



Quantifying and Leveraging Classification Uncertainty for Chest Radiograph Assessment

Florin C. Ghesu¹(✉), Bogdan Georgescu¹, Eli Gibson¹, Sebastian Guendel¹,
Mannudeep K. Kalra^{2,3}, Ramandeep Singh^{2,3}, Subba R. Digumarthy^{2,3},
Sasa Grbic¹, and Dorin Comaniciu¹

¹ Digital Technology and Innovation, Siemens Healthineers, Princeton, NJ, USA
florin.ghesu@siemens-healthineers.com

² Department of Radiology, Massachusetts General Hospital, Boston, MA, USA

³ Harvard Medical School, Boston, MA, USA

Abstract. The interpretation of chest radiographs is an essential task for the detection of thoracic diseases and abnormalities. However, it is a challenging problem with high inter-rater variability and inherent ambiguity due to inconclusive evidence in the data, limited data quality or subjective definitions of disease appearance. Current deep learning solutions for chest radiograph abnormality classification are typically limited to providing probabilistic predictions, relying on the capacity of learning models to adapt to the high degree of label noise and become robust to the enumerated causal factors. In practice, however, this leads to overconfident systems with poor generalization on unseen data. To account for this, we propose an automatic system that learns not only the probabilistic estimate on the presence of an abnormality, but also an explicit uncertainty measure which captures the confidence of the system in the predicted output. We argue that explicitly learning the classification uncertainty as an additional measure to the predicted output, is essential to account for the inherent variability characteristic of this data. Experiments were conducted on two datasets of chest radiographs of over 85,000 patients. Sample rejection based on the predicted uncertainty can significantly improve the ROC-AUC, e.g., by 8% to 0.91 with an expected rejection rate of under 25%. Eliminating training samples using uncertainty-driven bootstrapping, enables a significant increase in robustness and accuracy. In addition, we present a multi-reader study showing that the predictive uncertainty is indicative of reader errors.

1 Introduction

The interpretation of chest radiographs is an essential task in the practice of a radiologist, enabling the early detection of thoracic diseases [9, 12]. To accelerate and improve the assessment of the continuously increasing number of radiographs, several deep learning solutions have been recently proposed for the automatic classification of radiographic findings [4, 12, 13]. Due to large variations in

image quality or subjective definitions of disease appearance, there is a large inter-rate variability which leads to a high degree of label noise [9]. Modeling this variability when designing an automatic system for assessing this type of data is essential; an aspect which was not considered in previous work.

Using principles of information theory and subjective logic [6] based on the Dempster-Shafer framework for modeling of evidence [1], we present a method for training a system that generates both an image-level label and a classification uncertainty measure. We evaluate this system for classification of abnormalities on chest radiographs. The main contributions of this paper include:

1. describing a system for jointly learning classification probabilities and classification uncertainty in a parametric model;
2. proposing uncertainty-driven bootstrapping as a means to filter training samples with highest predictive uncertainty to improve robustness and accuracy;
3. comparing methods for generating stochastic classifications to model classification uncertainty;
4. presenting an application of this system to identify cases with uncertain classification, yielding more accurate classification on the remaining cases;
5. showing that the uncertainty measure can distinguish radiographs with correct and incorrect labels according to a multi-radiologist-consensus study.

2 Background and Motivation

2.1 Machine Learning for the Assessment of Chest Radiographs

The open access to the ChestX-Ray8 dataset [12] of chest radiographs has led to a series of recent publications that propose machine learning based systems for disease classification. With this dataset, Wang et al. [12] also report a first performance baseline of a deep neural network at an average area under receiver operating characteristic curve (ROC-AUC) of 0.75. These results have been further improved by using multi-scale image analysis [13], or by actively focusing the attention of the network on the most relevant sub-regions of the lungs [3]. State-of-the-art results on the official split of the ChestX-Ray8 dataset are reported in [4] (avg. ROC-AUC of 0.81), using a location-aware dense neural network. In light of these contributions, a recent study compares the performance of such an AI system and 9 practicing radiologists [9]. While the study indicates that the system can surpass human performance, it also highlights the high variability among different expert radiologists for the reading of chest radiographs. The reported average specificity of the readers is very high (over 95%), with an average sensitivity of $50\% \pm 8\%$. With such a large inter-rater variability, one may ask: How can real 'ground truth' data be obtained? Does the label noise affect the training? Current solutions do not consider this variability, which leads to models with overconfident predictions and limited generalization.

