outputs of the different models in the ensemble for a given input. This occurs because in an ensemble of independently trained models, the outputs from each model are less correlated as compared to MC-Dropout outputs, which are generated by sampling subsets of weights optimized jointly (Fort et al., 2020).

4.2 Lesion detection

Lesions were detected using a neural network trained to perform image segmentation (See Section 5.1 for details). Inference was performed on 2-D slices of the images, which were combined to form the segmented volume. For MC-Dropout and test-time augmentation, we used values of $N_{mcd} = 10$ and $M_{tta} = 10$. Larger values of N_{mcd} and M_{tta} did not result in any change in the precision and recall metrics on the test set (shown in Table 3). For test-time augmentation, random rotations between $[-45, 45]$ degrees were applied to compute predictions. For the ensemble configurations, we used five independently trained models ($N_{ensemble} = 5$) to compute the prediction on test images. We applied a threshold of 0.5 to convert the per-voxel probability output maps to binary lesion masks.

Post-processing was performed on the binary lesion masks using a morphological closing with a $3 \times 3 \times 3$ cube-shaped element followed by morphological opening with a $3 \times 3 \times 3$ cross-shaped element (Jansen et al., 2019). Within the post-processed volume, we performed connected component analysis to identify lesion volumes.

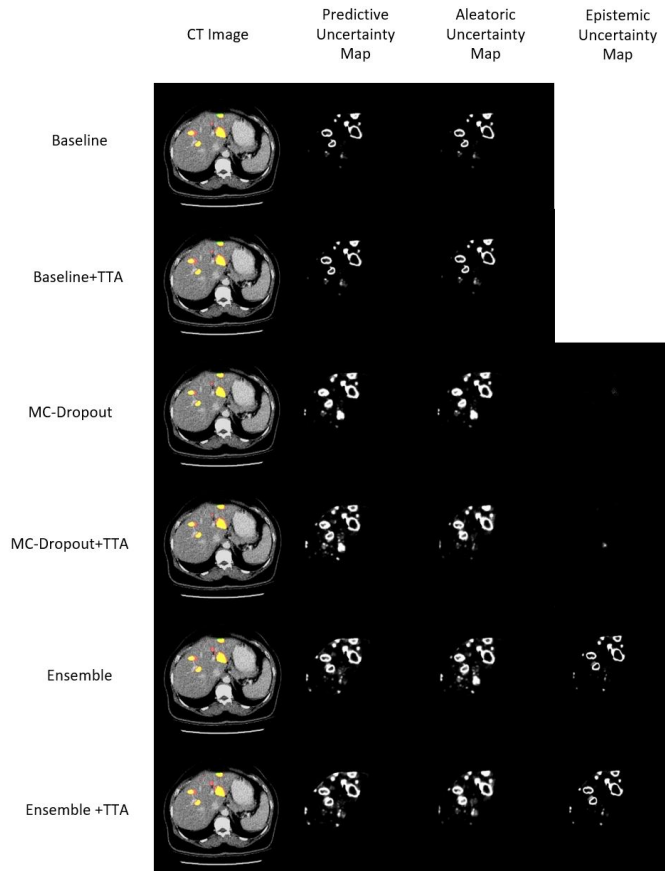


Figure 1: Uncertainty maps for different estimation techniques. In the left-most column, we show the CT image slice with the reference segmentation (green), predicted segmentation (red) and their overlap (yellow). Since model parameters are fixed for Baseline and Baseline+TTA configurations, the epistemic uncertainty is always 0.

To quantify the efficacy of our segmentation network at lesion detection, we used the precision, recall and F1 metrics. However, to compute these metrics, it is necessary to find correspondences between the detected lesion volumes and the lesion volumes present in the reference segmentation. Situations could arise where many-to-one or one-to-many correspondences between lesion volumes are found, and any method used to count the number of detected lesions must take this into account (Oguz et al., 2018).

We tackled the problem of counting detections by constructing a directed bipartite graph with 2 sets of nodes, i.e. the lesions in the predicted segmentation and reference segmentation respectively. We constructed edges between the lesions from the two sets of nodes and set the edge weight between them to be the Dice overlap between the lesion volumes. Edge construction was done in both directions, overlap for a single volume in the predicted segmentation was computed with each object in the reference segmentation and vice versa.

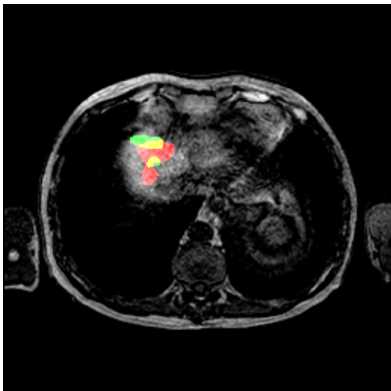
Once the graph edges are constructed, counting the number of true positive, false positive and false negative predictions is straightforward. Any volume in the predicted segmentation is considered a *true positive* detection if it has at least one outgoing edge with a non-zero weight. To account for many-to-one correspondences, we count true positives as the number of nodes in the reference lesion partition with at least one incoming edge to avoid multiple counting. A volume in the predicted segmentation is considered a *false positive* prediction if it has no edges with a non-zero weight. Similarly, a volume in the reference segmentation is considered a *false negative* if it has no incoming edges with a non-zero weight. In Figure 2, we show an example of a correspondence graph with a many-to-one lesion correspondence.

The detection metrics are computed as follows:

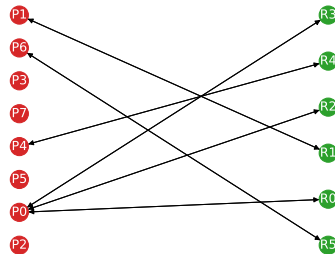
$$\text{precision} = \frac{\text{true positives}}{\text{true positives} + \text{false positives}} \tag{9}$$

$$\text{recall} = \frac{\text{true positives}}{\text{true positives} + \text{false negatives}} \tag{10}$$

$$\text{F1} = \frac{2 \times \text{precision} \times \text{recall}}{\text{precision} + \text{recall}} \tag{11}$$



(a) Prediction (red) and reference (green) segmentations (overlap in yellow), overlaid on a DCE MR image. In this image slice, it can be seen that the same predicted lesion volume overlaps with two lesions from the reference segmentations. In such a case, we count two detected lesions.



(b) Graph capturing lesion correspondences for the DCE MR image in Figure 2a in terms of overlap. The nodes on the left in red are lesions in the predicted segmentation, and those on the right in green are lesions from the reference segmentation. Edges with 0 weight have not been shown for clarity. It can be seen that three lesions (two of which can be seen in Figure 2a) in the reference segmentation correspond to a single lesion (P0) in the prediction.

Figure 2: Correspondences between predicted and reference segmentations

4.3 False-positive classification

The steps involved in filtering false positive predictions from the neural network prediction are shown in Figure 3. The pipeline was designed following guidelines on radiomics workflows detailed in Van Timmeren et al. (2020).

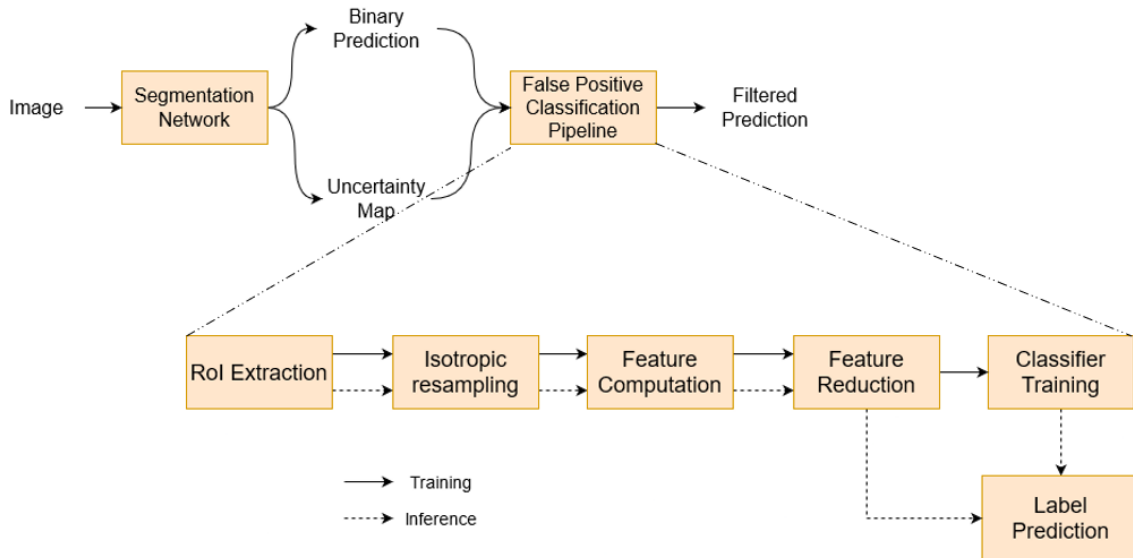


Figure 3: Pipeline to filter false positive lesion predictions from neural network predictions using uncertainty maps.

We used the lesion detections from the neural network training and validation patients as training data for our false-positive classifier and use the held-out test set to report the performance.

4.3.1 REGION OF INTEREST (ROI) EXTRACTION

We use the binarized neural network prediction to extract 3-D patches from the uncertainty map corresponding to the detected lesions. The lesion correspondence graph is used to assign a label (true positive or false positive) to each patch. The 3-D patch from the uncertainty map(s) and the ROI extracted from the binarized prediction are used for computing intensity and shape features, respectively.

4.3.2 ISOTROPIC RESAMPLING

We resample each 3-D uncertainty patch and ROI mask to isotropic spacing using 3^{rd} order B-spline and nearest-neighbor interpolation, respectively. Isotropic resampling was performed to ensure rotational invariance of texture features (Zwanenburg et al., 2020). For the LiTS dataset, patches were resampled along all axes to $0.76\text{mm} \times 0.76\text{mm} \times 0.76\text{mm}$, the median in-plane voxel spacing for the dataset. For the UMC dataset, we keep the in-plane voxel spacing unchanged and resample in the axial direction to achieve an isotropic spacing of $1.543\text{mm} \times 1.543\text{mm} \times 1.543\text{mm}$.

4.3.3 FEATURE COMPUTATION

For each 3-D patch from the uncertainty map, we computed features as described by the Imaging Biomarker Standardization Initiative (IBSI) (Zwanenburg et al., 2020). We used the PyRadiomics (Van Griethuysen et al., 2017) library to compute 99 radiomics features for each patch.

4.3.4 FEATURE REDUCTION

We designed our feature reduction step to produce a set of minimally correlated features. Hierarchical clustering was performed by constructing a distance matrix such that distance between a pair of features was inversely proportional to their correlation. In Figure 4 we show the hierarchical clusters formed for the MC-Dropout (Predictive) configuration of the LiTS dataset. The next step was forming *flat* clusters of features by choosing a threshold at which the dendrogram is *cut*. All leaves below a cut form a single *flat* cluster, and the feature with maximum mutual information with respect to the labels was chosen from each of the flat clusters (highlighted in bold in Figure 4). The same threshold (1.0) was used for both datasets and all configurations to produce an almost uniform number of features. The dendrogram threshold was chosen by visually examining hierarchical clusters for all configurations and datasets. The threshold was chosen to ensure approximately the same number and sizes of clusters. This could be achieved for most configurations, although some were outliers with fewer features. The idea of using a fixed threshold was to ensure the use of the same extent of correlation between features while converting the hierarchical clusters to flat clusters.

