



## Focused Active Learning for Histopathological Image Classification

Arne Schmidt<sup>a,\*</sup>, Pablo Morales-Álvarez<sup>b,1</sup>, Lee AD Cooper<sup>c</sup>, Lee A Newberg<sup>d</sup>, Andinet Enquobahrie<sup>d</sup>, Rafael Molina<sup>a</sup>, Aggelos K Katsaggelos<sup>e</sup>

<sup>a</sup>Department of Computer Science and Artificial Intelligence, University of Granada, Granada, 18010 Spain.

<sup>b</sup>Department of Statistics and Operation Research, University of Granada, Granada, 18010 Spain

<sup>c</sup>Department of Pathology, Northwestern University, Chicago, IL, 60611 USA.

<sup>d</sup>Kitware Inc., Carrboro, NC, 27510 USA.

<sup>e</sup>Department of Electrical Computer Engineering, Northwestern University, Evanston, IL, 60208 USA.

### ARTICLE INFO

#### Article history:

Received 07 December 2022

Received in final form XX December 2022

Accepted XX December 2022

Available online XX December 2022

Communicated by XX

2000 MSC: 68T37, 68T05, 68T01, 62F15

**Keywords:** Active Learning, Cancer Classification, Histopathological Images, Bayesian Deep Learning

### ABSTRACT

Active Learning (AL) has the potential to solve a major problem of digital pathology: the efficient acquisition of labeled data for machine learning algorithms. However, existing AL methods often struggle in realistic settings with artifacts, ambiguities, and class imbalances, as commonly seen in the medical field. The lack of precise uncertainty estimations leads to the acquisition of images with a low informative value. To address these challenges, we propose Focused Active Learning (FocAL), which combines a Bayesian Neural Network with Out-of-Distribution detection to estimate different uncertainties for the acquisition function. Specifically, the weighted epistemic uncertainty accounts for the class imbalance, aleatoric uncertainty for ambiguous images, and an OoD score for artifacts. We perform extensive experiments to validate our method on MNIST and the real-world Panda dataset for the classification of prostate cancer. The results confirm that other AL methods are 'distracted' by ambiguities and artifacts which harm the performance. FocAL effectively focuses on the most informative images, avoiding ambiguities and artifacts during acquisition. For both experiments, FocAL outperforms existing AL approaches, reaching a Cohen's kappa of 0.764 with only 0.69% of the labeled Panda data.

© 2025 Elsevier B. V. All rights reserved.

\*Corresponding author. This work was supported by the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska Curie grant agreement No 860627 (CLARIFY Project), the US National Institutes of Health National Library of Medicine grant R01LM013523, and the project PID2022-140189OB-C22 funded by MCIN / AEI / 10.13039 / 501100011033. PMA acknowledges grant C-EXP-153-UGR23 funded by Consejería de Universidad, Investigación e Innovación and by ERDF Andalusia Program.

*e-mail:* [arne@decsai.ugr.es](mailto:arne@decsai.ugr.es) (Arne Schmidt), [pablmorales@decsai.ugr.es](mailto:pablmorales@decsai.ugr.es) (Pablo Morales-Álvarez), [lee.cooper@northwestern.edu](mailto:lee.cooper@northwestern.edu) (Lee AD Cooper), [lee.newberg@kitware.com](mailto:lee.newberg@kitware.com) (Lee A Newberg), [andinet.enqu@kitware.com](mailto:andinet.enqu@kitware.com) (Andinet Enquobahrie), [rms@decsai.ugr.es](mailto:rms@decsai.ugr.es) (Rafael Molina), [a-katsaggelos@northwestern.edu](mailto:a-katsaggelos@northwestern.edu) (Aggelos K Katsaggelos)

### 1. Introduction

Artificial Intelligence (AI) methods have obtained impressive results in digital pathology and in some cases, AI models even outperformed expert pathologists in cancer classification (Zhang et al., 2019; Hekler et al., 2019; Ehteshami Bejnordi et al., 2017). The hope is that AI can make the diagnosis more accurate, objective, reproducible, and faster in the future (Dimitriou et al., 2019).

To achieve this goal, trained, specialized AI models for

each subtask are required, for example for the quantification of tumor-infiltrating lymphocytes in lung cancer (Shvetsov *et al.*, 2022), metastasis detection of breast cancer in lymph nodes (Ehteshami Bejnordi *et al.*, 2017; Schmidt *et al.*, 2022) or Gleason grading of prostate cancer (Bulten *et al.*, 2022; Otálora *et al.*, 2020). Openly available, labeled datasets are limited to certain subtasks, and for many future applications, the aggregation of large amounts of labeled data remains challenging because the annotation requires medical experts. This makes the labeling process time-consuming and expensive. A common approach to labeling is to divide a region or Whole Slide Image (WSI) into small patches that are individually labeled (Dimitriou *et al.*, 2019). The model is then trained to make local patch-level predictions that can be aggregated for the final diagnosis. The problem with supervised deep learning methods is the need for large amounts of detailed (patch-level) annotations for training to obtain a satisfying predictive performance. To alleviate this burden, *semi-supervised learning* (Li *et al.*, 2018; Marini *et al.*, 2021; Lu *et al.*, 2020; Schmidt *et al.*, 2022; Otálora *et al.*, 2020) and *multiple instance learning* (Campanella *et al.*, 2019; Chikontwe *et al.*, 2020; Li *et al.*, 2021) have become major fields of interest in the recent years. [Another very promising approach to efficiently handle labeling resources is active learning.](#)

*Active Learning (AL)* describes machine learning methods that actively query the most informative labels. In the AL setting, the AI model starts training with a small set of labeled images and iteratively selects images from a large pool of unlabeled data. These selected images are labeled in each iteration by an 'oracle,' in our application a medical expert. [AL has several benefits:](#) (i) The model training and dataset creation go hand-in-hand. The performance of the model is constantly monitored to assess if the collected labeled data is enough - or if more labeled data is needed to reach the desired performance. (ii) The model looks for the most informative images automatically in the acquisition step. In other approaches that require labeling (e.g. semi-supervised learning), finding these informative, salient images may require a lot of manual searching. (iii)

AL is very data-efficient, while other paradigms such as multiple instance learning often require large datasets to compensate for missing instance labels (Campanella *et al.*, 2019).

[Notice that AL methods can be combined with other paradigms such as semi-supervised and multiple instance learning to leverage the advantages of multiple different perspectives \(Gao \*et al.\*, 2020; Huang \*et al.\*, 2021\). To narrow down our contribution, in this article we focus on improving existing probabilistic models in a pure AL setting, but we make several suggestions to combine the proposed FocAL method with other approaches as future work in section 4.](#)

**Related work** In AI research, different AL strategies have been proposed to determine the most informative images. Early approaches used the uncertainty estimation of support vector machines (Joshi *et al.*, 2009), Gaussian processes (Li and Guo, 2013) or Gaussian random fields (Zhu *et al.*, 2003) to rate the image informativeness. With the rise of deep learning, the focus shifted to Bayesian Neural Networks (BNNs) for AL (Gal *et al.*, 2017), which was adapted several times for histopathological images (Rączkowski *et al.*, 2019; Carse and McKenna, 2019; Meirelles *et al.*, 2022). This approach has the advantage of a probabilistic uncertainty estimation which is not only used for acquisition, but it is also crucial for diagnostic predictions in medical applications. BNNs allow the application of several different uncertainty-based acquisition functions, such as *BALD* (Houlsby *et al.*, 2011), *Max Entropy* (Rączkowski *et al.*, 2019; Shannon, 1948), and *Mean Std* (Alex Kendall and Cipolla, 2017). Other publications focus on the user interface and server application of AL (Lee *et al.*, 2021; Marée *et al.*, 2016) rather than the AL model itself. In the existing literature, the uncertainty estimation is often only used to determine the amount of new information in each image. We extend this idea by using complementary uncertainty measures to *avoid* labeling uninformative, ambiguous, or artifactual images. In digital pathology, several data-related challenges like artifacts, ambiguities, and the typical huge class imbalance hinder the application of AL (see "Problem analysis" paragraph below). Our proposed method tackles these problems successfully by precise

uncertainty estimations which leads to improved performance.

BNNs are not only of interest for the AL acquisition, their capacity to estimate the predictive uncertainty is highly important in safety-critical areas like medicine (Kwon *et al.*, 2020) or autonomous driving (Alex Kendall and Cipolla, 2017). The uncertainty estimation helps to distinguish confident predictions from risky ones. In our case, we aim to decompose uncertainty into *epistemic uncertainty* and *aleatoric uncertainty* describing the model and data uncertainty, respectively (Der Kiureghian and Ditlevsen, 2009). Epistemic uncertainty describes uncertainty in model parameters that can be reduced by training with additional labeled data. Therefore it can serve as a measure of informativeness in the active learning process. Unfortunately, the epistemic uncertainty is not only high for informative, in-distribution images but also for OoD images. In fact, epistemic uncertainty has recently been used explicitly for OoD detection (Xiao *et al.*, 2019; Mukhoti *et al.*, 2021; Nguyen *et al.*, 2022). Aleatoric uncertainty describes irreducible uncertainty in the data due to ambiguities that cannot be improved with additional labeling. Studies have shown that training with ambiguous data can harm the performance of the algorithm considerably if not taken into account (Gao *et al.*, 2017; Bernhardt *et al.*, 2022). In the Panda challenge, label noise associated with the subjective grading assigned by pathologists was considered to be a major problem (Bulten *et al.*, 2022).

To estimate these uncertainties with BNNs, Kendall and Gal (2017) proposed a network with two final probabilistic layers, corresponding to the two uncertainty measures. A theoretically sound, more stable, and efficient approach (relying on a single probabilistic layer) was proposed by Kwon *et al.* (2020). We base our BNN for uncertainty estimations on the latter method due to the mentioned advantages. In Section 3.2 we outline how the uncertainty estimations can be interpreted in the context of clinical applications like pathology.

To avoid acquiring image patches with artifacts, we apply OoD detection. Commonly, OoD data refers to [blue](#) data that originate from a different distribution than the training data (which are called “in-distribution”) (Sun *et al.*, 2022). In the

context of AL and pathology, we define the in-distribution as the distribution of patches containing (cancerous or non-cancerous) tissue. All the images with artifacts (such as pen markings, tissue folds, blood, or ink) (Kanwal *et al.*, 2022) will be considered OoD. These artifacts are inevitable in real-world data and there are several reasons to exclude them from the distribution of interest for acquisition: (i) It is impossible to learn all possible artifacts explicitly due to their wide variability. We argue that a model should reliably classify tissue and predict a high uncertainty for everything it does not know. (ii) It harms the performance of AL algorithms to acquire images with artifacts, as we show empirically in Section 3. (iii) The model should focus on learning what *is* cancerous instead of everything that *is not* cancerous. By learning cancerous patterns it automatically learns what is not cancerous (everything else).