Principles of Uncertainty Estimation: One way to handle this challenge is to explicitly estimate the classification uncertainty from the data. Recent methods for uncertainty estimation in the context of deep learning rely on Bayesian

estimation theory [8] or ensembles [7] and demonstrate increased robustness to out-of-distribution data. However, these approaches come with significant computational limitations; associated with the high complexity of sampling parameter spaces of deep models for Bayesian risk estimation; or associated with the challenge of managing ensembles of deep models. Sensoy et al. [10] propose an efficient alternative based on the theory of subjective logic [6], training a deep neural network to estimate the sample uncertainty based on observed data.

3 Proposed Method

Following the work of Sensoy et al. [10] based on the Dempster-Shafer theory of evidence [1], we apply principles of subjective logic [6] to derive a binary classification model that can support the joint estimation of per-class probabilities (\hat{p}_+ ; \hat{p}_-) and predictive uncertainty \hat{u} . In this context, a decisional framework is defined through the assignment of so called belief masses from evidence collected from observed data to individual attributes, e.g., membership to a class [1,6]. Let us denote b^+ and b^- the belief values for the positive and negative class, respectively. The uncertainty mass u is defined as: $u = 1 - b^+ - b^-$, where $b^+ = e^+/E$ and $b^- = e^-/E$ with $e^+, e^- \geq 0$ denoting the per-class collected evidence and total evidence $E = e^+ + e^- + 2$. For binary classification, we propose to model the distribution of such evidence values using the beta distribution, defined by two parameters α and β as: $f(x; \alpha, \beta) = \frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)} x^{\alpha-1} (1-x)^{\beta-1}$, where Γ denotes the gamma function and $\alpha, \beta > 1$ with $\alpha = e^+ + 1$ and $\beta = e^- + 1$. In this context, the per-class probabilities can be derived as $p^+ = \alpha/E$ and $p^- = \beta/E$. Figure 1 visualizes the beta distribution for different α, β values.

A training dataset is provided: $\mathcal{D} = \{\mathbf{I}_k, y_k\}_{k=1}^N$, composed of N pairs of images \mathbf{I}_k with class assignment $y_k \in \{0, 1\}$. To estimate the per-class evidence values from the observed data, a deep neural network parametrized by θ can be applied, with: $[e_k^+, e_k^-] = \mathcal{R}(\mathbf{I}_k; \theta)$, where \mathcal{R} denotes the network response function. Using maximum likelihood estimation, one can learn the network parameters $\hat{\theta}$ by optimizing the Bayes risk of the class predictor p_k with a beta prior distribution:

$$\mathcal{L}_k^{data} = \int \|y_k - p_k\|^2 \frac{\Gamma(\alpha + \beta)}{\Gamma(\alpha)\Gamma(\beta)} p_k^{\alpha-1} (1 - p_k)^{\beta-1} dp_k, \tag{1}$$

where $k \in \{1, \dots, N\}$ denotes the index of the training example from dataset \mathcal{D} , p_k the predicted probability on the training sample k , and \mathcal{L}_k^{data} defines the goodness of fit. Using linearity properties of the expectation, Eq. 1 becomes:

$$\mathcal{L}_k^{data} = (y_k - \hat{p}_k^+)^2 + (1 - y_k - \hat{p}_k^-)^2 + \frac{\hat{p}_k^+(1 - \hat{p}_k^+) + \hat{p}_k^-(1 - \hat{p}_k^-)}{E_k + 1}, \tag{2}$$

where \hat{p}_k^+ and \hat{p}_k^- denote the network’s probabilistic prediction. The first two terms measure the goodness of fit, and the last term encodes the variance of the prediction [10].