In Table 1 we show reduced feature vector sizes (N^*) computed for all the configurations for both datasets.

Table 1: Reduced feature vector sizes across datasets and configurations

Configuration	Uncertainty Type	LiTS	UMC
		N^*	N^*
Baseline	Predictive	9	9
Baseline + TTA	Predictive	9	8
MC-Dropout	Predictive	6	7
	Aleatoric	5	7
	Epistemic	6	6
MC-Dropout + TTA	Predictive	5	7
	Aleatoric	5	7
	Epistemic	6	5
Ensemble	Predictive	8	8
	Aleatoric	9	9
	Epistemic	7	8
Ensemble + TTA	Predictive	8	8
	Aleatoric	9	8
	Epistemic	8	8

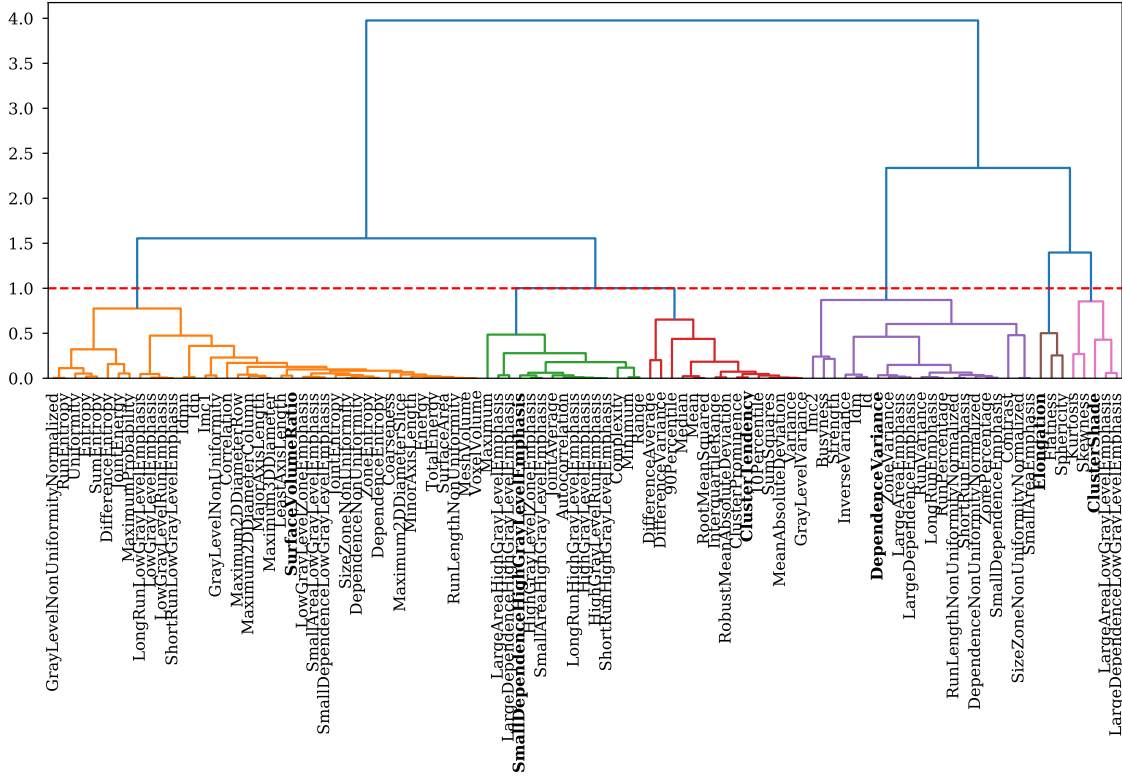


Figure 4: Hierarchical cluster formation based on the feature correlation matrix. The dashed red line shows the threshold at which the dendrogram is *cut*. A single feature is selected (based on maximum mutual information with respect to the label) from each of the flat clusters formed after thresholding (shown in bold).

4.3.5 CLASSIFIER TRAINING

We used features computed from uncertainty map patches, corresponding to lesion detections in the training and validation images for the neural network, to train an Extremely Randomized Trees (ERT) (Geurts et al., 2006) classifier to classify the detected lesion, as a true or false positive. This classifier is an example of a bagging predictor (Breiman, 1996), which performs the classification task using an ensemble of randomized base classifiers. The base classifier for the ERT is a decision tree, and randomization is introduced in each decision tree using two methods. First, random subsets of input features are used to split tree nodes. Second, thresholds for each of the input features are sampled randomly and the best of these is used as the splitting rule. Injecting randomness in the base classifiers decouple their prediction errors and reduce overfitting in the constructed ensemble.

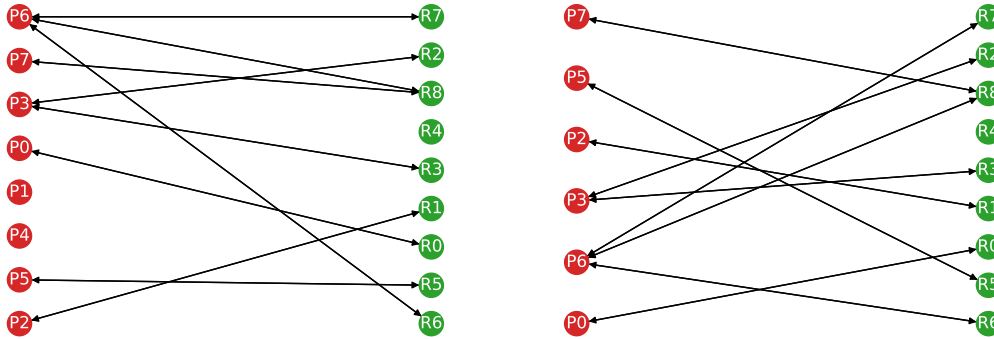
The hyperparameters for the classifier were chosen via five-fold cross-validation while optimizing for the area under the curve (AUC) metric, to achieve a good trade-off between sensitivity and specificity. The hyperparameters are shown in Table 6 in Appendix A. The number of training samples for both datasets and all uncertainty estimation methods is

shown in Table 2. We used the scikit-learn (Pedregosa et al., 2011) library to develop code related to classifier training.

Table 2: Training dataset size and class imbalance ratio for the LiTS and UMC datasets

Configuration	LiTS		UMC	
	Training Samples	TP:FP ratio	Training Samples	TP:FP ratio
Baseline	1356	0.684	289	1.125
Baseline + TTA	1282	0.765	233	1.118
MC-Dropout	1267	0.710	325	1.321
MC-Dropout + TTA	1564	0.494	576	0.714
Ensemble	1332	0.743	166	2.388
Ensemble + TTA	1328	0.722	145	2.536

4.3.6 LABEL PREDICTION



(a) Lesion correspondence graph constructed from the neural network prediction

(b) Lesion correspondence graph after label prediction

Figure 5: Modifying the lesion correspondence graph based on classifier predictions. All the false positives (detected structures P1, P4 in Figure 5a) are filtered out.

We used the trained classifier to predict the label (true positive or false positive) of each detected lesion for each patient in the held-out test set. For each patient, the predicted labels for each lesion in the neural network prediction were then used to modify the lesion correspondence graph (Section 4.2). If a lesion was classified as a false positive, the corresponding node in the graph was deleted. We show this in Figure 5.

5. Experiments

5.1 Neural network architecture and training

We used the 2-D U-Net (Ronneberger et al., 2015) architecture to create identical segmentation networks for both the datasets. Each convolution block consisted of a 3×3 weight

kernel, instance normalization and a leaky ReLU activation function. Downsampling was performed via max-pooling and up-sampling was performed via transposed convolutions. At the output layer, a 1×1 convolution was performed, followed by the application of the softmax function to obtain per-voxel class probabilities.

The networks for the UMC and LiTS datasets were trained for approximately 10K (batch size = 16) and 20K (batch size = 8) batch iterations based on early-stopping while monitoring the decrease in the validation loss. For both datasets, we used the ADAM (Kingma and Ba, 2015) optimizer with a learning rate of 10^{-4} and weight decay of 5×10^{-5} . To tackle class-imbalance, we used only slices containing a lesion to train the neural network. Furthermore, we used a weighted cross-entropy as the loss function, with the per-class weights inversely proportional to the fraction of voxels belonging to that class. For both datasets, data augmentation was performed by randomly rotating the input by an angle sampled from $[-45, 45]$ degrees.

For the MC-Dropout models, dropout layers were added at the outputs of the encoder and decoder at the lowest levels, similar to the Bayesian SegNet (Kendall et al., 2017). The dropout rate with the lowest validation loss was chosen. For the network trained using the UMC dataset, this was 0.3, and for the network trained using the LiTS dataset, this was 0.5. For test-time augmentation, we use rotations with the angles uniformly sampled from the range used during data augmentation.

The code for the neural network training and inference was developed using the PyTorch (Paszke et al., 2019) library.

For each uncertainty estimation method, we trained the false-positive classifier and used lesions from the neural network test set to evaluate the performance. For both datasets and each uncertainty estimation method, we trained five classifiers with different random seeds to report confidence intervals on our metrics.

5.2 Evaluation

We used the precision, recall and F1-score metrics to quantify the effect of our false-positive classification pipeline on lesion detection.

To isolate the role played by features computed from uncertainty estimates in reducing false positives, we trained the classifier using just the radiomics shape-based features. These features are computed from the binary lesion mask and, therefore, are independent of the per-voxel uncertainty estimates. Similar to Chlebus et al. (2018), we trained the classifier using radiomics features computed from image patches corresponding to lesion segmentations.

Furthermore, to study the impact of spatial aggregation of uncertainty on false-positive classification, we compare our false positive classification pipeline to the threshold-based filtering approach proposed by Nair et al. (2020). The classification threshold was determined using a hold-out validation set. The thresholds used were 0.9995 and 0.9999 for the UMC and LiTS datasets, respectively.

Additionally, we performed *cross testing* and *combined testing* to investigate the generalization properties of our classifier. In cross-testing, for each uncertainty estimation method, we used the classifier trained on one dataset to predict the type of detection in the test set of the other dataset. In combined testing, for each uncertainty estimation method, we

trained a classifier by merging training sets from both datasets and used it to predict the type of detection in the test set. Both these experiments used radiomics features computed from the uncertainty maps.

We measured feature importance using the leave-one-covariate-out (LOCO) (Lei et al., 2018) method, where feature importance was measured directly in terms of prediction performance on unseen data.