In OoD detection, early methods used the depth (Johnson *et al.*, 1998; Ruts and Rousseeuw, 1996) or distance (Knorr and Ng, 1998, 1999) of datapoints, represented by low-dimensional feature vectors. With the rise of deep learning, OoD metrics were often applied to the features extracted by a deep neural network (Abati *et al.*, 2019; Sun *et al.*, 2022; Lee *et al.*, 2018). In line with previous research, we utilize extracted feature vectors and implement a density-based OoD scoring method (Breunig *et al.*, 2000) to detect artifacts in the data.

**Problem Analysis** Although AL has a huge potential for digital pathology, we analyze several challenges that hinder its application in practice:

- Medical imaging problems like pathology often have a high class imbalance. For example, in prostate cancer grading, the highest Gleason patterns may be underrepresented which needs to be taken into account during acquisition. Other AL algorithms treat each class equally and are not able to acquire a sufficient number of images of this underrepresented class in our experiments (Section 3).
- Many patches are ambiguous. There may be patches for which even subspecialists disagree on their label, or patches containing multiple classes. Assigning labels to these patches is difficult and may be detrimental to the

quality of the dataset and the algorithm’s performance. This not only slows the labeling process down, but it can also add noise to the training data as only one label per patch is assigned. In fact, label noise associated with the subjective grading assigned by pathologists was considered one key problem in the Panda challenge (Bulten *et al.*, 2022).

- WSIs can contain many different artifacts, such as pen markings, tissue folds, ink, or cauterized tissue. Existing AL algorithms often assign a high informativeness to these patches although they do not contain important information for model training, as we show empirically in the experimental section 3.

We want to stress that similar problems of class imbalance, ambiguities, and artifacts are present in many other medical imaging applications, such as CT scans for hemorrhage detection (Wu *et al.*, 2021), dermatology images for skin cancer classification (Esteva *et al.*, 2017) or retinal images for the detection of retinopathy (Gulshan *et al.*, 2016-12).

**Contribution** To address these challenges we propose Focused Active Learning (FocAL), a probabilistic deep learning approach that focuses on the underrepresented malignant classes while ignoring artifacts and ambiguous images. More specifically, we combine a Bayesian Neural Network (BNN) with Out of Distribution (OoD) Detection to estimate the three major elements of the proposed acquisition function. The *weighted epistemic uncertainty* rates the image informativeness, taking the class imbalance into account. The *aleatoric uncertainty* is used to avoid ambiguous images for acquisition. The *OoD score* helps to ignore outliers (like artifacts) that do not contribute information for the classification of tissue. **While BNNs uncertainties, OoD scores, and AL based on an acquisition function are known concepts/approaches in machine learning literature, our contribution lies in combining them in a principled and sound manner to tackle class imbalance, ambiguities, and artifacts for medical images.** We show empirically that the proposed acquisition function help focus on labeling salient, informative images while other methods often fail to address this

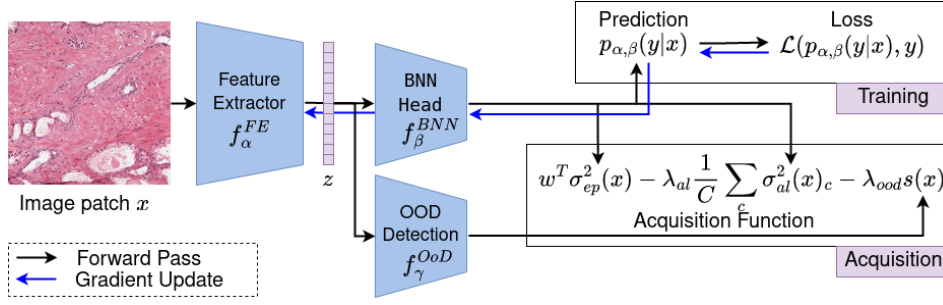
realistic data setting.

The article is structured as follows. We outline the theory of the proposed model, including the BNN and OoD components of the acquisition function in Section 2. In Section 3, we perform an illustrative MNIST experiment to analyze the behavior of existing AL approaches when artifacts and ambiguities are present. Furthermore, we demonstrate that each of our model components works as expected to avoid acquiring images with ambiguities and artifacts, overcoming the problems of the existing approaches. For the Panda prostate cancer [dataset](#), we perform an ablation study about the introduced hyperparameters, analyze the uncertainty estimations, and report in the final experiments that our method can reach a Cohen’s kappa of 0.763 with less than 1% of the labeled data (4400 labeled image patches). Finally, in Section 4 we conclude our article and give an outlook of future research.

## 2. Methods

Here we describe the three elements of FocAL: the feature extractor, the Bayesian Neural Network, and the Out-of-distribution score. The final paragraph of this section outlines the acquisition function and algorithm of the novel FocAL method. An overview of the model components is depicted in Figure 1.

**Active Learning (AL)** In AL we assume that at the beginning a small set of labeled data  $\mathcal{D}_{train} = \{x_i, y_i\}_{i=1, \dots, N}$  of images  $x_i$  and labels  $y_i$  and a pool of unlabeled data  $\mathcal{D}_{pool}$  is available. We assume each  $y$  to be a  $C$ -dimensional, one-hot encoded vector, where  $C$  stands for the number of classes. A machine learning model  $\mathcal{M}$  trains with a labeled set  $\mathcal{D}_{train}$  and then chooses a subset  $\mathcal{A} \subset \mathcal{D}_{pool}$  of unlabeled images to be labeled (by a specialist, such as a pathologist in the given application). The choice is made with the help of an *acquisition function*  $a(x, \mathcal{M})$  which estimates the informativeness of each image  $x$ . Probabilistic Active Learning The images with the highest acquisition scores are labeled and added to  $\mathcal{D}_{train}$ . Then, the model is retrained with the updated training set. This acquisition step is repeated iteratively such that the model performance increases



**Fig. 1. Model overview of the proposed FocAL method. It consists of three main components: feature extractor, BNN, and OoD detection (blue boxes). The figure shows how the different components are combined for training and acquisition.**

while more and more labeled data is aggregated.

**Feature Extraction** The feature extractor  $f_{\alpha}^{\text{FE}}$ , with model parameters  $\alpha$ , is the first component of the proposed model. We use a Convolutional Neural Network (CNN) to extract high-level 128-dimensional features  $z = f_{\alpha}^{\text{FE}}(x)$  from each image patch. The exact architecture of the feature extractor depends on the image data and task, see the implementation details for each experiment in Section 3. The feature extractor is trained during the AL process end-to-end with the BNN by gradient descent, see Figure 1. Although training the feature extractor is important for obtaining good final results, the emphasis of this article lies on the development of the BNN and OoD detection which perform the high-level reasoning, as described in the next paragraphs.


**Bayesian Neural Network (BNN)** The BNN model  $f_{\beta}^{\text{BNN}}$ , with model parameters  $\beta$ , allows probabilistic reasoning based on the extracted feature vectors  $Z = \{z_i\}$ . Note that the feature extractor and BNN head together could also be interpreted as a large, convolutional BNN with Bayesian layers near the output. Previous studies have shown that this combination of deterministic convolutions and Bayesian fully connected layers is the most effective way to introduce Bayesian uncertainty in the AL context (Zeng et al., 2018). Here, we treat the feature extractor and BNN separately, because the features are also used later for OoD detection. The BNN is not only able to make accurate classification predictions but can also estimate epistemic and aleatoric uncertainty. These estimated uncertainties will be further described below. They play a crucial role in the proposed acquisition function (see the last paragraph of this section).

The BNN in our model consists of two fully connected layers with 128 units and a final softmax output layer. In comparison to deterministic networks with weight parameters  $\omega$ , BNNs treat the model weights as random variables with a probability distribution  $p(\omega)$ . As the true posterior distribution  $p(\omega|\mathcal{D}_{\text{train}})$  is intractable, it has to be approximated. Following the success of similar approaches in recent studies (Gal et al., 2017; Alex Kendall and Cipolla, 2017; Kwon et al., 2020), we use variational inference to approximate the posterior distribution by a tractable variational distribution  $q_{\beta}(\omega)$ , where  $\beta$  describe the variational parameters of the distribution. Specifically, we define  $q$  as a product of independent Gaussian distributions over each model weight, parametrized by mean and variance. To approximate the real posterior, the minimization of the KL divergence  $KL(q_{\beta}(\omega)|p(\omega|\mathcal{D}_{\text{train}}))$  is achieved by maximizing the evidence lower bound (ELBO) utilizing the reparametrization trick (Kingma et al., 2015). Gradient descent allows the optimization of the BNN and the feature extractor end-to-end. We denote the ELBO loss function as  $\mathcal{L}(p_{\alpha,\beta}(y|x), y)$ . The predictive distribution  $p_{\alpha,\beta}(y|x)$  is obtained by applying the feature extractor  $z = f_{\alpha}^{\text{FE}}(x)$  and integrating the BNN through Monte Carlo sampling as:

$$\begin{aligned} p_{\beta}(y|z) &= \int p(y|\omega, z) q_{\beta}(\omega) d\omega \\ &\approx \frac{1}{T} \sum_{t=1}^T p(y|z, \omega_t), \end{aligned} \quad (1)$$

where we use Monte Carlo sampling by drawing  $T$  realizations  $\{\omega_t\}_{t=1,\dots,T}$  of the variational weight distribution. The  $\text{argmax}$  over classes of the vector  $p_{\beta}(y|z)$  defines the predicted class. For notational convenience, we will drop the parameters  $\alpha$  and





**Fig. 2. Illustration of aleatoric and epistemic uncertainty.** Two classes are depicted by blue circles and pink squares, separated by learned (solid) and optimal but unknown (dashed) decision boundaries. Samples not observed in training are shown in gray. A region with high aleatoric uncertainty (green) includes samples that are inseparable given available features. Uncertainty in this region cannot be reduced by further training with samples from this region. More discriminative features or additional features are needed to reduce this aleatoric uncertainty. A region with high epistemic uncertainty (brown) shows significant discrepancies between the learned and optimal decision boundaries due to lack of training samples in this region. Additional training with samples from this region can reduce this epistemic uncertainty and improve model accuracy.

$\beta$  when not needed. We show in Figure 1 an overview of the forward and backward pass with gradient descent.