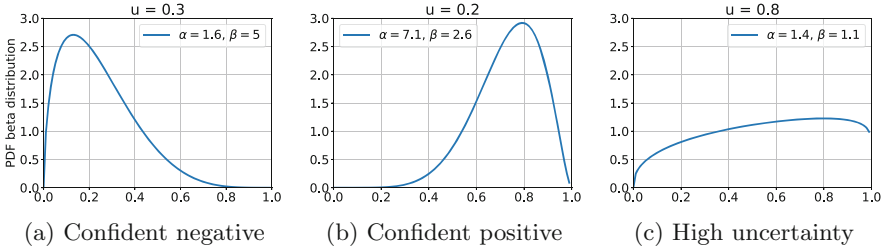


Fig. 1. Probability density function of the beta distribution: example parameters (α, β) modeling confident and uncertain predictions.

To ensure a high uncertainty value for data samples for which the gathered evidence is not conclusive for an accurate classification, an additional regularization term \mathcal{L}_k^{reg} is added to the loss. Using information theory, this term is defined as the relative entropy, i.e., the Kullback-Leibler divergence, between the beta distributed prior term and the beta distribution with total uncertainty $(\alpha, \beta = 1)$. In this way, cost deviations from the total uncertainty state, i.e., $u = 1$, which do not contribute to the data fit are accounted for [10]. With the additional term, the total cost becomes $\mathcal{L} = \sum_{k=1}^N \mathcal{L}_k$ with:

$$\mathcal{L}_k = \mathcal{L}_k^{data} + \lambda \text{KL} \left(f(\hat{p}_k; \tilde{\alpha}_k, \tilde{\beta}_k) \| f(\hat{p}_k; (1, 1)) \right), \quad (3)$$

where $\lambda \geq 0$, $\hat{p}_k = \hat{p}_k^+$, with $(\tilde{\alpha}_k, \tilde{\beta}_k) = (1, \beta_k)$ for $y_k = 0$ and $(\tilde{\alpha}_k, \tilde{\beta}_k) = (\alpha_k, 1)$ for $y_k = 1$. Removing additive constants and using properties of the logarithm function, one can simplify the regularization term to the following:

$$\mathcal{L}_k^{reg} = \log \frac{\Gamma(\tilde{\alpha}_k + \tilde{\beta}_k)}{\Gamma(\tilde{\alpha}_k)\Gamma(\tilde{\beta}_k)} + \sum_{x \in \{\tilde{\alpha}_k, \tilde{\beta}_k\}} (x - 1) \left(\psi(x) - \psi(\tilde{\alpha}_k + \tilde{\beta}_k) \right), \quad (4)$$

where ψ denotes the digamma function and $k \in \{1, \dots, N\}$. Using stochastic gradient descent, the total loss \mathcal{L} is optimized on the training set \mathcal{D} .

Sampling the Data Distribution: An important requirement to ensure training stability and to learn a robust estimation of evidence values is an adequate sampling of the data distribution. We empirically found dropout [11] to be a simple and very effective strategy to address this problem. In practice, dropout emulates an ensemble model combination driven by the random deactivation of neurons. Alternatively, one may use an explicit ensemble of M models $\{\theta_k\}_{k=1}^M$, each trained independently. Following the principles of deep ensembles [7], the per-class evidence can be computed from the ensemble estimates $\{e^{(k)}\}_{k=1}^M$ via averaging. In our work, we found dropout to be as effective as deep ensembles.

Uncertainty-driven Bootstrapping: Given the predictive uncertainty measure \hat{u} , we propose a simple and effective algorithm for filtering the training set with the target of reducing label noise. A fraction of training samples with

highest uncertainty are eliminated and the model is retrained on the remaining data. Instead of sample elimination, robust M-estimators may be applied, using a per-sample weight that is inversely proportional to the predicted uncertainty. The hypothesis is that by focusing the training on ‘confident’ labels, we increase the robustness of the classifier and improve its performance on unseen data.

4 Experiments

Dataset and Setup: The evaluation is based on two datasets, the ChestX-Ray8 [12] and PLCO [2]. Both datasets provide a series of AP/PA chest radiographs with binary labels on the presence of different radiological findings, e.g., granuloma, pleural effusion, or consolidation. The ChestX-Ray8 dataset contains 112,120 images from 30,805 patients, covering 14 findings extracted from radiological reports using natural language processing (NLP) [12]. In contrast, the PLCO dataset was constructed as part of a screening trial, containing 185,421 images from 56,071 patients and covering 12 different abnormalities.