Table 3: Lesion detection metrics for the LiTS and UMC datasets before and after false-positive reduction

Configuration	Input features	LiTS						UMC					
		Precision		Recall		F1		Precision		Recall		F1	
		Before	After	Before	After	Before	After	Before	After	Before	After	Before	After
Baseline	Predictive Uncertainty	0.318	0.766 ± 0.005	0.700	0.510 ± 0.006	0.438	0.612 ± 0.005	0.245	0.608 ± 0.014	0.717	0.608 ± 0.006	0.365	0.608 ± 0.010
	Image	0.318	0.727 ± 0.004	0.700	0.538 ± 0.012	0.438	0.618 ± 0.009	0.245	0.551 ± 0.008	0.717	0.604 ± 0.000	0.365	0.578 ± 0.004
	Binary mask	0.318	0.691 ± 0.005	0.700	0.510 ± 0.000	0.438	0.587 ± 0.002	0.245	0.616 ± 0.022	0.717	0.603 ± 0.000	0.365	0.609 ± 0.010
Baseline + TTA	Predictive Uncertainty	0.335	0.744 ± 0.020	0.690	0.492 ± 0.003	0.451	0.592 ± 0.006	0.289	0.605 ± 0.015	0.774	0.645 ± 0.006	0.421	0.624 ± 0.008
	Image	0.335	0.634 ± 0.004	0.690	0.492 ± 0.003	0.451	0.554 ± 0.003	0.289	0.577 ± 0.004	0.774	0.611 ± 0.007	0.421	0.593 ± 0.003
	Binary mask	0.335	0.725 ± 0.008	0.690	0.516 ± 0.004	0.451	0.603 ± 0.003	0.289	0.659 ± 0.005	0.774	0.626 ± 0.005	0.421	0.642 ± 0.005
MC-Dropout	Predictive Uncertainty	0.373	0.841 ± 0.014	0.720	0.506 ± 0.004	0.491	0.632 ± 0.006	0.261	0.465 ± 0.013	0.774	0.668 ± 0.008	0.39	0.548 ± 0.006
	Aleatoric Uncertainty	0.373	0.740 ± 0.014	0.720	0.496 ± 0.01	0.491	0.594 ± 0.011	0.261	0.479 ± 0.004	0.774	0.683 ± 0.006	0.39	0.563 ± 0.004
	Epistemic Uncertainty	0.373	0.838 ± 0.001	0.720	0.516 ± 0.004	0.491	0.639 ± 0.003	0.261	0.641 ± 0.003	0.774	0.506 ± 0.006	0.39	0.565 ± 0.005
	Image	0.373	0.695 ± 0.007	0.720	0.470 ± 0.006	0.491	0.561 ± 0.003	0.261	0.494 ± 0.002	0.774	0.679 ± 0.000	0.39	0.572 ± 0.001
	Binary mask	0.373	0.790 ± 0.004	0.720	0.534 ± 0.004	0.491	0.637 ± 0.004	0.261	0.563 ± 0.000	0.774	0.679 ± 0.000	0.39	0.615 ± 0.000
	Predictive Uncertainty	0.3	0.732 ± 0.004	0.78	0.514 ± 0.006	0.433	0.604 ± 0.005	0.204	0.504 ± 0.007	0.83	0.758 ± 0.006	0.327	0.605 ± 0.005
MC-Dropout + TTA	Aleatoric Uncertainty	0.3	0.734 ± 0.005	0.78	0.514 ± 0.004	0.433	0.605 ± 0.004	0.204	0.535 ± 0.002	0.83	0.758 ± 0.006	0.327	0.627 ± 0.003
	Epistemic Uncertainty	0.3	0.743 ± 0.003	0.780	0.6 ± 0.0	0.433	0.664 ± 0.001	0.204	0.6 ± 0.006	0.83	0.668 ± 0.008	0.327	0.632 ± 0.005
	Image	0.3	0.705 ± 0.006	0.780	0.522 ± 0.003	0.433	0.600 ± 0.003	0.204	0.496 ± 0.011	0.83	0.784 ± 0.007	0.327	0.607 ± 0.008
	Binary mask	0.3	0.712 ± 0.005	0.780	0.592 ± 0.009	0.433	0.646 ± 0.006	0.204	0.499 ± 0.009	0.83	0.789 ± 0.005	0.327	0.611 ± 0.008
	Predictive Uncertainty	0.385	0.804 ± 0.011	0.790	0.558 ± 0.009	0.518	0.659 ± 0.008	0.375	0.754 ± 0.008	0.792	0.543 ± 0.006	0.509	0.632 ± 0.005
	Aleatoric Uncertainty	0.385	0.806 ± 0.004	0.790	0.530 ± 0.000	0.518	0.639 ± 0.001	0.375	0.546 ± 0.010	0.792	0.725 ± 0.015	0.509	0.622 ± 0.006
Ensemble	Epistemic Uncertainty	0.385	0.811 ± 0.008	0.790	0.558 ± 0.006	0.518	0.661 ± 0.007	0.375	0.616 ± 0.029	0.792	0.547 ± 0.000	0.509	0.579 ± 0.013
	Image	0.385	0.779 ± 0.008	0.790	0.544 ± 0.004	0.518	0.641 ± 0.002	0.375	0.524 ± 0.013	0.792	0.679 ± 0.000	0.509	0.591 ± 0.008
	Binary mask	0.385	0.799 ± 0.004	0.790	0.542 ± 0.003	0.518	0.646 ± 0.004	0.375	0.674 ± 0.007	0.792	0.672 ± 0.013	0.509	0.673 ± 0.004
	Predictive Uncertainty	0.379	0.807 ± 0.005	0.800	0.578 ± 0.006	0.514	0.674 ± 0.006	0.417	0.765 ± 0.014	0.811	0.528 ± 0.000	0.551	0.625 ± 0.005
	Aleatoric Uncertainty	0.379	0.813 ± 0.012	0.800	0.554 ± 0.004	0.514	0.659 ± 0.006	0.417	0.544 ± 0.022	0.811	0.702 ± 0.023	0.551	0.612 ± 0.010
	Epistemic Uncertainty	0.379	0.834 ± 0.014	0.800	0.580 ± 0.006	0.514	0.684 ± 0.005	0.417	0.805 ± 0.014	0.811	0.528 ± 0.000	0.551	0.638 ± 0.004
Ensemble + TTA	Image	0.379	0.775 ± 0.007	0.800	0.564 ± 0.004	0.514	0.653 ± 0.004	0.417	0.597 ± 0.035	0.811	0.701 ± 0.035	0.551	0.644 ± 0.008
	Binary mask	0.379	0.760 ± 0.015	0.800	0.582 ± 0.003	0.514	0.659 ± 0.007	0.417	0.770 ± 0.011	0.811	0.683 ± 0.005	0.551	0.724 ± 0.007

6. Results

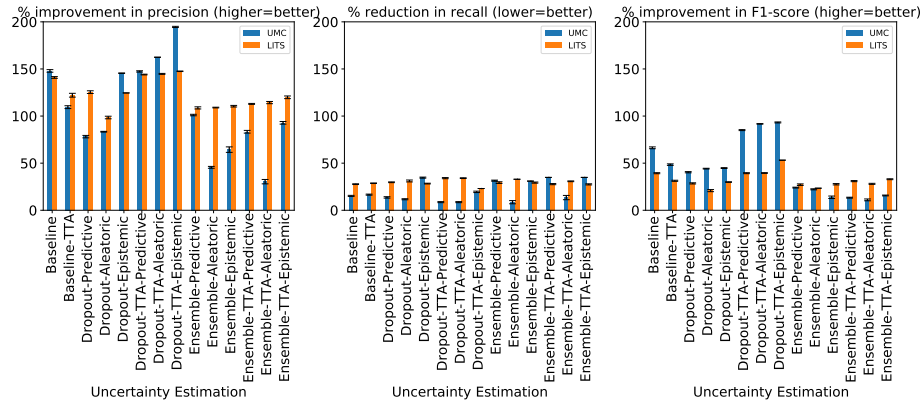
In Figure 6 we show the relative change in the lesion detection metrics of precision, recall, and F1 after false-positive classification for both datasets and all the uncertainty estimation methods using radiomics features computed from uncertainty maps (Figure 6a), image intensities (Figure 6b), and binary masks (Figure 6c). We show these detection metrics (before and after classification) in Table 3.

The relative change in lesion detection metrics for all uncertainty methods and both dataset, using the threshold-based method described in Nair et al. (2020) is shown in Figure 7.

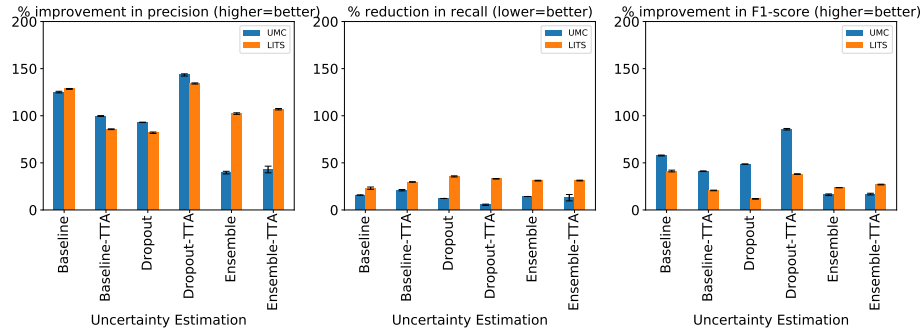
In Table 4 we show the changes in the lesion detection metrics (with respect to Table 3) when MR-CT cross-testing.

In Table 5 we show the changes in the lesion detection metrics (with respect to Table 3) when combined testing was performed.

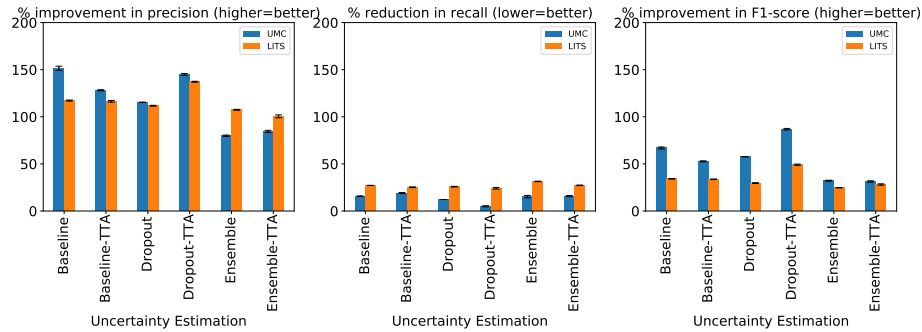
The feature importance scores for all uncertainty estimation methods have been shown in Appendix B (Figures 10 and 11).



(a) Radiomics features computed from uncertainty map



(b) Radiomics features computed from image intensities



(c) Radiomics features computed from binary mask

Figure 6: Relative change in precision, recall, and F1 metrics after false-positive classification. Figure 6a shows the relative change in the lesion detection metrics when radiomics features computed from uncertainty maps are used to perform false-positive classification. Figures 6b and 6c show the relative changes in lesion detection metrics when radiomics features are computed from image patches and binary masks, respectively.