**Bayesian Uncertainty Estimations** In Addition to the class prediction (eq. 1), the BNN is able to estimate the uncertainty, measured by the predictive covariance matrix  $\text{Cov}_{p(y^*|z^*, Z, Y)}(y^*)$ . This variance can be further decomposed into epistemic and aleatoric uncertainty.

*Epistemic uncertainty* (model uncertainty) measures the uncertainty introduced by the model parameters  $\omega$  and can be reduced with more labeled training data. A high epistemic uncertainty indicates a high informativeness of a given image. Unfortunately, using only the epistemic uncertainty for acquisition can lead to an unwanted outcome. If the data is contaminated with outliers such as artifacts, this can result in acquiring only outliers that do not contribute any value towards learning the classes of interest, as shown empirically in the experimental section 3.

*Aleatoric uncertainty* (data uncertainty) captures the uncer-

tainty inherent in the data due to ambiguities. It can not be reduced with more labeled data. Images with a high aleatoric uncertainty should be avoided during acquisition as the chance of mislabeling (due to ambiguous image content) or inherent data noise is higher for these images.

The estimating epistemic and aleatoric uncertainty, which are intuitively explained in Figure 2, can be calculated in different ways. Here, we follow the approach of Kwon *et al.* (2020), which does not require additional parameters, is numerically stable, and has a strong theoretical background. The covariance matrix is decomposed into

$$\begin{aligned} \text{Cov}_{p(y^*|z^*, Z, X)}(y^*) &= \underbrace{\frac{1}{T} \sum_{t=1}^T \{p(y^*|z^*, \omega_t) - \hat{p}(y^*|z^*)\} \otimes^2}_{:= \text{epistemic uncertainty}} \\ &+ \underbrace{\frac{1}{T} \sum_{t=1}^T \text{diag}\{p(y^*|z^*, \omega_t)\} - p(y^*|z^*, \omega_t) \otimes^2}_{:= \text{aleatoric uncertainty}} \end{aligned} \quad (2)$$

where  $\hat{p}(y^*|z^*) = \frac{1}{T} \sum_{t=1}^T p(y^*|z^*, \omega_t)$ ,  $\text{diag}\{p(y^*|z^*, \omega_t)\}$  is the diagonal matrix formed by the vector entries of  $p(y^*|z^*, \omega_t)$  in the diagonal, and the outer product  $v \otimes^2 = vv^T$ .

Note that both epistemic and aleatoric uncertainties are given as  $C \times C$  covariance matrices with values of the *uncertainty per class* on the diagonal. We define the  $C$ -dimensional vectors of class-wise uncertainties as  $\sigma_{ep}^2$  for the epistemic and  $\sigma_{al}^2$  for the aleatoric uncertainty.

**Out-of-Distribution (OoD) Detection** For the OoD detection, we use an unsupervised, density-based model  $f_\gamma^{\text{OoD}}$  with parameters  $\gamma$ <sup>1</sup>, based on the extracted features  $z$ . Instead of having a binary decision (in/out of distribution), we want to score each feature vector of the unlabeled images with a Local Outlier Factor (LOF) (Breunig *et al.*, 2000). The LOF is based on the  $k$ -nearest neighbors  $N_k(z)$  of a vector  $z$  and the *local reachability density*  $lrd_k(z)$ , a density measure based on the distance to the  $k$ -nearest-neighbors.

<sup>1</sup>Note that the parameters  $\gamma$  consist of the locations and densities of the currently labeled feature vectors. These are not model parameters in the strict sense but we follow this notation for coherence.

The LOF of a vector  $z$  is defined as

$$\text{LOF}_k(z) = \frac{1}{|N_k(z)|} \sum_{z' \in N_k(z)} \frac{\text{lrd}_k(z')}{\text{lrd}_k(z)} \quad (3)$$

with hyperparameter  $k$  which should be set to the minimum amount of expected datapoints in a cluster (Breunig *et al.*, 2000). Intuitively, the LOF is high if a feature vector lies in a region with a lower density than its neighbors (indicating an outlier). If the region of a feature vector has the same density as its neighbors, its LOF is close to 1. The upper bound depends on the characteristics of the data, i.e. the distances between feature vectors. Empirically we observed that scaling the LOF by 0.1 leads to an OoD score that is in same the range as the other uncertainty measures (epistemic and aleatoric uncertainty). Therefore, we define the outlier scoring function as

$$s(x) = 0.1 \text{LOF}_k(f^{\text{FE}}(x)). \quad (4)$$

Note that the scaling factor does not introduce an additional hyperparameter. It is inherently tuned by manipulating the weighting factor  $\lambda_{ood}$  in eq. 5 for which we perform experiments in section 3.2. The hyperparameter  $k$  can be set to a rough estimate of the minimum number of initial images  $x \in \mathcal{D}_{train}$  that are not affected by ambiguities or artifacts. We set it to  $k = 10$  for MNIST and  $k = 50$  for the Panda dataset.

**Focused Active Learning (FocAL)** We propose an acquisition function that combines the uncertainty-based measures and the OoD scoring discussed above:

$$a(x, \mathcal{M}) = \underbrace{w^T \sigma_{ep}^2(x)}_{\text{weighted ep. unc.}} - \lambda_{al} \underbrace{\frac{1}{C} \sum_c \sigma_{al}^2(x)_c}_{\text{aleatoric unc.}} - \lambda_{ood} \underbrace{s(x)}_{\text{OoD score}} \quad (5)$$

with the (calculated) class weight vector  $w = [w_1, w_2, \dots, w_C]^T$  (see eq. 6) and hyperparameters  $\lambda_{al} \in \mathbb{R}^+$ ,  $\lambda_{ood} \in \mathbb{R}^+$ . The images with the highest scores of the acquisition function are selected for labeling in each step. Each component fulfills a specific task in the acquisition process:

*1) Weighted Epistemic Uncertainty* With the BNN we can calculate the informativeness of each unlabeled image measured by the epistemic uncertainty. This measure has a sound theoretical background and a proven track record in practice. The

advantage of BNNs is that this approach estimates the epistemic uncertainty for each class independently. Existing approaches based on epistemic uncertainty often just take the sum over all classes. For our proposed model, we want to emphasize the informativeness of underrepresented classes. Therefore, we multiply the epistemic uncertainty of each class with a class-weight  $w_c$  which is calculated by

$$w_c = \frac{N_{train}}{N_c * C} \quad (6)$$

with  $N_{train}$  being the number of labeled images,  $N_c$  being the number of labeled images of class  $c$  and  $C$  the number of classes. Initially, the class weights are calculated based on all available labels  $\mathcal{D}_{train}^0$ . Then, the class weights are recalculated at each acquisition step  $s$ , depending on the given label distribution of the current set  $\mathcal{D}_{train}^s$ . This allows the algorithm to automatically adjust to class imbalances in  $\mathcal{D}_{train}$ .

*2) Aleatoric Uncertainty* The BNN measures the aleatoric uncertainty of each unlabeled image. We down-weight the informativeness of images based on their aleatoric uncertainty estimate to avoid labeling ambiguous patches. Although in the existing literature the aleatoric uncertainty is described as a measure of data uncertainty, we found that it does not capture data uncertainty for images with a different appearance (OoD). Therefore, an additional measure for the OoD images is necessary.

*3) OoD Score* To avoid the acquisition of outliers we apply an OoD algorithm on extracted image features. We down-weight the informativeness of images with a high OoD score. This allows the network to focus on the in-distribution data and acquire informative image patches.

The active learning procedure is summarized in Algorithm 1. After the acquisition steps are completed, the trained models are not only able to give accurate classification predictions for each new test image but also the epistemic and aleatoric uncertainty and OoD score which is very useful for the pathologist in the diagnostic process. In the regions where all three uncertainty measures are low, the prediction is reliable and the pathologist can trust the classification result.

**Algorithm 1** FocAL algorithm

---

**Input:** Start training set  $\mathcal{D}_{train}^0$ , pool of unlabeled data  $\mathcal{D}_{pool}^0$ , models  $f_{\alpha}^{FE}, f_{\beta}^{BNN}, f_{\gamma}^{OoD}$ , number of acquisition steps  $S$ .  
**Output:** Optimal model parameters  $\alpha, \beta, \gamma$ ; training dataset  $\mathcal{D}_{train}^S$

**for**  $s = 0$  to  $S$  **do**  
  Train  $f_{\alpha}^{FE}, f_{\beta}^{BNN}$  with  $\mathcal{D}_{train}^s$ .  
  Predict features  $Z_{train} \leftarrow f_{\alpha}^{FE}(X_{train})$   
  Update  $f_{\gamma}^{OoD}$  with  $Z_{train}$   
  Predict features  $Z_{pool} \leftarrow f_{\alpha}^{FE}(X_{pool})$   
  Estimate unc.  $\sigma_{ep}^2(X_{pool}), \sigma_{al}^2(X_{pool}) \leftarrow f_{\beta}^{BNN}(Z_{pool})$   
  Estimate OoD scores  $s(X_{pool}) \leftarrow f_{\gamma}^{OoD}(Z_{pool})$  (eq. 4)  
  Select acq. set  $A^s$  with  $a(X_{pool}, \{f_{\alpha}^{FE}, f_{\beta}^{BNN}, f_{\gamma}^{OoD}\})$  (eq. 5)  
  Label  $A^s$   
  Add  $A^s$  to labeled data  $\mathcal{D}_{train}^{s+1} \leftarrow \mathcal{D}_{train}^s \cup A^s$   
  Remove  $A^s$  from pool  $\mathcal{D}_{pool}^{s+1} \leftarrow \mathcal{D}_{pool}^s \setminus A^s$   
**end for**  
**return** Optimal model parameters  $\alpha, \beta, \gamma$ ; training dataset  $\mathcal{D}_{train}^S$ .

---

### 3. Experiments

For the empirical validation, we use two publicly available datasets, Panda and MNIST. In the MNIST dataset, we artificially introduce ambiguities and artifacts to demonstrate the functionality of the different model components. The proposed FocAL strategy avoids ambiguities, and artifacts and outperforms other approaches. In the second experiment with the Panda dataset, we apply the model to real-world data. We perform a study about the introduced hyperparameters, analyze the uncertainty estimations of the model, and compare different AL methods that were used in the recent literature:

**RA** (Gal et al., 2017; Carse and McKenna, 2019; Rączkowski et al., 2019): Random Acquisition (RA) is a simple baseline method that uses a uniform distribution over the images instead of an informativeness measure.

**EN** (Gal et al., 2017; Zeng et al., 2018; Rączkowski et al., 2019): The maximum entropy (EN) is used for acquisition. As entropy is a measure of new information, the most informative images should be obtained.