For performance comparison, we selected location-aware dense networks [4] as baseline. This method achieves state-of-the-art results on this problem, with a reported average ROC-AUC of **0.81** (significantly higher than that of competing methods: 0.75 [12] and 0.77 [13]) on the official split of the ChestX-Ray8 dataset and a ROC-AUC of **0.88** on the official split of the PLCO dataset. To evaluate our method, we identified testing subsets with higher confidence labels from multi-radiologist studies. For PLCO, we randomly selected 565 test images and had 2 board-certified expert radiologists read the images – updating the labels to the majority vote of the 3 opinions (incl. the original label). For ChestX-Ray8, a subset of 689 test images was randomly selected and read by 4 board-certified radiologists. The final label was decided following a consensus discussion. For both datasets, the remaining data was split at patient level in 90% training and 10% validation. All images were rescaled to 256×256 using bilinear interpolation.

System Training: We constructed our learning model from the DenseNet-121 architecture [5]. A dropout layer with a dropout rate of 0.5 was inserted after the last convolutional layer. We also investigated the benefits of using deep ensembles to improve the sampling ($M = 5$ models trained on random subsets of 80% of the training data; we refer to this with the keyword [ens]). A fully connected layer with ReLU activation units maps to the two outputs α and β . We used a systematic grid search to find the optimal configuration of training meta-parameters: learning rate (10^{-4}), regularization factor ($\lambda = 1$; decayed to 0.1 and 0.001 after 1/3, respectively 2/3 of the epochs), training epochs (around 12, using an early stop strategy with a patience of 3 epochs) and a batch size of 128. The low number of epochs is explained by the large size of the dataset.

Uncertainty-driven Sample Rejection: Given a model trained for the assessment of an arbitrary finding, one can directly estimate the prediction uncertainty $\hat{u} = 2/(\alpha + \beta) \in [0, 1]$. This is an additional measure to the predicted probability,

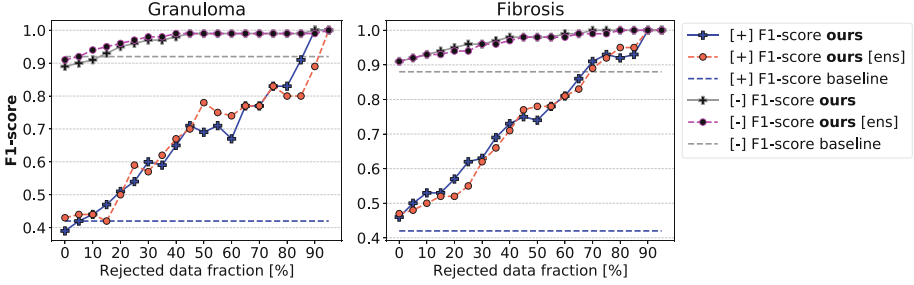


Fig. 2. Evolution of the F1-scores for the positive (+) and negative (-) classes relative to the sample rejection threshold - determined using the estimated uncertainty. We show the performance for granuloma and fibrosis based on the PLCO dataset [2]. The baseline (horizontal dashed lines) is determined using the method from [4] (working point at max. average of per-class F1 scores). Decision threshold for our method is fixed at 0.5.

Table 1. Comparison between the reference method [4] and several versions of our method calibrated at sample rejection rates of 0%, 10%, 25% and 50% (based on the PLCO dataset [2]). Lesion refers to lesions of the bones or soft tissue.

Finding	ROC-AUC				
	Guendel et al. [4]	Ours [0%]	Ours [10%]	Ours [25%]	Ours [50%]
Granuloma	0.83	0.85	0.87	0.90	0.92
Fibrosis	0.87	0.88	0.90	0.92	0.94
Scarring	0.82	0.81	0.84	0.89	0.93
Lesion	0.82	0.83	0.86	0.88	0.90
Cardiac Ab	0.93	0.94	0.95	0.96	0.97
Average	0.85	0.86	0.89	0.91	0.93

with increased values on out-of-distribution cases under the given model. One can use this measure for sample rejection, i.e., set a threshold u_t and steer the system to not output its prediction on all cases with an expected uncertainty larger than u_t . Instead, these cases are labeled with the message “*Do not know for sure; process case manually*”. In practice this leads to a significant increase in accuracy compared to the state-of-the-art on the remaining cases, as reported in Table 1 and Fig. 2. For example, for the identification of granuloma, a rejection rate of 25% leads to an increase of over 20% in the micro-average F1 score. On the same abnormality, a 50% rejection rate leads to a F1 score over 0.99 for the prediction of negative cases. We found no significant difference in average performance when using ensembles (see Fig. 2).