7. Discussion

The results in Figure 6 show that the chosen uncertainty estimation method, by itself, did not contribute much towards reducing false positives. However, its impact was seen

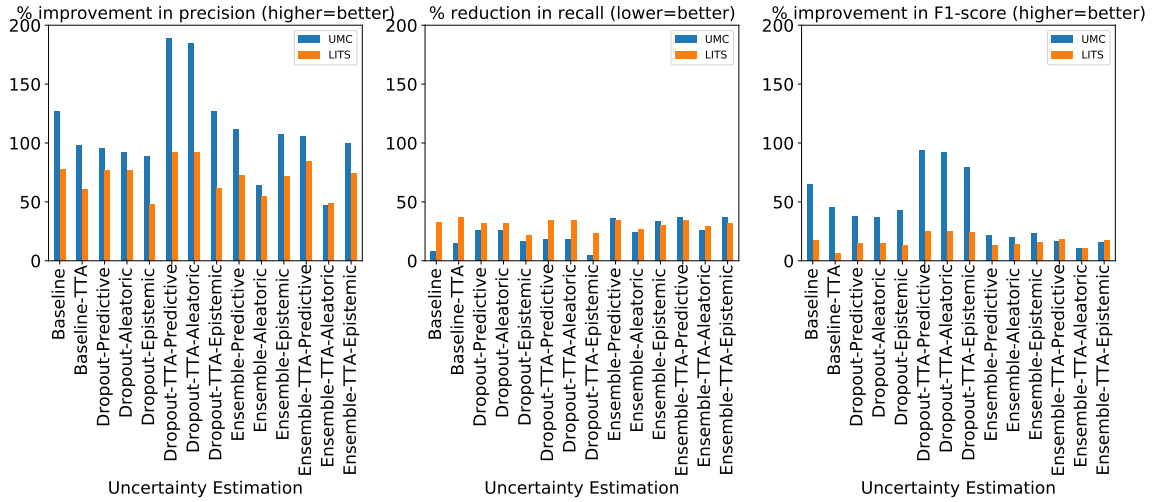


Figure 7: Relative change in the lesion detection metrics when a threshold on the log-sum of uncertainty values over the predicted lesion volume was used to filter false-positive lesions (Nair et al., 2020)

Table 4: Change in lesion detection metrics with cross-testing. The Kolmogorov-Smirnov test was used to check for statistical significance ($\alpha = 0.05$) of the change in lesion detection metrics, statistically significant changes have been marked with *

Configuration	Uncertainty Type	LiTS			UMC		
		Precision (% increase)	Recall (% increase)	F1 (% increase)	Precision (% increase)	Recall (% increase)	F1 (% increase)
Baseline	Predictive	-6.35*	-17.64*	-13.48*	-19.03*	11.80*	-6.09*
Baseline + TTA	Predictive	-15.07*	-7.31*	-10.76*	-2.72	5.84*	1.24
MC-Dropout	Predictive	-30.35*	3.16*	-12.66*	27.86*	2.26*	15.65*
	Aleatoric	-24.33*	8.87*	-7.42*	47.42*	-0.55	22.97*
	Epistemic	3.98	-22.09*	-13.89*	-8.39*	23.13*	6.91*
MC-Dropout + TTA	Predictive	-1.63*	-3.89*	-2.98*	-37.33*	6.97*	-24.92*
	Aleatoric	-0.92	-5.44*	-3.63*	-38.77*	6.97*	-25.62*
	Epistemic	1.37	-25.66*	-15.68*	-52.57*	24.29*	-32.94*
Ensemble	Predictive	-5.72*	2.15	-1.23*	-1.44	10.41*	5.12*
	Aleatoric	-32.99*	29.43*	-5.57*	0.61	-1.04	-0.07
	Epistemic	-10.51*	-0.71	-5.26*	13.79*	6.89*	10.11*
Ensemble + TTA	Predictive	-20.45*	-10.72*	-15.07*	7.24*	7.85*	7.60*
	Aleatoric	-22.14*	4.33	-8.56*	4.33	-3.22	1.05
	Epistemic	-14.26*	-9.31	-11.41*	0.31	7.85*	4.72*

on the neural network training, detection performance and feature selection for the false-positive classifier. Figures 6a, 6b, and 6c show similar trends for all uncertainty estimation methods. The uncertainty estimated by any given method does not play a major role in reducing false positives, since a similar performance is observed when radiomics features from binary masks and image intensities are used to train the classifier.

Figure 6 shows that the false-positive classification pipeline had a considerable impact on the performance metrics of the neural network for all uncertainty estimation methods and type of input features. While precision and F1-score improved, the extent of improvement depended on the class imbalance present in the data to train the classifier. Table 2 shows that the Ensemble and Ensemble+TTA configurations of the UMC dataset have a larger degree of class imbalance and a smaller training set size than the others. This was reflected

Table 5: Change in lesion detection metrics with a classifier trained with a combined training set. The Kolmogorov-Smirnov test was used to check for statistical significance ($\alpha = 0.05$) of the change in lesion detection metrics, statistically significant changes have been marked with *

Configuration	Uncertainty Type	LiTS			UMC		
		Precision (% increase)	Recall (% increase)	F1 (% increase)	Precision (% increase)	Recall (% increase)	F1 (% increase)
Baseline	Predictive	3.16*	-2.76	-0.49	-14.26*	9.32*	-3.91*
	Predictive	-0.81	-0.41	-0.54	4.53	5.26*	4.88*
MC-Dropout	Predictive	-4.77*	-4.74*	-4.74*	-5.99	10.16*	0.07
	Aleatoric	8.51*	-2.82	1.43	-8.44*	7.73*	-2.41
	Epistemic	-4.25*	-4.26*	-4.26*	-30.63*	52.99*	-0.11
MC-Dropout + TTA	Predictive	1.63*	-6.61*	1.11	-16.75*	7.46*	-8.52*
	Aleatoric	13.01*	-5.06*	1.63	-20.47*	4.47*	-11.76*
	Epistemic	-6.32*	-17.00*	12.54*	-31.84*	24.29*	-13.31*
Ensemble	Predictive	2.65	-1.75	-0.02	-6.80*	13.19*	3.84*
	Aleatoric	-1.54	-3.40*	-1.50*	-5.26*	2.60	-2.00
	Epistemic	6.04*	-1.07	1.69	9.10*	6.20*	7.62*
Ensemble + TTA	Predictive	0.40	2.07	1.37	2.85*	5.71*	4.53*
	Aleatoric	-1.41	0.00	-0.6	-0.77	-2.15	-1.22
	Epistemic	-0.66	1.37	0.54	-11.35*	12.14*	1.49

in the smaller improvement in the F1-score metric for these configurations. Future work could improve upon this, by implementing strategies that make classifiers more robust against class imbalance, like SMOTE (Chawla et al., 2002) or ADASYN (He et al., 2008). However, Figure 6c and Table 3, show a better performance for the UMC Ensemble and Ensemble+TTA when radiomics features computed from binary masks were used. This shows that shape-based features are more robust to class imbalance.

As expected, there is also a slight drop in recall, due to some true positive lesions being classified as false positives and filtered out. A trade-off between precision and recall can be made, for which the optimal classifier threshold may be chosen by the end-user.

Across both datasets, for almost all uncertainty estimation methods and types, the shape-based features of *Sphericity* or *Flatness* or *Elongation* were consistently ranked highest with respect to the LOCO feature importance scores. In Figure 6c we see almost no difference in the relative change in the metrics when intensity and texture features are excluded from the classifier training. The trends in Figures 6a, 6b, and 6c show us that features computed from the per-voxel uncertainty estimates do not play a major role in reducing false positives in this setting. In Figure 8, we show an example of how classifying a detected lesion as a false-positive depends on its size. This can be also seen in the relative improvement in the F1 score, which is almost the same for all uncertainty types for a given uncertainty estimation method.

The trends in Figure 7 show that threshold-based classification using the log-sum of per-voxel uncertainty estimates is an effective method to reduce false positives. Compared to Figure 6, the trends in the relative change of precision, recall, and F1 metrics are similar for the UMC dataset, while for the LiTS dataset they are worse. The log-sum computed over a predicted lesion strongly correlates with its size. This leads to most small segmentations being classified as false positives. This trend is also reported by Nair et al. (2020). Therefore, the log-sum aggregate is a proxy for lesion size, thereby strengthening our claim that shape-based features, and not uncertainty estimates themselves, play an important role in false-positive classification. We show evidence of this correlation in Appendix C. The presence

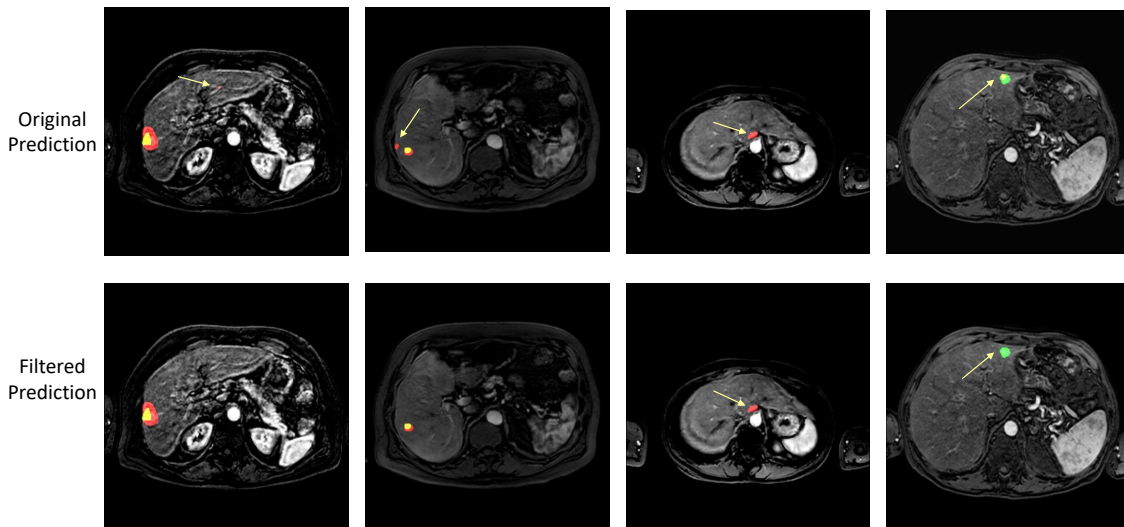
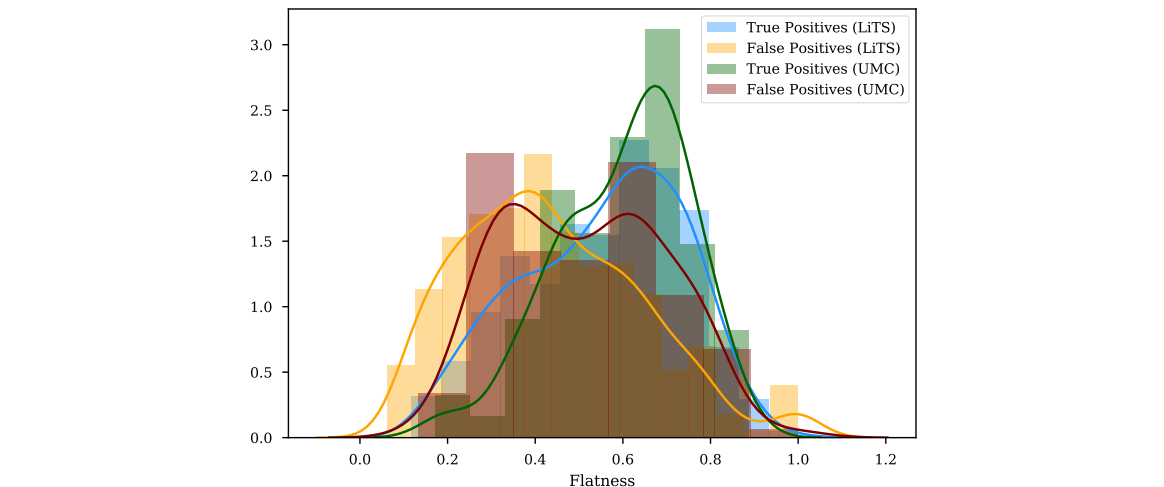
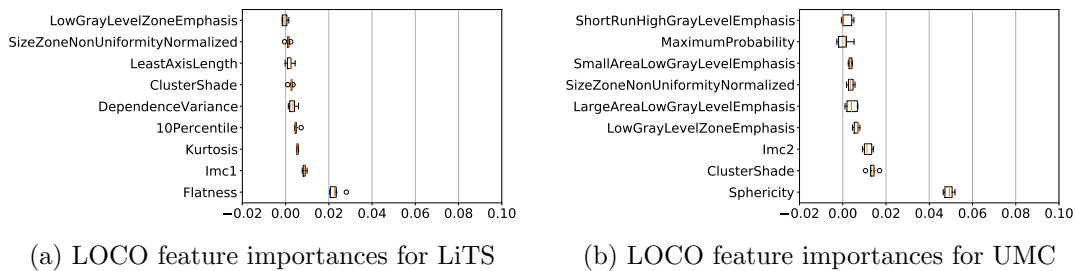