**BALD** (Houlsby et al., 2011; Gal et al., 2017; Carse and McKenna, 2019; Rączkowski et al., 2019): Acquisition with Bayesian Active Learning by Disagreement (BALD). The idea of BALD is to select the images that maximize the mutual in-

formation between predictions and model posterior. This is one of the most popular methods adapted in recent literature.

**MS** (Gal et al., 2017): The Mean Std (MS) measures the uncertainty by the average standard deviation of the predictive distribution. The idea is to acquire images with the least confident predictions.

**EP** (Nguyen et al., 2019): BNN using only the epistemic uncertainty as calculated in equation 2. This method is similar to FocAL, but without weighting the epistemic uncertainty and without the aleatoric uncertainty and OoD scoring.

**FocAL**: The proposed FocAL method as described in Section 2.

The reason to choose these baselines is that they all share the rationale of “acquiring the samples with highest uncertainty”, and they just differ on how this uncertainty is calculated in practice. This is relevant because the goal of our work is precisely to show that these measurements of uncertainty are not enough to deal with ambiguities and artifacts present in real world data, such as in the medical imaging domain. In contrast, we will see that FocAL copes with this scenario more satisfactorily.

#### 3.1. MNIST

The goal of this experiment is to illustrate the functionality of FocAL in a controlled environment with an intuitive dataset with artificial artifacts and ambiguities.

**Dataset** The well-known MNIST dataset (Deng, 2012) contains 60,000 training and 10,000 test images of handwritten digits with 28x28 greyscale pixels. Of the original training split, we randomly sample 2000 images of which 20 images are initially labeled ( $\mathcal{D}_{train}$ ) while 1980 images remain initially unlabeled ( $\mathcal{D}_{pool}$ ). This relatively low number of images is chosen for better visualization of the data distribution (Fig. 4 and 5). In each acquisition step, 10 images are acquired (labeled) until  $\mathcal{D}_{train}$  contains 200 labeled images. We also sample 200 validation images (from the training split) to use a reasonably small validation set in the context of limited labeled data (Oliver et al., 2018). For testing, we use the original test split of 10000 images. Furthermore, we adjust this dataset to mimic the problems in digital pathology that we want to tackle. The class imbalance is obtained by reducing the original 10 classes to only 3 classes:

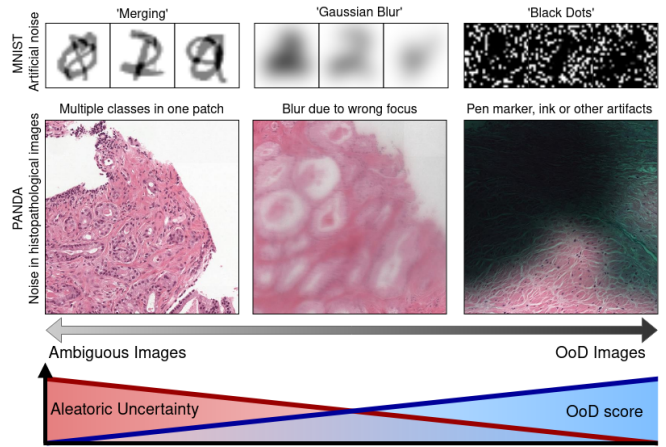


Digit '0', digit '1', and 'all other digits'. The classes '0' and '1' represent the malignant classes (10% portion of the whole dataset each) while 'all other digits' represent the healthy tissue (80% portion of the whole dataset).

**Artifacts and Ambiguities** The artificial artifacts and ambiguities are obtained by adding perturbations to the input images, as depicted in Figure 3. We use three different perturbations that mimic artifacts and ambiguities in histopathological images: *Black dots* are randomly added to 75% of the total image pixels by setting the greyscale value to 0. This simulates pen marker or ink in histopathological images that can cover large parts of image patches. *Gaussian blur* filter with a standard deviation of  $\sigma = 4$  is used to simulate the blur caused by wrong focus. *Merging* by randomly blending one image with another image of a different class together leads to ambiguous images with two different plausible labels (while maintaining the original one-hot encoded label). This simulates ambiguities by the presence of two cancerous classes in one image patch or edge cases with unclear ground truth label. Each of the three perturbations is applied to a total of 200 unlabeled images.

**Implementation Details** As the images of MNIST are very small, we use a simple feature extractor consisting of one convolutional layer with 4 filters (stride 3x3), max pooling (stride 2x2), and one fully connected layer with 128 units. The BNN consists of two fully connected layers with 128 units each and a final softmax layer with three output units, corresponding to the three classes. We use the cross-entropy loss and the Adam optimizer (Kingma and Ba, 2015) for 1000 epochs before each acquisition step. The learning rate is set to 0.0001 and multiplied by 0.5 if the validation accuracy does not increase for 50 epochs. The combination of a high number of epochs and learning rate reduction assures complete convergence at each acquisition step. We experimentally set the weight factors to  $\lambda_{al} = 0.5$  and  $\lambda_{ood} = 2.0$  since it showed the best results (tested: 0.5, 1.0, and 2.0 for each hyperparameter). Note, that an extensive ablation study of these hyperparameters is included in Section 3.2 for the Panda dataset.


**Data Distribution** First, we empirically analyze the data dis-



**Fig. 3.** Images with ambiguities and artifacts that should be avoided during AL acquisition. The top row shows MNIST images with the artificial noise types 'Merging', 'Gaussian Blur', and 'Black Dots'. They simulate the artifacts and ambiguities encountered in histopathological images (bottom row) in the Panda dataset. The left Panda patch contains two different classes (Gleason Grade 3 and 4), the middle patch is blurry due to wrong microscope focus, and the right patch is covered by pen marker, obscuring most tissue parts. Although a clear categorization is difficult, we propose the following scale: The images on the left side show ambiguities but the images are in-distribution because their appearance (color distribution and shapes) is normal. The images on the right side can be considered OoD because the color distribution and shapes substantially differ from the 'normal' images of interest. The blurry images are in between these two extremes as the color distribution and appearance is slightly OoD and they contain ambiguities due to blurry edges and patterns. We will see that with the proposed FocAL method, the shown images are avoided thanks to the aleatoric uncertainty and OoD score.

tribution with respect to the artificial artifacts and ambiguities. For this purpose, we plot the feature vectors  $z$  after the complete AL process using the FocAL method. We reduce the features to a two-dimensional distribution with t-SNE and depict the data in Figure 4. The distribution empirically reflects the categorization shown in Fig. 3. The images with 'black dots' are OoD because they are far away from the data distribution of interest. The images with 'Gaussian blur' are partially OoD and the 'merged' images are completely in-distribution. Similarly, ambiguities can be identified. The 'merged' images and a part of the images with 'Gaussian blur' are ambiguous and therefore close to the class boundaries. The images with 'Black dots' are not ambiguous. Apart from this data-related observation, the figure shows that the model learns to separate the classes during the active learning procedure. This class-separation is a necessary step for a good final classification.

**Acquisition** Figure 5 shows the acquisition behavior of FocAL and competing methods. It confirms that the FocAL model



**Fig. 4.** Feature distribution of the 2000 images ( $X_{train}$  and  $X_{pool}$ ) after the last acquisition step with 200 labeled images. Each point represents the feature vector  $z$  of an MNIST image, reduced to two dimensions by t-SNE. The distribution supports our categorization of artifacts and ambiguities (Fig. 3). The images with 'black dots' (depicted as squares) are OoD while the 'merged' images are ambiguous and therefore close to the class boundaries. Blurred images show both characteristics (OoD and ambiguities) as some images are far away from the distribution of interest, while others lie close to the class boundaries.

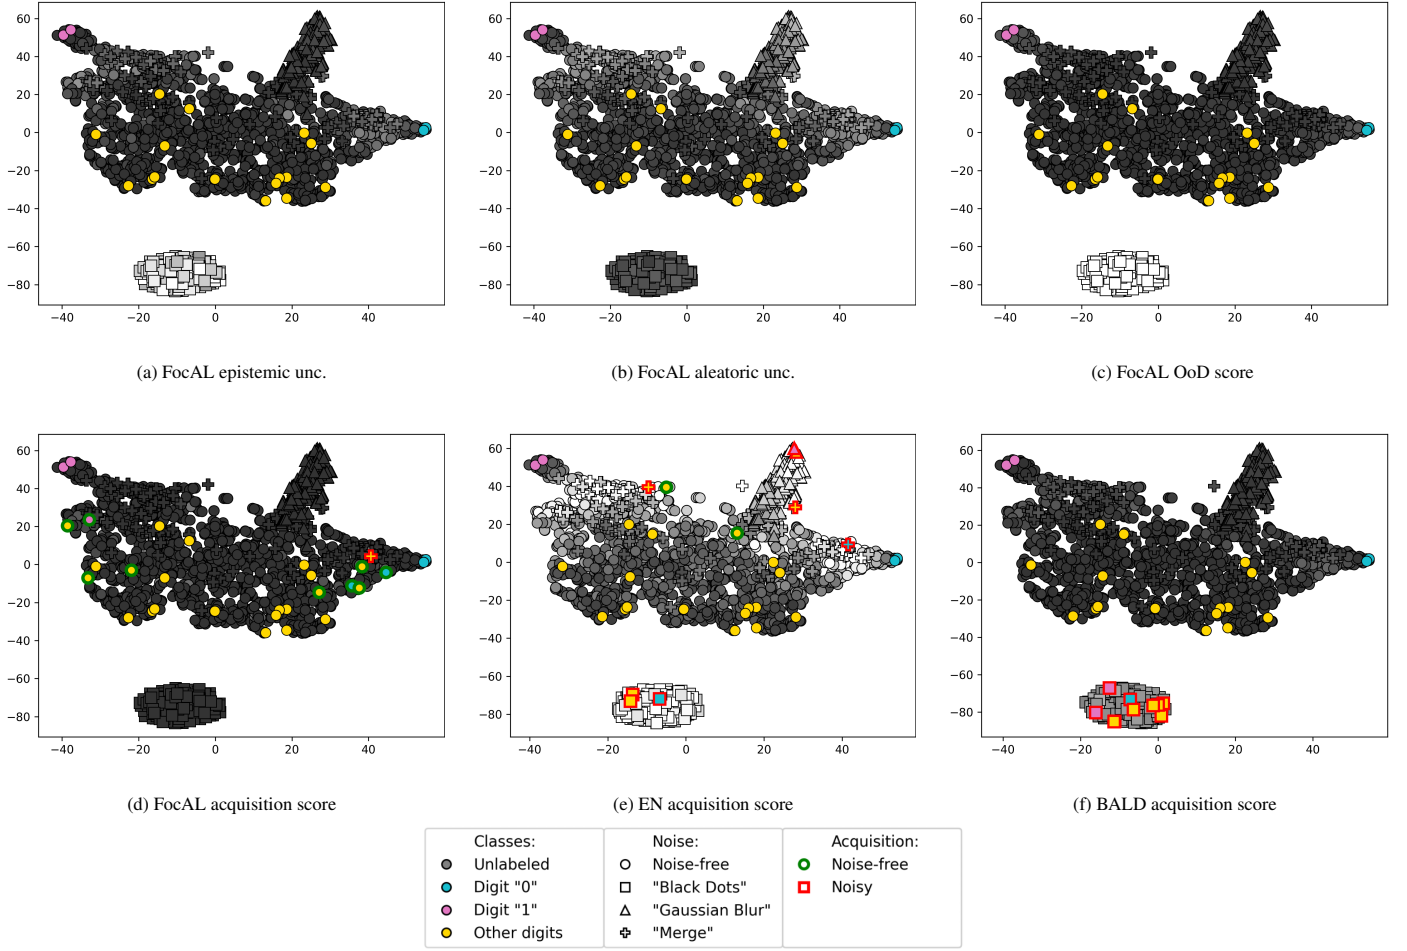
components work as expected in practice. The weighted episodic uncertainty of the FocAL method (5a) is high (bright greyscale color) for images at the class boundaries, but also for noisy images (especially of the noise type 'Black dots' and 'Merging'). This means that it is a good measure of informativeness, but 'distracted' easily by artifacts and ambiguities. The aleatoric uncertainty (5b) captures images with 'Merging' and 'Gaussian blur' while the OoD score (5c) highlights the images with 'Black dots'. These images with ambiguities and artifacts are avoided during acquisition. As a result, the FocAL method (5d) acquires only 1 ambiguous image while 9 acquired images are informative and contain several images of the minority classes '0' and '1'. Furthermore, the acquisition analysis highlights the problems of existing AL methods. EN in Figure 5e and BALD in Figure 5f are highly 'distracted' by images with ambiguities and artifacts and do not acquire any informative data in this step. Ambiguous images are close to class boundaries while OoD images have a 'novel' image content due to their different appearance. Therefore, the acquisition scores of other AL methods for these images are usually high. Although here we depict only the data distribution of features from the first acquisition step, these observations are representative for all further acquisition steps as well.