System versus Reader Uncertainty: To provide an insight into the meaning of the uncertainty measure and its correlation with the difficulty of cases, we evaluated our system on the detection of pleural effusion (excess accumulation

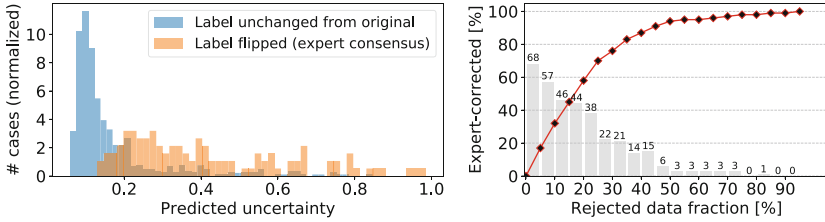
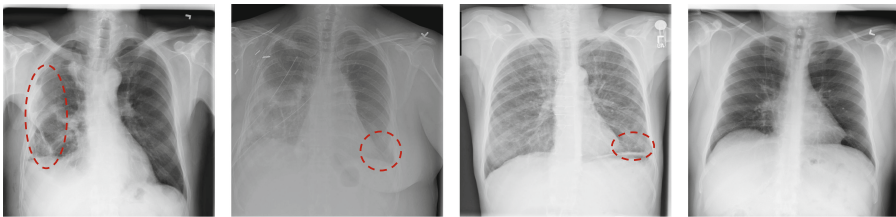


Fig. 3. **Left:** Predictive uncertainty distribution on 689 ChestX-Ray test images; a higher uncertainty is associated with cases of the critical set, which required a label correction according to expert committee. **Right:** Plot showing the capacity of the algorithm to eliminate cases from the critical set via sample rejection. Bars indicate the percentage of critical cases for each batch of 5% rejected cases.

of fluid in the pleural cavity) based on the ChestX-Ray8 dataset. In particular, we analyzed the test set of 689 cases that were relabeled using an expert committee of 4 experts. We defined a so called *critical set*, that contains only cases for which the label was changed after the expert reevaluation. According to the committee, this set contained not only easy examples for which probably the NLP algorithm has failed to properly extract the correct labels from the radiographic report; but also difficult cases for which the evidence of effusion was very subtle. In Fig. 3 (left), we empirically show that the uncertainty estimates of our algorithm correlate with the committee’s decision to change the label. Specifically, for unchanged cases, our algorithm displayed very low uncertainty estimates (average 0.16) at an average AUC of 0.976 (rejection rate of 0%). In contrast, on cases in the critical set, the algorithm showed higher uncertainties distributed between 0.1 and the maximum value of 1 (average 0.41). This indicates the ability of the algorithm to recognize the cases where annotation errors occurred in the first place (through NLP or human reader error). In Fig. 3 (right) we show how cases of the critical set can be effectively filtered out using sample rejection. Qualitative examples are shown in Fig. 4.



(a) $\hat{u}, \hat{p} = 0.90, 0.45$ (b) $\hat{u}, \hat{p} = 0.93, 0.48$ (c) $\hat{u}, \hat{p} = 0.54, 0.65$ (d) $\hat{u}, \hat{p} = 0.11, 0.05$

Fig. 4. ChestX-Ray8 test images assessed for pleural effusion (\hat{u} : est. uncertainty, \hat{p} : output probability; with affected regions circled in red). Figures 4a, 4b and 4c show positive cases of the critical set with high predictive uncertainty – possibly explained by the atypical appearance of accumulated fluid in 4a, and poor quality of image 4b. Figure 4d shows a high confidence case with no pleural effusion. (Color figure online)