Figure 8: Influence of shape-based features on false-positive classification. The top row of the figure shows DCE-MR images from the UMC dataset overlaid with lesion segmentations computed by the segmentation network for the MC-Dropout configuration. The bottom row shows the corresponding lesion segmentation obtained after false-positive classification. The reference lesion segmentations are coloured green, the predicted lesion segmentations are red, and their overlap is yellow. The bright yellow arrow points to the predicted lesion in the original and filtered predictions. The first two columns show examples of small false positives filtered out by the classifier, while the relatively larger detections are correctly classified as true positives and retained. The third column shows an example of a slightly larger false positive misclassified as a true positive. The fourth column shows an example of a small detection overlapping with a true lesion incorrectly filtered out by the classifier.

of smaller true positive lesions in the LiTS dataset could be a reason for the worsening of performance with respect to the false positive classification pipeline.

The results on the cross-testing (Table 4) and combined-testing (Table 5) show that, although the underlying task is the same, the trained classifier cannot be shared between datasets. This is likely because the learned features are similar, but the thresholds learned do not generalize across datasets. In Figure 9 we show while correlated features (*Flatness* and *Sphericity*) play an important role in classifying lesion predictions in each of the datasets, the classifier produces poor results when used in a *cross* setting.

A limitation of our work was the inability to study the role uncertainty estimation can play in recovering false negatives. This was because the first step in the pipeline, identifying regions of interest, did not include regions for undetected lesions. Moreover, the uncertainty in the region of false negatives was low, indicating that the neural networks were incorrect with high confidence; a consequence of miscalibration. Miscalibration refers to the effect that the confidence assigned to an outcome by a classifier does not correspond to the



(c) Distribution plot of the *Flatness* feature in the training sets of LiTS and UMC data. *Flatness* and *Sphericity*, which are correlated features, are the most important features for the LiTS and UMC datasets respectively. We show that the different distributions of the *Flatness* features in the training sets of the two datasets lead to different optimal thresholds, which cause degradation in the metrics when we perform cross-testing.

Figure 9: Feature importance (ascending order, top to bottom) and distributions for the *Flatness* feature for the Baseline configuration

eventual prediction accuracy. This has been shown to occur in deep neural networks, where recent architectural advances have lead to improved classification performance but poorly modelled probability estimates (Guo et al., 2017). In our work, this was substantiated by the fact that the false-positive classifiers did not assign a high importance to the intensity-based features computed from the uncertainty maps, but rather focussed on shape-based features computed from the binary lesion masks (Figure 6c, Figure 10, and Figure 11). This phenomenon was demonstrated by Jungo et al. (2020), who showed that voxel-wise uncertainties were insufficient for detecting segmentation failures. Improving neural network calibration would improve lesion detection metrics and might produce uncertainty estimates that are more informative for false positive reduction.

Future work could consider a number of different steps to further improve our results. Our false-positive classification pipeline² can be used with any probabilistic classifier. Therefore, analysing the influence of uncertainty estimates computed by more recent

2. <https://github.com/ishaanb92/FPCPipeline>

methods (Van Amersfoort et al., 2020; Liu et al., 2022) on false-positive classification is an interesting research direction. Jungo et al. (2020) show that the calibration error of segmentation models goes down as training set size is increased. Therefore, the effect of training set size on our false-positive classification pipeline may be studied in a future work. Sander et al. (2020) train a second neural network using the uncertainty map computed by the segmentation network to correct local errors. A similar method to detect false positives is an interesting future direction because explicit feature computation is avoided. Finally, steps taken to make models robust to class-imbalance, such as selective sampling of slices or using weighted versions of loss functions during training, while improving detection metrics (especially recall), might harm calibration. Poor calibration is caused by poorly modelled probabilities, which makes it difficult to trust the uncertainty metrics computed using these probabilities. This is reflected in our results, where uncertainty estimates by themselves are not meaningful in reducing false-positive predictions. Therefore, an important direction for future research is developing segmentation models robust to class-imbalance, while still producing well-calibrated probabilities.

Our results on the LiTS and UMC datasets showed that per-voxel uncertainty estimates did not play a major role in false-positive classification. Similar to Jungo et al. (2020), we observed that different uncertainty estimation methods affected segmentation performance via their influence on the training dynamics and their regularization effects. Our results show that model ensembles perform the best with respect to the F1-score on both datasets.

8. Conclusion

We studied the efficacy of features computed from uncertainty estimates at reducing false positives by developing a classifier-based pipeline. We found that the relative improvement in the lesion detection metrics is mainly influenced by the class imbalance in the data used to train the classifier and the distribution of various shape-based features for all the uncertainty estimation methods we studied.

Acknowledgments

This work was financially supported by the project IMPACT (Intelligence based iMprovement of Personalized treatment And Clinical workflow supporT) in the framework of the EU research programme ITEA3 (Information Technology for European Advancement).

Ethical Standards

The work follows appropriate ethical standards in conducting research and writing the manuscript, following all applicable laws and regulations regarding treatment of animals or human subjects. The UMCU Medical Ethical Committee has reviewed this study and informed consent was waived due to its retrospective nature.

Conflicts of Interest

The authors declare no conflict of interest.

References

- Ishaan Bhat, Hugo J. Kuijf, Veronika Cheplygina, and Josien P.W. Pluim. Using Uncertainty Estimation To Reduce False Positives In Liver Lesion Detection. In *2021 IEEE 18th International Symposium on Biomedical Imaging (ISBI)*, pages 663–667, April 2021. doi: 10.1109/ISBI48211.2021.9434119. ISSN: 1945-8452.
- Patrick Bilic, Patrick Ferdinand Christ, Eugene Vorontsov, Grzegorz Chlebus, Hao Chen, Qi Dou, Chi-Wing Fu, Xiao Han, Pheng-Ann Heng, Jürgen Hesser, Samuel Kadoury, Tomasz Konopczynski, Miao Le, Chunming Li, Xiaomeng Li, Jana Lipková, John Lowengrub, Hans Meine, Jan Hendrik Moltz, Chris Pal, Marie Piraud, Xiaojuan Qi, Jin Qi, Markus Rempfler, Karsten Roth, Andrea Schenk, Anjany Sekuboyina, Eugene Vorontsov, Ping Zhou, Christian Hülsemeyer, Marcel Beetz, Florian Ettliger, Felix Gruen, Georgios Kaissis, Fabian Lohöfer, Rickmer Braren, Julian Holch, Felix Hofmann, Wieland Sommer, Volker Heinemann, Colin Jacobs, Gabriel Efrain Humpire Mamani, Bram van Ginneken, Gabriel Chartrand, An Tang, Michal Drozdal, Avi Ben-Cohen, Eyal Klang, Marianne M. Amitai, Eli Konen, Hayit Greenspan, Johan Moreau, Alexandre Hostettler, Luc Soler, Refael Vivanti, Adi Szeskin, Naama Lev-Cohain, Jacob Sosna, Leo Joskowicz, and Björn H. Menze. The Liver Tumor Segmentation Benchmark (LiTS). *arXiv:1901.04056 [cs]*, January 2019. URL <http://arxiv.org/abs/1901.04056>. arXiv: 1901.04056.
- Charles Blundell, Julien Cornebise, Koray Kavukcuoglu, and Daan Wierstra. Weight uncertainty in neural networks. In *Proceedings of the 32nd International Conference on International Conference on Machine Learning - Volume 37, ICML’15*, page 1613–1622. JMLR.org, 2015.
- Leo Breiman. Bagging Predictors. *Machine Learning*, 24(2):123–140, August 1996. ISSN 1573-0565. doi: 10.1023/A:1018054314350. URL <https://doi.org/10.1023/A:1018054314350>.
- Robin Camarasa, Daniel Bos, Jeroen Hendrikse, Paul Nederkoorn, M. Eline Kooi, Aad van der Lugt, and Marleen de Bruijne. A quantitative comparison of epistemic uncertainty maps applied to multi-class segmentation. *Machine Learning for Biomedical Imaging*, 1, 2021.
- Daniel C. Castro, Ian Walker, and Ben Glocker. Causality matters in medical imaging. *Nature Communications*, 11(1):3673, July 2020. ISSN 2041-1723. doi: 10.1038/s41467-020-17478-w.
- Nitesh V. Chawla, Kevin W. Bowyer, Lawrence O. Hall, and W. Philip Kegelmeyer. Smote: Synthetic minority over-sampling technique. *J. Artif. Int. Res.*, 16(1):321–357, June 2002. ISSN 1076-9757.
- Travers Ching, Daniel S. Himmelstein, Brett K. Beaulieu-Jones, Alexandr A. Kalinin, Brian T. Do, Gregory P. Way, Enrico Ferrero, Paul-Michael Agapow, Michael Zietz, Michael M. Hoffman, Wei Xie, Gail L. Rosen, Benjamin J. Lengerich, Johnny Israeli,