**Model Comparison** Figure 6a confirms the previous observation (Fig. 5e and 5f) that other methods acquire many images with artifacts and ambiguities, even more than the model with random acquisition (RA). Figure 6b shows that the acquisition of many images with ambiguities and artifacts harms the test performance, as the model EN acquires the most images with ambiguities and artifacts and shows the worst performance. The other AL methods also acquire a substantial amount of images with ambiguities and artifacts and their performance is on par or even below the baseline model with random acquisition RA. Note that in other studies AL methods outperform random acquisition but many are conducted on clean, highly curated datasets. This often does not apply to real-world data (like histopathological images). The proposed FocAL method effectively avoids acquiring uninformative images which leads to the overall best performance.

To compare to the supervised baseline, we try two different settings: one trained with all 2000 images (including the 600 images with artifacts and ambiguities) and another one with only the 1400 images without perturbations. Again, we report the average results over 5 independent runs. The supervised model trained without artifacts and ambiguities performed substantially better (accuracy 0.965; mean f1 score: 0.940) than the model trained with all 2000 images (accuracy 0.939; mean f1 score: 0.901). This confirms our hypothesis that avoiding artifacts and ambiguities is essential for a good performance. We observe that the FocAL method reaches the supervised performance (with all 2000 images) with only 90 labeled training images (4.5% of all images) and even outperforms this supervised performance due to the successful avoidance of perturbed images. Therefore, FocAL not only alleviates the labeling process - **but it can also** save resources usually necessary for data curation. The performance of the supervised model trained on the highly curated dataset with 1400 images is reached by FocAL with 190 labeled images (9.5% of all training data).

### 3.2. Panda

The Panda dataset is a large open dataset for the classification of prostate cancer. We use this dataset for hyperparameter tun-

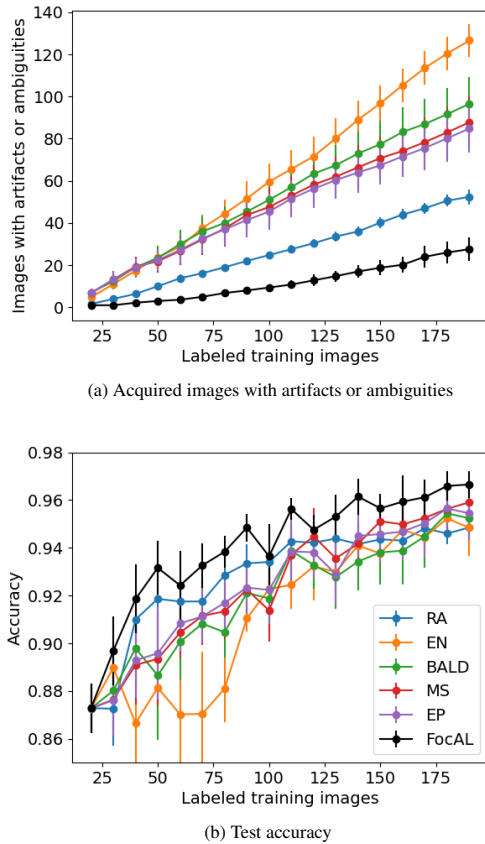


**Fig. 5.** Feature distribution, uncertainty, and acquisition scores for the first acquisition step (best viewed with zoom), similar to Figure 4. Labeled images are dots filled with turquoise (Digit "0"), pink (Digit "1"), or yellow (other digits). Unlabeled images are dots with greyscale color, representing the uncertainty or acquisition score (the higher the brighter). The datapoints with a green edge represent noise-free images (good for training) and datapoints with a red edge represent images with artifacts, blur or ambiguities. For the proposed FocAL method, the epistemic uncertainty (a) measures image informativeness, but it is easily distracted by artifacts and ambiguities. These uninformative images can be captured by a high aleatoric uncertainty (b) or a high OoD score (c). Therefore, in the final FocAL acquisition (d), 9 noise-free images are acquired (and only 1 ambiguous image). The competing methods EN (e) and BALD (f) in comparison acquire almost only images with artifacts or ambiguities in this step which adds less information to the training.

ing, analysis of the different uncertainty measures, and finally for a comparison of the different AL methods.

**Dataset** The Panda dataset (Bulten et al., 2022) consists of 10,616 WSIs and was presented at the MICCAI 2020 conference as a Kaggle challenge. Two institutes participated in labeling the WSIs. The images from Radboud University Medical Center come with detailed label masks of all tissue parts in the Gleason Grading (GG) scheme. The classes are 'Non-cancerous' (NC), 'Gleason 3', 'Gleason 4', and 'Gleason 5' depending on the architectural growth patterns of the tumor. The second institute, the Karolinska Institute, only assigned binary (cancer vs. healthy) labels and we therefore disregard their im-

ages for our experiments. After sorting out corrupted images, a total of 5058 WSIs are left of which we used 1000 WSIs for testing and 30 WSIs for validation that were randomly chosen. The WSIs were divided into 50% overlapping 512x512 patches. To use a multi-scale feature extractor (see Implementation paragraph below), we determine the class of each patch by its center segment (256x256 square), depending on the majority class of the pixels if at least 5% of the pixels are annotated as cancerous. If less than 5% of the pixels are annotated as cancerous, the patch is assigned the non-cancerous label. If a patch contains more than 95% of background (according to the annotation mask), it is disregarded. Therefore, the dataset is already



**Fig. 6. Results of the MNIST Experiments with mean and standard error of five independent runs. In 6a the total acquired noisy images are plotted and in 6b the accuracy. The proposed FocAL algorithm effectively avoids acquiring images with ambiguities and artifacts and shows the strongest performance.**

curated because several artifacts are already excluded in this step. Note that in other experiments, where the dataset curation is more difficult, the advantage of FocAL might be even bigger.

For the empirical validation, we design two different experiments, *small-Panda* for studying the hyperparameter setting due to computational constraints and *big-Panda* with all available images for the final experiments. The big-Panda setup includes all available 4028 WSIs from the Radboud center for training. In this setup, we start with 400 initial images. To ensure having patches of all classes, we use the following procedure: for each class, we randomly select 20 WSIs in which the class is predominant. The information of predominant class is available in the clinical information of prostate cancer images. Then, we sample 5 patches of the predominant class from all selected WSIs. Therefore we have  $4 \text{ classes} \times 5 \text{ patches} \times 20 \text{ WSIs}$ , which leads to 400 patches with 100 patches per


class. In practice, pathologists can reproduce this procedure by assessing the WSIs by predominant classes and annotate 5 patches per WSI for the dominant class. In each acquisition step, we acquire 400 more patches. In the small-Panda experiment, we take a subset of 200 WSIs for training. For the AL initial set, we follow a similar procedure as in big-Panda. We randomly choose 20 WSIs with 5 labeled patches each, resulting in 100 labeled patches in total (equally distributed over the classes with 25 patches per class). We perform 16 acquisition steps of 50 patches each.

The common metric to measure the classification of prostate cancer is Cohen’s quadratic kappa which measures the similarity of the ground truth and the predictions, taking the class order into account (misclassifying Gleason 5 as Gleason 4 has less impact than misclassifying Gleason 5 as NC).

**Implementation** For the feature extraction, we use the EfficientNetB3 model (Tan and Le, 2019) as a CNN backbone in a multi-scale architecture. Remember that the patch classes were assigned based on the center  $256 \times 256$  square of each patch with resolution  $512 \times 512$  (see dataset paragraph above). The center square ( $256 \times 256$ ) is cropped and fed into the CNN. At the same time, the complete  $512 \times 512$  patch is resized to  $256 \times 256$  and fed into a parallel CNN. The two feature vectors (of 1536 dimensions each) are concatenated, followed by a dropout layer and a fully connected layer with 128 units. This approach has the following two advantages. First, the WSI can be segmented with a high level of detail because the classification is performed for relatively small patch centers of  $256 \times 256$ . Second, the context (surrounding tissue) can still be taken into account by the model.

The complete feature extractor has fewer than 22 million trainable parameters (for comparison, a ResNet50 (He et al., 2016) has over 23 million parameters). The BNN consisted, as in the MNIST experiment, of two fully connected layers with 128 units and a final softmax layer with one output unit per class. We train the model with the Adam optimizer (Kingma and Ba, 2015) for 200 epochs for each acquisition step. The learning rate is set to  $1e-4$  for the first 100 epochs, then re-





**Fig. 7. Hyperparameter Tuning of the Focal method.** The two hyperparameters  $\lambda_{al}$  and  $\lambda_{ood}$  are analyzed which weigh the aleatoric uncertainty and OoD detection, respectively. The best performing model with  $\lambda_{al} = 0.5$ ,  $\lambda_{ood} = 1.0$  is used for the final experiments. If the weight  $\lambda_{al}$  is chosen too high, the performance of the model drops because the avoidance of ambiguities is given too much importance.

duced to  $1e-5$  for the other 100 epochs.