Uncertainty-driven Bootstrapping: We also investigated the benefit of using bootstrapping based on the uncertainty measure on the example of plural effusion (ChestX-Ray8). We report performance as [*AUC*; *F1-score* (pos. class); *F1-score* (neg. class)]. After training our method, the baseline performance was measured at [0.89; 0.60; 0.92] on testing. We then eliminated 5%, 10% and 15% of training samples with highest uncertainty, and retrained in each case on the remaining data. The metrics improved to [0.90; 0.68; 0.92]_{5%}, [0.91; 0.67; 0.94]_{10%} and [**0.93; 0.69; 0.94**]_{15%} on the test set. This is a significant increase, demonstrating the potential of this strategy to improve the robustness to label noise.

5 Conclusion

In conclusion, this paper presents an effective method for the joint estimation of class probabilities and classification uncertainty in the context of chest radiograph assessment. Extensive experiments on two large datasets demonstrate a significant accuracy increase if sample rejection is performed based on the estimated uncertainty measure. In addition, we highlight the capacity of the system to distinguish radiographs with correct and incorrect labels according to a multi-radiologist-consensus user study, using the uncertainty measure only.

The authors thank the National Cancer Institute for access to NCI's data collected by the Prostate, Lung, Colorectal and Ovarian (PLCO) Cancer Screening Trial. The statements contained herein are solely those of the authors and do not represent or imply concurrence or endorsement by NCI.

Disclaimer. The concepts and information presented in this paper are based on research results that are not commercially available.

References

1. Dempster, A.P.: A generalization of bayesian inference. *J. Royal Stat. Soc.: Ser. B (Methodol.)* **30**(2), 205–232 (1968)
2. Gohagan, J.K., Prorok, P.C., Hayes, R.B., Kramer, B.S.: The prostate, lung, colorectal and ovarian (PLCO) cancer screening trial of the National Cancer Institute: history, organization, and status. *Control. Clin. Trials* **21**(6), 251–272 (2000)
3. Guan, Q., Huang, Y., Zhong, Z., Zheng, Z., Zheng, L., Yang, Y.: Diagnose like a radiologist: attention guided convolutional neural network for thorax disease classification. [arXiv:1801.09927](https://arxiv.org/abs/1801.09927) (2018)
4. Guendel, S., et al.: Learning to recognize abnormalities in chest X-rays with location-aware dense networks. [arXiv:1803.04565](https://arxiv.org/abs/1803.04565) (2018)
5. Huang, G., Liu, Z., v. d. Maaten, L., Weinberger, K.Q.: Densely connected convolutional networks. In: *CVPR*, pp. 2261–2269 (2017)
6. Jøsang, A.: *Subjective Logic: A Formalism for Reasoning Under Uncertainty*, 1st edn. Springer, Heidelberg (2016). <https://doi.org/10.1007/978-3-319-42337-1>
7. Lakshminarayanan, B., Pritzel, A., Blundell, C.: Simple and scalable predictive uncertainty estimation using deep ensembles. In: *NIPS*, pp. 6402–6413 (2017)
8. Molchanov, D., Ashukha, A., Vetrov, D.: Variational dropout sparsifies deep neural networks. In: *ICML*, pp. 2498–2507 (2017)

9. Rajpurkar, P., et al.: Deep learning for chest radiograph diagnosis: a retrospective comparison of the CheXNeXt algorithm to practicing radiologists. *PLoS Med.* **15**(11), e1002686 (2018)
10. Sensoy, M., Kaplan, L., Kandemir, M.: Evidential deep learning to quantify classification uncertainty. In: *NIPS*, pp. 3179–3189 (2018)
11. Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., Salakhutdinov, R.: Dropout: a simple way to prevent neural networks from overfitting. *JMLR* **15**(1), 1929–1958 (2014)
12. Wang, X., Peng, Y., Lu, L., Lu, Z., Bagheri, M., Summers, R.: ChestX-Ray8: hospital-scale chest X-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases. In: *CVPR*, pp. 3462–3471 (2017)
13. Yao, L., Prosky, J., Poblenz, E., Covington, B., Lyman, K.: Weakly supervised medical diagnosis and localization from multiple resolutions. [arXiv:1803.07703](https://arxiv.org/abs/1803.07703) (2018)