- Jack Lanchantin, Stephen Woloszynek, Anne E. Carpenter, Avanti Shrikumar, Jinbo Xu, Evan M. Cofer, Christopher A. Lavender, Srinivas C. Turaga, Amr M. Alexandari, Zhiyong Lu, David J. Harris, Dave DeCaprio, Yanjun Qi, Anshul Kundaje, Yifan Peng, Laura K. Wiley, Marwin H. S. Segler, Simina M. Boca, S. Joshua Swamidass, Austin Huang, Anthony Gitter, and Casey S. Greene. Opportunities and obstacles for deep learning in biology and medicine. *Journal of The Royal Society Interface*, 15(141):20170387, April 2018. ISSN 1742-5689, 1742-5662. doi: 10.1098/rsif.2017.0387.
- Grzegorz Chlebus, Andrea Schenk, Jan Hendrik Moltz, Bram van Ginneken, Horst Karl Hahn, and Hans Meine. Automatic liver tumor segmentation in CT with fully convolutional neural networks and object-based postprocessing. *Scientific Reports*, 8(1):15497, December 2018. ISSN 2045-2322. doi: 10.1038/s41598-018-33860-7. URL <http://www.nature.com/articles/s41598-018-33860-7>.
- Ozan Ciga, Tony Xu, Sharon Nofech-Mozes, Shawna Noy, Fang-I. Lu, and Anne L. Martel. Overcoming the limitations of patch-based learning to detect cancer in whole slide images. *Scientific Reports*, 11(1):8894, April 2021. ISSN 2045-2322. doi: 10.1038/s41598-021-88494-z. URL <https://www.nature.com/articles/s41598-021-88494-z>. Number: 1 Publisher: Nature Publishing Group.
- Stefan Depeweg, Jose Miguel Hernandez-Lobato, Finale Doshi-Velez, and Steffen Udluft. Decomposition of uncertainty in bayesian deep learning for efficient and risk-sensitive learning. In *Proceedings of the 35th International Conference on Machine Learning (ICML)*, volume 80, Stockholm, Sweden, 2018.
- Terrance DeVries and Graham W. Taylor. Leveraging Uncertainty Estimates for Predicting Segmentation Quality. July 2018. URL <http://arxiv.org/abs/1807.00502>. arXiv: 1807.00502.
- Zach Eaton-Rosen, Felix Bragman, Sotirios Bisdas, Sébastien Ourselin, and M. Jorge Cardoso. Towards Safe Deep Learning: Accurately Quantifying Biomarker Uncertainty in Neural Network Predictions. In Alejandro F. Frangi, Julia A. Schnabel, Christos Davatzikos, Carlos Alberola-López, and Gabor Fichtinger, editors, *Medical Image Computing and Computer Assisted Intervention – MICCAI 2018*, Lecture Notes in Computer Science, pages 691–699, Cham, 2018. Springer International Publishing. ISBN 978-3-030-00928-1. doi: 10.1007/978-3-030-00928-1_78.
- Stanislav Fort, Huiyi Hu, and Balaji Lakshminarayanan. Deep Ensembles: A Loss Landscape Perspective. *arXiv:1912.02757 [cs, stat]*, June 2020. URL <http://arxiv.org/abs/1912.02757>. arXiv: 1912.02757.
- Yarin Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *Proceedings of the 33rd International Conference on Machine Learning, ICML 2016*, volume 48 of *JMLR Workshop and Conference Proceedings*, pages 1050–1059. JMLR.org, 2016.
- Pierre Geurts, Damien Ernst, and Louis Wehenkel. Extremely randomized trees. *Mach. Learn.*, 63(1):3–42, April 2006. ISSN 0885-6125. doi: 10.1007/s10994-006-6226-1. URL <https://doi.org/10.1007/s10994-006-6226-1>.

- Simon Graham, Hao Chen, Jevgenij Gamper, Qi Dou, Pheng-Ann Heng, David Snead, Yee Wah Tsang, and Nasir Rajpoot. MILD-Net: Minimal information loss dilated network for gland instance segmentation in colon histology images. *Medical Image Analysis*, 52:199–211, February 2019. ISSN 1361-8415. doi: 10.1016/j.media.2018.12.001.
- Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q. Weinberger. On calibration of modern neural networks. In *Proceedings of the 34th International Conference on Machine Learning - Volume 70*, ICML’17, page 1321–1330. JMLR.org, 2017.
- Haibo He, Yang Bai, Eduardo A. Garcia, and Shutao Li. Adasyn: Adaptive synthetic sampling approach for imbalanced learning. In *2008 IEEE International Joint Conference on Neural Networks (IEEE World Congress on Computational Intelligence)*, pages 1322–1328, 2008. doi: 10.1109/IJCNN.2008.4633969.
- Carl-Johan Hoel, Tommy Tram, and Jonas Sjöberg. Reinforcement learning with uncertainty estimation for tactical decision-making in intersections. In *2020 IEEE 23rd International Conference on Intelligent Transportation Systems (ITSC)*, pages 1–7, 2020. doi: 10.1109/ITSC45102.2020.9294407.
- M J A Jansen, H J Kuijf, W B Veldhuis, F J Wessels, M S van Leeuwen, and J P W Pluim. Evaluation of motion correction for clinical dynamic contrast enhanced MRI of the liver. *Physics in Medicine & Biology*, 62(19):7556–7568, September 2017. ISSN 1361-6560. doi: 10.1088/1361-6560/aa8848.
- Mariëlle J. A. Jansen, Hugo J. Kuijf, Maarten Niekel, Wouter B. Veldhuis, Frank J. Wessels, Max A. Viergever, and Josien P. W. Pluim. Liver segmentation and metastases detection in MR images using convolutional neural networks. *Journal of Medical Imaging*, 6(4):1 – 10, 2019. doi: 10.1117/1.JMI.6.4.044003. URL <https://doi.org/10.1117/1.JMI.6.4.044003>.
- Alain Jungo, Fabian Balsiger, and Mauricio Reyes. Analyzing the Quality and Challenges of Uncertainty Estimations for Brain Tumor Segmentation. *Frontiers in Neuroscience*, 14:282, April 2020. ISSN 1662-453X. doi: 10.3389/fnins.2020.00282.
- Davood Karimi, Qi Zeng, Prateek Mathur, Apeksha Avinash, Sara Mahdavi, Ingrid Spadinger, Purang Abolmaesumi, and Septimiu E. Salcudean. Accurate and robust deep learning-based segmentation of the prostate clinical target volume in ultrasound images. *Medical Image Analysis*, 57:186–196, October 2019. ISSN 1361-8415. doi: 10.1016/j.media.2019.07.005.
- Alex Kendall, Vijay Badrinarayanan, and Roberto Cipolla. Bayesian SegNet: Model Uncertainty in Deep Convolutional Encoder-Decoder Architectures for Scene Understanding. In *Proceedings of the British Machine Vision Conference 2017*, page 57, London, UK, 2017. British Machine Vision Association. ISBN 978-1-901725-60-5. doi: 10.5244/C.31.57. URL <http://www.bmva.org/bmvc/2017/papers/paper057/index.html>.
- Diederik P. Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015.

- Diederik P. Kingma, Tim Salimans, and Max Welling. Variational dropout and the local reparameterization trick. In *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 2*, NIPS'15, page 2575–2583, Cambridge, MA, USA, 2015. MIT Press.
- Armen Der Kiureghian and Ove Ditlevsen. Aleatory or epistemic? Does it matter? *Structural Safety*, 31(2):105–112, March 2009. ISSN 0167-4730. doi: 10.1016/j.strusafe.2008.06.020.
- S. Klein, M. Staring, K. Murphy, M.A. Viergever, and J.P.W. Pluim. Elastix : a toolbox for intensity-based medical image registration. *IEEE Transactions on Medical Imaging*, 29(1):196–205, 2010. ISSN 0278-0062. doi: 10.1109/TMI.2009.2035616.
- Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017*, pages 6402–6413, 2017.
- Jing Lei, Max G'Sell, Alessandro Rinaldo, Ryan J. Tibshirani, and Larry Wasserman. Distribution-Free Predictive Inference for Regression. *Journal of the American Statistical Association*, 113(523):1094–1111, July 2018. ISSN 0162-1459, 1537-274X. doi: 10.1080/01621459.2017.1307116. URL <https://www.tandfonline.com/doi/full/10.1080/01621459.2017.1307116>.
- Christian Leibig, Vaneeda Allken, Murat Seçkin Ayhan, Philipp Berens, and Siegfried Wahl. Leveraging uncertainty information from deep neural networks for disease detection. *Scientific Reports*, 7(1):17816, December 2017. ISSN 2045-2322. doi: 10.1038/s41598-017-17876-z.
- Geert Litjens, Thijs Kooi, Babak Ehteshami Bejnordi, Arnaud Arindra Adiyoso Setio, Francesco Ciompi, Mohsen Ghafoorian, Jeroen A. W. M. van der Laak, Bram van Ginneken, and Clara I. Sánchez. A survey on deep learning in medical image analysis. *Medical Image Analysis*, 42:60 – 88, 2017. ISSN 1361-8415. doi: <https://doi.org/10.1016/j.media.2017.07.005>.
- Jeremiah Zhe Liu, Shreyas Padhy, Jie Ren, Zi Lin, Yeming Wen, Ghassen Jerfel, Zack Nado, Jasper Snoek, Dustin Tran, and Balaji Lakshminarayanan. A Simple Approach to Improve Single-Model Deep Uncertainty via Distance-Awareness, May 2022. URL <http://arxiv.org/abs/2205.00403>. arXiv:2205.00403 [cs, stat].
- Antonio Loquercio, Mattia Segu, and Davide Scaramuzza. A General Framework for Uncertainty Estimation in Deep Learning. *IEEE Robotics and Automation Letters*, 5(2): 3153–3160, April 2020. ISSN 2377-3766, 2377-3774. doi: 10.1109/LRA.2020.2974682.
- David J. C. MacKay. A Practical Bayesian Framework for Backpropagation Networks. *Neural Computation*, 4(3):448–472, 05 1992. ISSN 0899-7667. doi: 10.1162/neco.1992.4.3.448.

- Alireza Mehrtash, William M. Wells, Clare M. Tempany, Purang Abolmaesumi, and Tina Kapur. Confidence calibration and predictive uncertainty estimation for deep medical image segmentation. *IEEE Transactions on Medical Imaging*, page 1–1, 2020. ISSN 1558-254X. doi: 10.1109/tmi.2020.3006437. URL <http://dx.doi.org/10.1109/TMI.2020.3006437>.
- Raghav Mehta, Thomas Christinck, Tanya Nair, Paul Lemaitre, Douglas Arnold, and Tal Arbel. Propagating Uncertainty Across Cascaded Medical Imaging Tasks for Improved Deep Learning Inference. In *Uncertainty for Safe Utilization of Machine Learning in Medical Imaging and Clinical Image-Based Procedures*, volume 11840, pages 23–32. Springer International Publishing, Cham, 2019. ISBN 978-3-030-32688-3 978-3-030-32689-0. doi: 10.1007/978-3-030-32689-0_3.
- Tanya Nair, Doina Precup, Douglas L. Arnold, and Tal Arbel. Exploring uncertainty measures in deep networks for Multiple sclerosis lesion detection and segmentation. *Medical Image Analysis*, 59:101557, 2020. ISSN 1361-8415. doi: <https://doi.org/10.1016/j.media.2019.101557>. URL <https://www.sciencedirect.com/science/article/pii/S1361841519300994>.
- Radford M. Neal. *Bayesian Learning for Neural Networks*. Springer-Verlag, Berlin, Heidelberg, 1996. ISBN 0387947248.
- Matthew Ng, Fumin Guo, Labonny Biswas, Steffen E. Petersen, Stefan K. Piechnik, Stefan Neubauer, and Graham Wright. Estimating Uncertainty in Neural Networks for Cardiac MRI Segmentation: A Benchmark Study. *arXiv:2012.15772 [cs, eess]*, December 2020. URL <http://arxiv.org/abs/2012.15772>. arXiv: 2012.15772.
- Ipek Oguz, Aaron Carass, Dzung L. Pham, Snehashis Roy, Nagesh Subbana, Peter A. Calabresi, Paul A. Yushkevich, Russell T. Shinohara, and Jerry L. Prince. Dice Overlap Measures for Objects of Unknown Number: Application to Lesion Segmentation. In Alessandro Crimi, Spyridon Bakas, Hugo Kuijf, Bjoern Menze, and Mauricio Reyes, editors, *Brainlesion: Glioma, Multiple Sclerosis, Stroke and Traumatic Brain Injuries*, volume 10670, pages 3–14. Springer International Publishing, Cham, 2018. ISBN 978-3-319-75237-2 978-3-319-75238-9. doi: 10.1007/978-3-319-75238-9_1. Series Title: Lecture Notes in Computer Science.
- Yaniv Ovadia, Emily Fertig, Jie Ren, Zachary Nado, D. Sculley, Sebastian Nowozin, Joshua Dillon, Balaji Lakshminarayanan, and Jasper Snoek. Can you trust your model’s uncertainty? Evaluating predictive uncertainty under dataset shift. In *Advances in Neural Information Processing Systems 32*, pages 13991–14002. Curran Associates, Inc., 2019.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019.

- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-Net: Convolutional Networks for Biomedical Image Segmentation. In *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, pages 234–241, Cham, 2015. Springer International Publishing. ISBN 978-3-319-24574-4.
- Abhijit Guha Roy, Sailesh Conjeti, Nassir Navab, and Christian Wachinger. Inherent Brain Segmentation Quality Control from Fully ConvNet Monte Carlo Sampling. In *Medical Image Computing and Computer Assisted Intervention – MICCAI 2018*, pages 664–672, Cham, 2018. Springer International Publishing. ISBN 978-3-030-00928-1.
- Jörg Sander, Bob D. de Vos, Jelmer M. Wolterink, and Ivana Išgum. Towards increased trustworthiness of deep learning segmentation methods on cardiac MRI. In *Medical Imaging 2019: Image Processing*, volume 10949, pages 324 – 330. International Society for Optics and Photonics, SPIE, 2019. doi: 10.1117/12.2511699. URL <https://doi.org/10.1117/12.2511699>.
- Jörg Sander, Bob D. de Vos, and Ivana Išgum. Automatic segmentation with detection of local segmentation failures in cardiac MRI. *Scientific Reports*, 10(1):21769, December 2020. ISSN 2045-2322. doi: 10.1038/s41598-020-77733-4. Number: 1 Publisher: Nature Publishing Group.
- Suman Sedai, Bhavna Antony, Dwarikanath Mahapatra, and Rahil Garnavi. Joint Segmentation and Uncertainty Visualization of Retinal Layers in Optical Coherence Tomography Images Using Bayesian Deep Learning. In Danail Stoyanov, Zeike Taylor, Francesco Ciompi, Yanwu Xu, Anne Martel, Lena Maier-Hein, Nasir Rajpoot, Jeroen van der Laak, Mitko Veta, Stephen McKenna, David Snead, Emanuele Trucco, Mona K. Garvin, Xin Jan Chen, and Hrvoje Bogunovic, editors, *Computational Pathology and Ophthalmic Medical Image Analysis*, volume 11039, pages 219–227. Springer International Publishing, Cham, 2018. ISBN 978-3-030-00948-9 978-3-030-00949-6. doi: 10.1007/978-3-030-00949-6_26. Series Title: Lecture Notes in Computer Science.
- Philipp Seeböck, José Ignacio Orlando, Thomas Schlegl, Sebastian M. Waldstein, Hrvoje Bogunović, Sophie Klimescha, Georg Langs, and Ursula Schmidt-Erfurth. Exploiting Epistemic Uncertainty of Anatomy Segmentation for Anomaly Detection in Retinal OCT. *IEEE Transactions on Medical Imaging*, 39(1):87–98, January 2020. ISSN 0278-0062, 1558-254X. doi: 10.1109/TMI.2019.2919951. URL <http://arxiv.org/abs/1905.12806>. arXiv: 1905.12806.
- L. Smith and Y. Gal. Understanding Measures of Uncertainty for Adversarial Example Detection. In *UAI*, 2018.
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: A Simple Way to Prevent Neural Networks from Overfitting. *Journal of Machine Learning Research*, 15(56):1929–1958, 2014.

- Chao Tao, Ke Chen, Lin Han, Yulan Peng, Cheng Li, Zhan Hua, and Jiangli Lin. New one-step model of breast tumor locating based on deep learning. *Journal of X-Ray Science and Technology*, 27(5):839–856, 2019. ISSN 1095-9114. doi: 10.3233/XST-190548.
- Joost Van Amersfoort, Lewis Smith, Yee Whye Teh, and Yarin Gal. Uncertainty Estimation Using a Single Deep Deterministic Neural Network. In *Proceedings of the 37th International Conference on Machine Learning*, pages 9690–9700. PMLR, November 2020. URL <https://proceedings.mlr.press/v119/van-amersfoort20a.html>. ISSN: 2640-3498.
- Joost J.M. Van Griethuysen, Andriy Fedorov, Chintan Parmar, Ahmed Hosny, Nicole Aucoin, Vivek Narayan, Regina G.H. Beets-Tan, Jean-Christophe Fillion-Robin, Steve Pieper, and Hugo J.W.L. Aerts. Computational Radiomics System to Decode the Radiographic Phenotype. *Cancer Research*, 77(21):e104–e107, 2017. ISSN 0008-5472. doi: 10.1158/0008-5472.CAN-17-0339. Publisher: American Association for Cancer Research .eprint: <https://cancerres.aacrjournals.org/content/77/21/e104.full.pdf>.
- Janita E. Van Timmeren, Davide Cester, Stephanie Tanadini-Lang, Hatem Alkadhi, and Bettina Baessler. Radiomics in medical imaging—“how-to” guide and critical reflection. *Insights into Imaging*, 11(1):91, December 2020. ISSN 1869-4101. doi: 10.1186/s13244-020-00887-2. URL <https://insightsimaging.springeropen.com/articles/10.1186/s13244-020-00887-2>.
- Guotai Wang, Wenqi Li, Michael Aertsen, Jan Deprest, Sébastien Ourselin, and Tom Vercauteren. Aleatoric uncertainty estimation with test-time augmentation for medical image segmentation with convolutional neural networks. *Neurocomputing*, 338:34 – 45, 2019. ISSN 0925-2312. doi: <https://doi.org/10.1016/j.neucom.2019.01.103>.
- S. Kevin Zhou, Hayit Greenspan, Christos Davatzikos, James S. Duncan, Bram Van Ginneken, Anant Madabhushi, Jerry L. Prince, Daniel Rueckert, and Ronald M. Summers. A review of deep learning in medical imaging: Imaging traits, technology trends, case studies with progress highlights, and future promises. *Proceedings of the IEEE*, 109(5): 820–838, 2021. doi: 10.1109/JPROC.2021.3054390.
- Alex Zwanenburg, Stefan Leger, Martin Vallières, and Steffen Löck. Image biomarker standardisation initiative. *Radiology*, 295(2):328–338, May 2020. ISSN 0033-8419, 1527-1315. doi: 10.1148/radiol.2020191145. URL <http://arxiv.org/abs/1612.07003>. arXiv: 1612.07003.

Appendix A. Classifier Hyperparameters

Configuration	Input features	LiTS			UMC				
		Number of trees	Minimum samples for splitting	Minimum samples per leaf	Splitting criterion	Number of trees	Minimum samples for splitting	Minimum samples per leaf	Splitting criterion
Baseline	Predictive Uncertainty	250	10	1	Entropy	250	4	1	Entropy
	Image	1000	2	4	Gini	750	12	1	Gini
	Binary mask	750	10	4	Gini	250	2	8	Gini
Baseline + TTA	Predictive Uncertainty	250	10	4	Gini	750	12	1	Entropy
	Image	500	12	1	Entropy	250	4	1	Gini
	Binary mask	1000	2	8	Entropy	250	12	1	Entropy
Dropout	Predictive Uncertainty	500	10	2	Entropy	250	10	1	Entropy
	Aleatoric Uncertainty	250	2	4	Entropy	250	2	4	Gini
	Epistemic Uncertainty	250	8	2	Gini	750	2	1	Entropy
Dropout + TTA	Image	250	12	2	Entropy	1000	10	1	Gini
	Binary mask	750	10	4	Entropy	250	2	1	Entropy
	Predictive Uncertainty	750	12	4	Entropy	750	8	1	Entropy
Ensemble	Aleatoric Uncertainty	1000	2	4	Entropy	500	8	2	Entropy
	Epistemic Uncertainty	250	8	2	Entropy	250	12	2	Entropy
	Image	250	12	1	Entropy	250	2	2	Entropy
Ensemble + TTA	Binary mask	250	12	4	Entropy	250	12	2	Entropy
	Predictive Uncertainty	250	12	1	Entropy	500	4	1	Gini
	Aleatoric Uncertainty	250	8	2	Entropy	750	2	1	Entropy
Ensemble	Epistemic Uncertainty	500	2	4	Gini	250	2	1	Gini
	Image	250	8	2	Gini	1000	4	1	Entropy
	Binary mask	750	10	4	Entropy	250	12	4	Entropy
Ensemble + TTA	Predictive Uncertainty	250	12	1	Gini	250	10	4	Gini
	Aleatoric Uncertainty	500	12	2	Entropy	500	4	1	Entropy
	Epistemic Uncertainty	500	10	2	Entropy	750	2	10	Entropy
Ensemble + TTA	Image	250	2	4	Entropy	750	2	1	Entropy
	Binary mask	750	10	2	Gini	1000	2	4	Gini

Table 6: Classifier hyperparameters for the LiTS and UMC datasets

Appendix B. Feature Importance scores

Feature importance scores from the trained classifiers for LiTS and UMC dataset are shown in Figures 10 and 11.

Appendix C. Analysis of threshold-based classification

Figures 12 and 13 show the strong correlation between the log-sum of per-voxel uncertainties and lesion size for both datasets. Nair et al. (2020) reported that the log-sum computed for smaller lesions is higher, and most false positives are small, therefore log-sum can be used to successfully filter them out. Therefore, the uncertainty log-sum is a proxy for lesion size, as we can see from the high negative correlation between the log-sum and lesion diameter. Therefore, even though it seems that log-sum of the uncertainty over the lesion volume is used to filter false positives, given the high correlation with respect to the size, our observation about the influence of lesion shape on false-positive classification holds.