**Hyperparameter Tuning** First, we perform experiments regarding the newly introduced hyperparameters. We analyze  $\lambda_{al}$  and  $\lambda_{ood}$  of the Focal acquisition function (eq. 5) that weight the importance of avoiding ambiguities and OoD images, respectively. For this purpose, we use the small-Panda setup (as described in the dataset paragraph above). Figure 7 shows the results for different hyperparameter settings. If the factors are too high, the model performance decreases, as the worst performances are given for the models with  $(\lambda_{al} = 2.0, \lambda_{ood} = 1.0)$  and  $(\lambda_{al} = 2.0, \lambda_{ood} = 2.0)$ . We assume that in this case, the model focuses too much on the avoidance of ambiguities such that the novelty (measured by the epistemic uncertainty) is not given enough importance. Especially a high  $\lambda_{al}$  can harm the performance, as this measure sometimes mistakenly shows high values for informative images at the class boundary. We choose the model with  $\lambda_{al} = 0.5$ ,  $\lambda_{ood} = 1.0$  for the final experiments because it is the overall best-performing model. Note that a comparable performance was obtained for the models with  $(\lambda_{al} = 0.5, \lambda_{ood} = 2.0)$  and  $(\lambda_{al} = 1.0, \lambda_{ood} = 2.0)$ . Overall the performance is robust for all models with  $\lambda_{al} < 2$ .

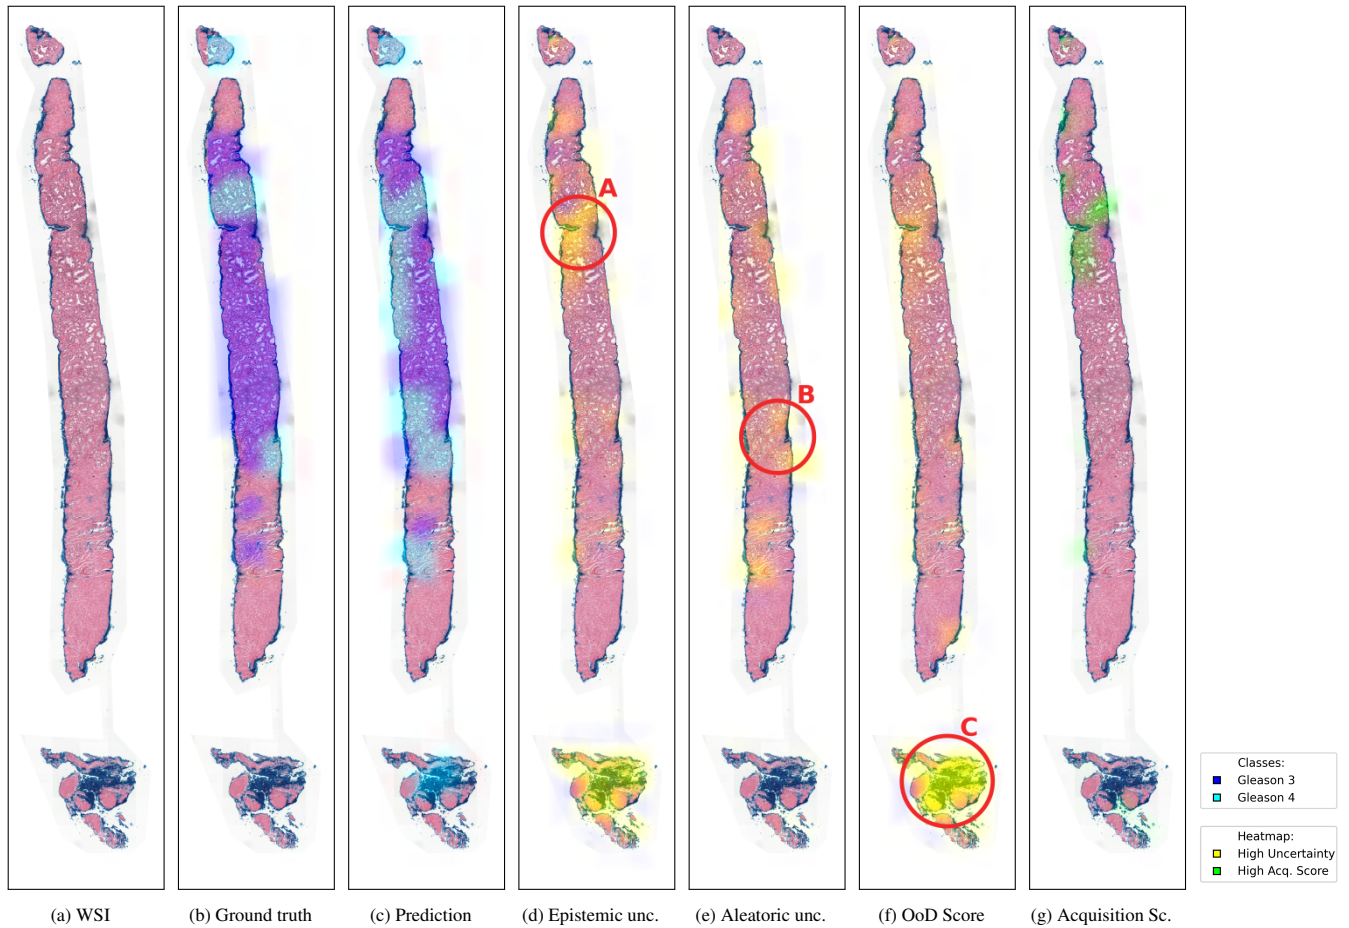
**Uncertainty Estimations** Figure 8 illustrates the uncertainty estimates in the first acquisition step of the Focal method. It shows that each component of the acquisition function works as

expected. The area with a high acquisition score (8g) is based on a high epistemic uncertainty in the circle A in Figure 8d and a low aleatoric uncertainty and OoD score in Figures 8e and 8f, respectively. Indeed, the area contains cancerous tissue and the model shows some misclassifications here (Gleason 4 instead of Gleason 3), as is clear by comparing Figures 8b and 8c. Therefore, labeling these patches can improve the overall model performance. Other parts of the image show a high aleatoric uncertainty, for example, circle B in Figure 8e. This indicates ambiguous patches and therefore, these images are avoided. The acquisition score in this area is low. In circle C in Figure 8f we see a region containing mainly artifacts of dark ink. These artifacts are detected by the OoD detection and therefore also avoided in the final acquisition, although the epistemic uncertainty is high in this region. As the images are produced at an early stage of the active learning process, the model’s predictions (8c) are not yet accurate which is also reflected by high uncertainty values (8d - 8f).

In the diagnostic process, these uncertainties are important to identify unreliable predictions. A high epistemic uncertainty means that the model has to be further trained on this specific tissue type. A high aleatoric uncertainty indicates data ambiguities that might lead to a wrong class prediction. The OoD score shows that the indicated region contains artifacts or content that is substantially different from the learned data distribution. Low uncertainties - of all measures - indicate that the model probably made a correct classification in those areas.

**Acquisition** In Figure 9 we depict five patches with the highest acquisition score of the methods EN, BALD, and Focal. They are taken from the third acquisition step but represent the general tendency that can be observed throughout the active learning process. The proposed Focal method avoids artifacts and ambiguities and acquires representative patches containing cancerous tissue. All five patches with the highest acquisition score contain either Gleason 3 or Gleason 4 in this acquisition step. Overall, in the complete AL process (big-Panda setup), the Focal method acquired the highest number of Gleason 5 patches. It is the most severe grade and at the same time the most un-






**Fig. 8.** Visualization of predictions and uncertainties for a test slide at the first acquisition step of the Panda dataset (after training on the initial 400 patches). It was split into 120 overlapping patches and all heatmaps were produced by cubic interpolation of the patch center predictions. Uncolored parts correspond to non-cancerous tissue (in 8b, 8c) and areas without (or with low) uncertainties (in 8d - 8g). The red circle A in Fig. 8d marks an area with high epistemic uncertainty. As the aleatoric uncertainty and OoD score are low in this area, this results in a high acquisition score (green in Fig. 8g). The area of circle A is informative. The red circle B in Fig. 8e shows an area with high aleatoric uncertainty due to slight blur and ambiguities. Therefore, the final acquisition score is low in this area. The red circle C in Fig. 8f shows an area with artifacts (blue ink) that results in a high OoD score. Although the epistemic uncertainty is high in this area, the acquisition of these non-cancerous and uninformative patches is avoided.


derrepresented one. In total 525 patches of Gleason 5 are acquired on average in the three runs (while EN acquired 403 and all other methods below 400 each). This shows, that the class imbalance was successfully addressed. The EN acquisition assigns a high score for ambiguous patches that contain multiple different classes, like patches 9b and 9e. Both patches include glands of both Gleason 3 and Gleason 4, but for the classification task, only one label per patch is assigned, which is Gleason 4 for patch 9b and Gleason 3 for patch 9e. We argue that these ambiguities can slow down the labeling process because the pathologists take longer for the decision in comparison to the annotation of representative, non-ambiguous patches. The BALD method, which is commonly used and has shown impressive results on clean datasets (Gal et al., 2017), fails to find

informative patches. All five patches with the highest acquisition score contain artifacts and none of them contains cancerous tissue. Similar observations can be made for the MS and EP methods. All three methods (BALD, MS and EP) are highly 'distracted' by artifacts resulting in an acquired dataset in which many patches contain little or no tissue at all.

**Model Comparison** The proposed Focal algorithm shows an overall strong performance as reported in Figure 10 with a final Cohen's quadratic kappa of 0.764 with 4400 image patches corresponding to only 0.69% of all patches of the dataset. After the second acquisition step (with 1200 labeled patches and more), the result is constantly better than RA, MS, EP, and BALD, because the acquisition of artifacts and ambiguities is actively avoided. As the Focal acquisition selects representa-



**Fig. 9.** Acquired image patches. Each row shows the five patches with the highest acquisition score of the methods EN, BALD, and FocAL at the third acquisition step (after training with 800 labeled patches). Below each patch, we report the class and underline cancerous classes to highlight them. While the EN strategy favors ambiguous patches (like patches 9b and 9e) and BALD gets distracted by artifacts that do not contain cancerous tissue, the proposed FocAL method acquires informative patches that represent the cancerous classes well.