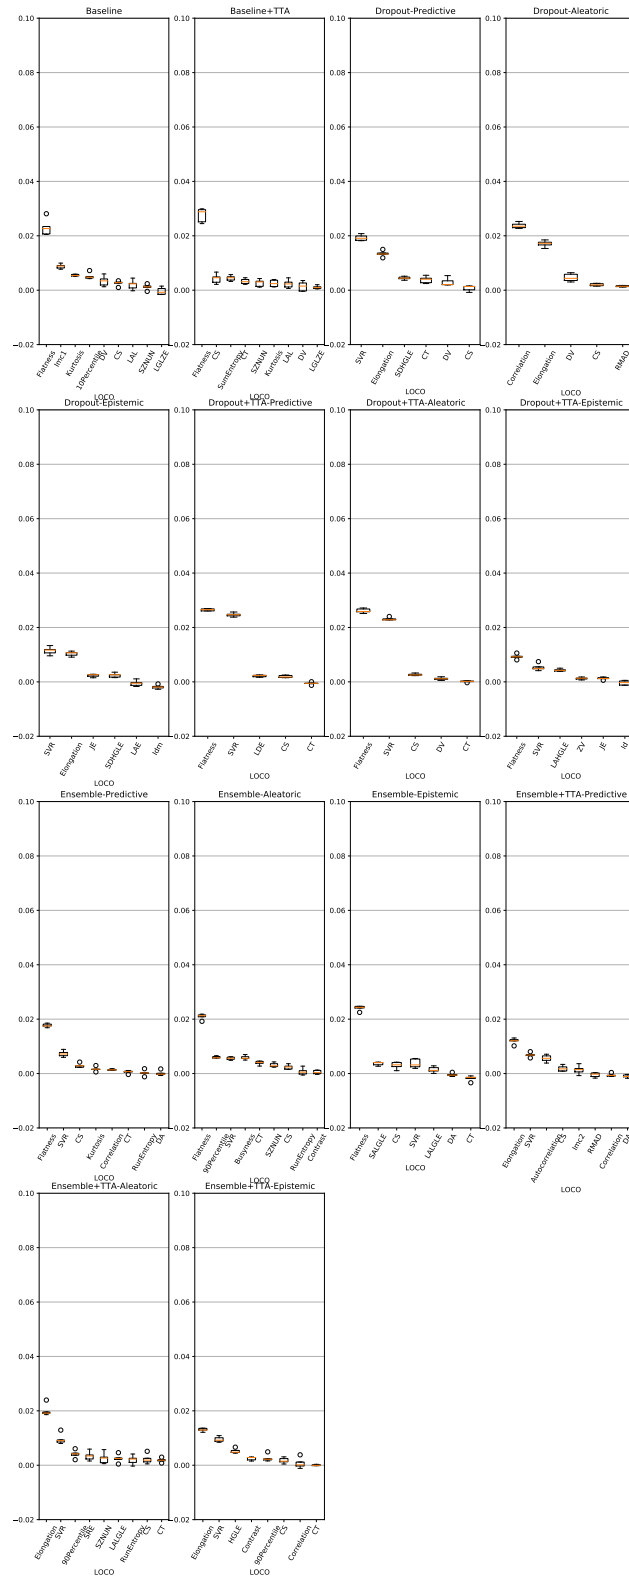


Figure 10: Feature importance scores (LOCO) for LiTS classifiers

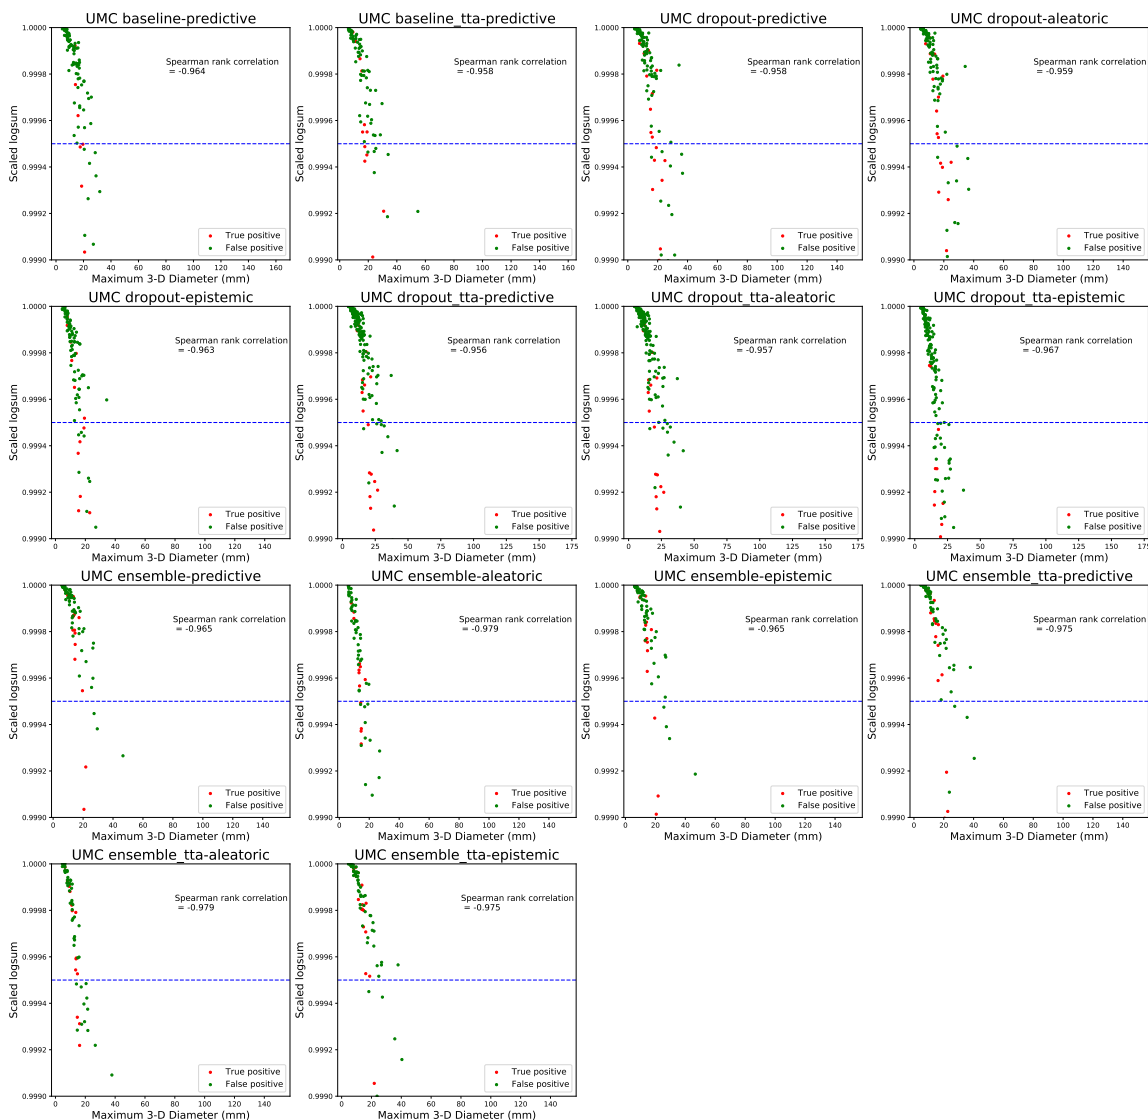


Figure 12: Scatter plot of the (scaled) log-sum aggregate calculated over predicted lesions, and maximum lesion diameter for the UMC test-set. The spearman rank correlation is reported. The detected lesions above the dotted blue line are classified as false positives.

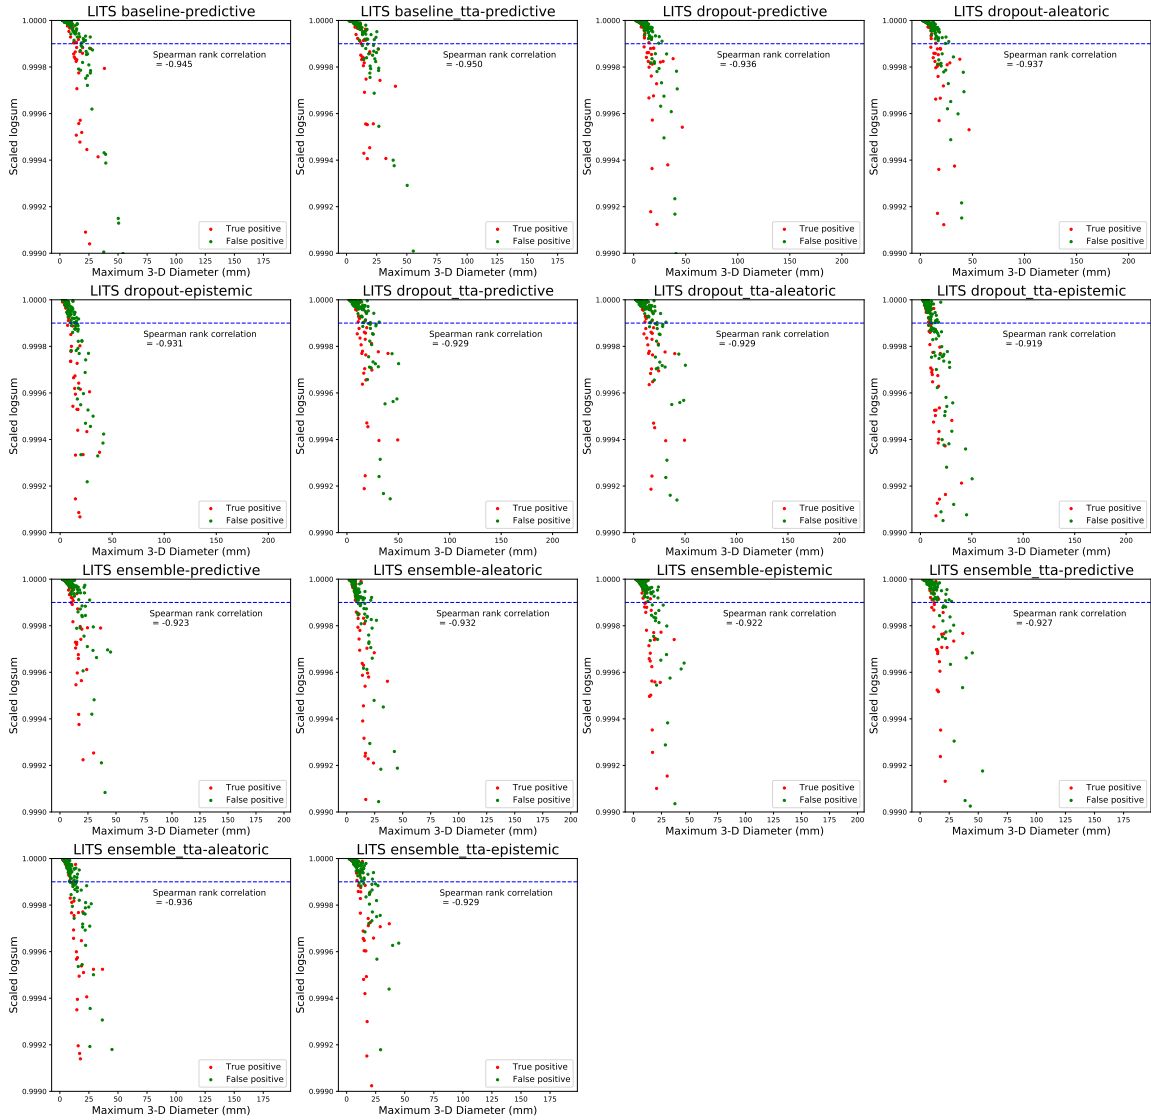


Figure 13: Scatter plot of the (scaled) log-sum aggregate calculated over predicted lesions, and maximum lesion diameter for the LiTS test-set. The spearman rank correlation is reported. The detected lesions above the dotted blue line are classified as false positives.