**Fig. 10.** Model Comparison on the Panda dataset (setting big-Panda). Several existing methods (BALD, MS, and EP) are below or on par with random acquisition (RA). This is the consequence of the acquisition of artifacts in the active learning process. The EN method acquires fewer artifacts, but more ambiguous images. This seems to be less harmful to the final result but leads to a more difficult labeling process and a less representative dataset. The proposed FocAL method actively avoids artifacts and ambiguities and reaches a satisfying performance with a final Cohen's quadratic kappa of 0.764.

tive patches of the classes, including many images of the most severe grade (Gleason 5) which is highly underrepresented, the created dataset is of high quality. Therefore, our algorithm has successfully addressed the challenges of histopathological labeling.

The methods BALD, MS, and EP, which are acquiring the highest number of artifacts, show a weak performance in the first acquisition steps. Their performance remains significantly below the baseline of random acquisition (RA) until 2400 patches are acquired. Afterwards, their performance is comparable to RA, but not significantly better. This confirms the previous findings of the MNIST experiments: these commonly applied acquisition methods perform well on clean datasets but artifacts have a major impact on their performance.

The EN method is less affected by the acquisition of artifacts but acquires more ambiguous images as already seen in Figure 9 and the MNIST experiment 5. Interestingly, it appears that this doesn't significantly affect the predictive performance on the Panda dataset, since the final Cohen's quadratic kappa values are on par with those of FocAL across various acquisitions. Nevertheless, it's important to highlight some drawbacks of the EN method. In the MNIST experiment, this approach exhibited the poorest performance and gathered the greatest number of patches with ambiguities and artifacts. Consequently, it may not be suitable for other applications. Furthermore, the extensive acquisition of ambiguous images in the Panda dataset can

pose challenges in the posterior labeling process: the annotation by pathologists might take longer and the chances of wrong annotations are higher. Additionally, the resulting dataset may hold less value as it primarily comprises edge cases specific to this model, rather than representative patches from each class. The final Cohen’s quadratic kappa value of EN is 0.759 (FocAL: 0.764). The supervised model with access to all 633.235 patch labels reaches a Cohen’s quadratic kappa of 0.841 for this task.

Finally, notice that the numerical differences against baselines in Figure 10 are not large. In this sense, it is important to consider that active learning in the histopathological image domain remains a tough challenge for which other approaches in literature have shown only minor improvements in comparison to the random and entropy baselines, see for example (Meirelles *et al.*, 2022; Carse and McKenna, 2019).

**Model Limitations** Although FocAL successfully addresses the analyzed problems, there are some limitations worth discussing. We mimic the AL process with an already labeled dataset but it needs to be validated in a human study in the future. In practice, new problems but also advantages of FocAL might appear that are not notable in the experiments with an already labeled dataset. The aleatoric uncertainty is a widely adopted uncertainty measure for data uncertainty and overall captures ambiguities in the data well, but it sometimes shows false-positive values assigned to images close to the class boundary, as seen in the MNIST experiment (Figure 5). Indeed, we showed in the hyperparameter tuning for Panda (Figure 7) that the weight of the aleatoric uncertainty  $\lambda_{al}$  must be carefully chosen. To further improve the model, a more precise uncertainty estimation for ambiguities could be a promising direction. Another limitation is that the model is currently trained from scratch for each acquisition step, as adapted from Gal *et al.* (2017). Incrementally updating the model parameters at each step can reduce the overall training time, especially in the later acquisition steps.

#### 4. Conclusions and future work

In the presented experiments with the MNIST and Panda datasets, the existing probabilistic AL approaches did not work as expected when exposed to ambiguities and artifacts. The widely used BALD algorithm, for example, acquires large amounts of images with artifacts, leading to performance that is often on par with or even below that of a random acquisition strategy. Furthermore, the resulting dataset is not a good representative of the classes of interest. Our proposed FocAL method addresses this issue by using precise uncertainty measures combined with OoD detection to avoid these ambiguities and artifacts while accounting for the class imbalance. In our experiments, we showed that each model component works as expected and that the overall results improve considerably. The acquired images are representative of the classes of interest and form a high-quality dataset. To further improve the model in the future, it would be interesting to combine the proposed FocAL model with other approaches in existing literature. One promising direction is the combination of AL methods with semi-supervised learning, for example, using a consistency loss (Gao *et al.*, 2020) or output discrepancy (Huang *et al.*, 2021). Another addition to our model could be a specific term promoting diversity, as done in the CoreSet approach (Sener and Savarese, 2018). Furthermore, it would be interesting to analyze AL methods for other types of medical images, such as CT scans, dermatology images, or retinal images, with regard to artifacts, ambiguities, and class imbalance. It is likely that state-of-the-art methods such as BALD encounter similar problems, and the proposed FocAL method could provide a possible solution. In addition to these future applications, human studies are needed to validate the method in a real labeling process.

#### References

- Abati, D., Porrello, A., Calderara, S., Cucchiara, R., 2019. Latent space autoregression for novelty detection, in: Conference on Computer Vision and Pattern Recognition - CVPR, IEEE Computer Society. pp. 481–490. doi:10.1109/CVPR.2019.00057.
- Alex Kendall, V.B., Cipolla, R., 2017. Bayesian SegNet: Model uncertainty in deep convolutional encoder-decoder architectures for scene understanding, in: British Machine Vision Conference - BMVC, pp. 57.1–57.12. doi:10.5244/C.31.57.
- Bernhardt, M., Coelho de Castro, D., Tanno, R., Schwaighofer, A., Tezcan, K., Monteiro, M., Bannur, S., Lungren, M., Nori, A., Glocker, B.,

- Alvarez-Valle, J., Oktay, O., 2022. Active label cleaning for improved dataset quality under resource constraints. *Nature Communications* 13. doi:10.1038/s41467-022-28818-3.
- Breunig, M.M., Kriegel, H.P., Ng, R.T., Sander, J., 2000. LOF: Identifying density-based local outliers, in: *ACM SIGMOD International Conference on Management of Data*, pp. 93–104. doi:10.1145/342009.335388.
- Bulten, W., Kartasalo, K., Chen, P.H.C., Ström, P., Pinckaers, H., Nagpal, K., Cai, Y., Steiner, D.F., van Boven, H., Vink, R., Hulsbergen-van de Kaa, C., van der Laak, J., Amin, M.B., Evans, A.J., van der Kwast, T., Allan, R., Humphrey, P.A., Grönberg, H., Samarasinghe, H., Delahunt, B., Tsuzuki, T., Häkkinen, T., Egevad, L., Demkin, M., Dane, S., Tan, F., Valkonen, M., Corrado, G.S., Peng, L., Mermel, C.H., Ruusuvaari, P., Litjens, G., Eklund, M., the PANDA challenge consortium, Brilhante, A., Çakır, A., Farré, X., Geronatsiou, K., Molinić, V., Pereira, G., Roy, P., Saile, G., Salles, P.G.O., Schaafsma, E., Tschui, J., Billock-Lima, J., Pereira, E.M., Zhou, M., He, S., Song, S., Sun, Q., Yoshihara, H., Yamaguchi, T., Ono, K., Shen, T., Ji, J., Roussel, A., Zhou, K., Chai, T., Weng, N., Grechka, D., Shugaev, M.V., Kiminya, R., Kovalev, V., Voynov, D., Malyshev, V., Lapo, E., Campos, M., Ota, N., Yamaoka, S., Fujimoto, Y., Yoshioka, K., Juvonen, J., Tukiainen, M., Karlsson, A., Guo, R., Hsieh, C.L., Zubarev, I., Bukhar, H.S.T., Li, W., Li, J., Speier, W., Arnold, C., Kim, K., Bae, B., Kim, Y.W., Lee, H.S., Park, J., 2022. Artificial intelligence for diagnosis and gleason grading of prostate cancer: the PANDA challenge. *Nature Medicine* 28, 154–163. doi:10.1038/s41591-021-01620-2.
- Campanella, G., Hanna, M.G., Geneslaw, L., Mirafior, A., Werneck Krauss Silva, V., Busam, K.J., Brogi, E., Reuter, V.E., Klimstra, D.S., Fuchs, T.J., 2019. Clinical-grade computational pathology using weakly supervised deep learning on whole slide images. *Nature Medicine* 25, 1301–1309. doi:10.1038/s41591-019-0508-1.
- Carse, J., McKenna, S., 2019. Active learning for patch-based digital pathology using convolutional neural networks to reduce annotation costs, in: *Digital Pathology*. volume 11435, pp. 20–27. doi:10.1007/978-3-030-23937-4\_3.
- Chikontwe, P., Kim, M., Nam, S.J., Go, H., Park, S.H., 2020. Multiple instance learning with center embeddings for histopathology classification, in: *Medical Image Computing and Computer Assisted Intervention – MICCAI*, pp. 519–528. doi:10.1007/978-3-030-59722-1\_50.
- Deng, L., 2012. The mnist database of handwritten digit images for machine learning research. *IEEE Signal Processing Magazine* 29, 141–142.
- Der Kiureghian, A., Ditlevsen, O., 2009. Aleatory or epistemic? does it matter? *Structural safety* 31, 105–112. Publisher: Elsevier.
- Dimitriou, N., Arandjelović, O., Caie, P.D., 2019. Deep learning for whole slide image analysis: An overview. *Frontiers in Medicine* 6, 264. doi:10.3389/fmed.2019.00264.
- Ehteshami Bejnordi, B., Veta, M., Johannes van Diest, P., van Ginneken, B., Karsssemeijer, N., Litjens, G., van der Laak, J.A.W.M., Consortium, a.t.C., 2017. Diagnostic assessment of deep learning algorithms for detection of lymph node metastases in women with breast cancer. *JAMA* 318, 2199–2210. doi:10.1001/jama.2017.14585.
- Esteva, A., Kuprel, B., Novoa, R.A., Ko, J., Swetter, S.M., Blau, H.M., Thrun, S., 2017. Dermatologist-level classification of skin cancer with deep neural networks. *Nature* 542, 115–118. doi:10.1038/nature21056.
- Gal, Y., Islam, R., Ghahramani, Z., 2017. Deep bayesian active learning with image data. *International Conference on Machine Learning - ICML*, 1183–1192.
- Gao, B.B., Xing, C., Xie, C.W., Wu, J., Geng, X., 2017. Deep label distribution learning with label ambiguity. *IEEE Transactions on Image Processing* 26, 2825–2838. doi:10.1109/TIP.2017.2689998.
- Gao, M., Zhang, Z., Yu, G., Arik, S.Ö., Davis, L.S., Pfister, T., 2020. Consistency-based semi-supervised active learning: Towards minimizing labeling cost, in: *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part X* 16, Springer. pp. 510–526.
- Gulshan, V., Peng, L., Coram, M., Stumpe, M.C., Wu, D., Narayanaswamy, A., Venugopalan, S., Widner, K., Madams, T., Cuadros, J., Kim, R., Raman, R., Nelson, P.C., Mega, J.L., Webster, D.R., 2016-12. Development and validation of a deep learning algorithm for detection of diabetic retinopathy in retinal fundus photographs. *Journal of the American Medical Association -AMA* 316, 2402–2410. doi:10.1001/jama.2016.17216.
- He, K., Zhang, X., Ren, S., Sun, J., 2016. Deep residual learning for image recognition, in: *Conference on Computer Vision and Pattern Recognition - CVPR*, pp. 770–778. doi:10.1109/CVPR.2016.90.
- Hekler, A., Utikal, J.S., Enk, A.H., Solass, W., Schmitt, M., Klode, J., Schaden-dorf, D., Sondermann, W., Franklin, C., Bestvater, F., Flaig, M.J., Krah, D., von Kalle, C., Fröhling, S., Brinker, T.J., 2019. Deep learning outperformed 11 pathologists in the classification of histopathological melanoma images. *European Journal of Cancer* 118, 91–96. doi:10.1016/j.ejca.2019.06.012.
- Houlsby, N., Huszar, F., Ghahramani, Z., Lengyel, M., 2011. Bayesian active learning for classification and preference learning. *CoRR* abs/1112.5745. URL: <http://arxiv.org/abs/1112.5745>, arXiv:1112.5745.
- Huang, S., Wang, T., Xiong, H., Huan, J., Dou, D., 2021. Semi-supervised active learning with temporal output discrepancy, in: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 3447–3456.
- Johnson, T., Kwok, I., Ng, R.T., 1998. Fast computation of 2-dimensional depth contours., in: *International Conference on Knowledge Discovery and Data Mining*, pp. 224–228.
- Joshi, A.J., Porikli, F., Papanikolopoulos, N., 2009. Multi-class active learning for image classification, in: *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2372–2379. doi:10.1109/CVPR.2009.5206627.
- Kanwal, N., Pérez-Bueno, F., Schmidt, A., Molina, R., Engan, K., 2022. The devil is in the details: Whole slide image acquisition and processing for artifacts detection, color variation, and data augmentation: A review. *IEEE Access* 10. doi:10.1109/ACCESS.2022.3176091.
- Kendall, A., Gal, Y., 2017. What uncertainties do we need in bayesian deep learning for computer vision?, in: *Neural Information Processing Systems - NeurIPS*.
- Kingma, D.P., Ba, J., 2015. Adam: A method for stochastic optimization .
- Kingma, D.P., Salimans, T., Welling, M., 2015. Variational dropout and the local reparameterization trick, in: *International Conference on Neural Information Processing Systems - NIPS*, pp. 2575–2583.
- Knorr, E.M., Ng, R.T., 1998. Algorithms for mining distance-based outliers in large datasets., in: Gupta, A., Shmueli, O., Widom, J. (Eds.), *International Conference on Very Large Databases - VLDB*, pp. 392–403.
- Knorr, E.M., Ng, R.T., 1999. Finding intensional knowledge of distance-based outliers., in: Atkinson, M.P., Orlowska, M.E., Valduriez, P., Zdonik, S.B., Brodie, M.L. (Eds.), *International Conference on Very Large Databases - VLDB*, pp. 211–222.
- Kwon, Y., Won, J.H., Kim, B.J., Paik, M.C., 2020. Uncertainty quantification using bayesian neural networks in classification: Application to biomedical image segmentation. *Computational Statistics & Data Analysis* 142, 106816. doi:10.1016/j.csda.2019.106816.
- Lee, K., Lee, K., Lee, H., Shin, J., 2018. A simple unified framework for detecting out-of-distribution samples and adversarial attacks, in: *International Conference on Neural Information Processing Systems - NeurIPS*, pp. 7167–7177.
- Lee, S., Amgad, M., Mobadersany, P., McCormick, M., Pollack, B.P., Elfandy, H., Hussein, H., Gutman, D.A., Cooper, L.A., 2021. Interactive classification of whole-slide imaging data for cancer researchers. *Cancer Research* 81, 1171–1177. doi:10.1158/0008-5472.CAN-20-0668.
- Li, B., Li, Y., Elieci, K.W., 2021. Dual-stream multiple instance learning network for whole slide image classification with self-supervised contrastive learning, in: *Conference on Computer Vision and Pattern Recognition - CVPR*, pp. 14313–14323. doi:10.1109/CVPR46437.2021.01409.
- Li, J., Speier, W., Ho, K.C., Sarma, K.V., Gertych, A., Knudsen, B.S., Arnold, C.W., 2018. An EM-based semi-supervised deep learning approach for semantic segmentation of histopathological images from radical prostatectomies. *Comput. Medical Imaging Graph.* 69, 125–133. doi:10.1016/j.compmedimag.2018.08.003.
- Li, X., Guo, Y., 2013. Adaptive active learning for image classification, in: *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 859–866. doi:10.1109/CVPR.2013.116.
- Lu, M.Y., Chen, R.J., Mahmood, F., 2020. Semi-supervised breast cancer histology classification using deep multiple instance learning and contrast predictive coding (conference presentation), in: *Medical Imaging 2020: Digital Pathology*, p. 18. doi:10.1117/12.2549627.
- Marini, N., Otálora, S., Müller, H., Atzori, M., 2021. Semi-supervised training of deep convolutional neural networks with heterogeneous data and few local annotations: An experiment on prostate histopathology image classification. *Medical Image Analysis* 73, 102165. doi:10.1016/j.media.2021.102165.
- Marée, R., Rollus, L., Stévens, B., Hoyoux, R., Louppe, G., Vandaele, R., Begon, J.M., Kainz, P., Geurts, P., Wehenkel, L., 2016. Collaborative analysis of multi-gigapixel imaging data using cytamine. *Bioinformatics* 32, 1395–



1401.

- Meirelles, A.L., Kurc, T., Saltz, J., Teodoro, G., 2022. Effective active learning in digital pathology: A case study in tumor infiltrating lymphocytes. *Computer Methods and Programs in Biomedicine* 220, 106828. doi:10.1016/j.cmpb.2022.106828.
- Mukhoti, J., Kirsch, A., van Amersfoort, J., Torr, P.H.S., Gal, Y., 2021. Deep deterministic uncertainty: A simple baseline. URL: 10.48550/ARXIV.2102.11582.
- Nguyen, A.T., Lu, F., Munoz, G.L., Raff, E., Nicholas, C., Holt, J., 2022. Out of distribution data detection using dropout bayesian neural networks. URL: 10.48550/ARXIV.2202.08985.
- Nguyen, V.L., Destercke, S., Hüllermeier, E., 2019. Epistemic uncertainty sampling, in: Kralj Novak, P., Šmuc, T., Džeroski, S. (Eds.), *International Conference on Discovery Science*, pp. 72–86. doi:10.1007.
- Oliver, A., Odena, A., Raffel, C., Cubuk, E.D., Goodfellow, I.J., 2018. Realistic evaluation of deep semi-supervised learning algorithms, in: *International Conference on Neural Information Processing Systems - NIPS*, Curran Associates Inc., pp. 3239–3250.
- Otálora, S., Marini, N., Müller, H., Atzori, M., 2020. Semi-weakly supervised learning for prostate cancer image classification with teacher-student deep convolutional networks 12446 LNCS, 193–203.
- Ruts, I., Rousseeuw, P.J., 1996. Computing depth contours of bivariate point clouds. *Computational Statistics & Data Analysis* 23, 153–168. doi:10.1016/S0167-9473(96)00027-8.
- Rączkowski, L., Możejko, M., Zambonelli, J., Szczurek, E., 2019. ARA: accurate, reliable and active histopathological image classification framework with bayesian deep learning. *Scientific Reports* 9, 14347. doi:10.1038/s41598-019-50587-1.
- Schmidt, A., Silva-Rodriguez, J., Molina, R., Naranjo, V., 2022. Efficient cancer classification by coupling semi supervised and multiple instance learning. *IEEE Access* 10, 9763–9773. doi:10.1109/ACCESS.2022.3143345.
- Sener, O., Savarese, S., 2018. Active learning for convolutional neural networks: A core-set approach, in: *International Conference on Learning Representations*.
- Shannon, C.E., 1948. A mathematical theory of communication. *The Bell System Technical Journal* 27, 379–423.
- Shvetsov, N., Grønnesby, M., Pedersen, E., Møllersen, K., Busund, L.T.R., Schwienbacher, R., Bongo, L.A., Kilvaer, T.K., 2022. A pragmatic machine learning approach to quantify tumor-infiltrating lymphocytes in whole slide images. *Cancers* 14, 2974. doi:10.3390/cancers14122974.
- Sun, Y., Ming, Y., Zhu, X., Li, Y., 2022. Out-of-distribution detection with deep nearest neighbors, in: *International Conference on Machine Learning - ICML*.
- Tan, M., Le, Q., 2019. EfficientNet: Rethinking model scaling for convolutional neural networks, in: *International Conference on Machine Learning - ICML*, pp. 6105–6114.
- Wu, Y., Schmidt, A., Hernández-Sánchez, E., Molina, R., Katsaggelos, A.K., 2021. Combining attention-based multiple instance learning and gaussian processes for CT hemorrhage detection, in: *Medical Image Computing and Computer Assisted Intervention – MICCAI*. volume 12902, pp. 582–591.
- Xiao, T., Gomez, A., Gal, Y., 2019. Wat heb je gezegd? detecting out-of-distribution translations with variational transformers.
- Zeng, J., Lesnikowski, A., Alvarez, J.M., 2018. The relevance of bayesian layer positioning to model uncertainty in deep bayesian active learning. *Neural Information Processing Systems - NeurIPS*.
- Zhang, Z., Chen, P., McGough, M., Xing, F., Wang, C., Bui, M., Xie, Y., Sapkota, M., Cui, L., Dhillon, J., Ahmad, N., Khalil, F.K., Dickinson, S.I., Shi, X., Liu, F., Su, H., Cai, J., Yang, L., 2019. Pathologist-level interpretable whole-slide cancer diagnosis with deep learning. *Nature Machine Intelligence* 1, 236–245. doi:10.1038/s42256-019-0052-1.
- Zhu, X., Lafferty, J., Ghahramani, Z., 2003. Combining active learning and semi-supervised learning using gaussian fields and harmonic functions, in: *International Conference on Machine Learning - ICML*, pp. 58–65.