

Received 2 February 2024, accepted 20 February 2024, date of publication 1 March 2024, date of current version 7 March 2024.

Digital Object Identifier 10.1109/ACCESS.2024.3372423

RESEARCH ARTICLE

VI-PANN: Harnessing Transfer Learning and Uncertainty-Aware Variational Inference for Improved Generalization in Audio Pattern Recognition

JOHN FISCHER¹, MARKO ORESCANIN¹, (Member, IEEE), AND ERIC ECKSTRAND²

¹Department of Computer Science, Naval Postgraduate School, Monterey, CA 93943, USA

²Data Science and Analytics Group, Naval Postgraduate School, Monterey, CA 93943, USA

Corresponding author: John Fischer (john.fischer@nps.edu)

This work was supported in part by the Office of Naval Research under Grant N-00014-23-WX-00868.

ABSTRACT Transfer learning (TL) is an increasingly popular approach to training deep learning (DL) models that leverages the knowledge gained by training a foundation model on diverse, large-scale datasets for use on downstream tasks where less domain- or task-specific data is available. The literature is rich with TL techniques and applications; however, the bulk of the research makes use of deterministic DL models which are often uncalibrated and lack the ability to communicate a measure of epistemic (model) uncertainty in prediction. Unlike their deterministic counterparts, Bayesian DL (BDL) models are often well-calibrated, provide access to epistemic uncertainty for a prediction, and are capable of achieving competitive predictive performance. In this study, we propose variational inference pre-trained audio neural networks (VI-PANNs). VI-PANNs are a variational inference variant of the popular ResNet-54 architecture which are pre-trained on AudioSet, a large-scale audio event detection dataset. We evaluate the quality of the resulting uncertainty when transferring knowledge from VI-PANNs to other downstream acoustic classification tasks using the ESC-50, UrbanSound8K, and DCASE2013 datasets. We demonstrate, for the first time, that it is possible to transfer calibrated uncertainty information along with knowledge from upstream tasks to enhance a model's capability to perform downstream tasks.

INDEX TERMS Audio event detection, AudioSet, Bayesian deep learning, transfer learning, uncertainty decomposition, uncertainty quantification.

I. INTRODUCTION

Transfer learning (TL) leverages knowledge gained from large foundation models to enhance performance on downstream tasks. In the audio domain, the feasibility of TL has been demonstrated through the successful application of TL techniques in numerous applications ranging from music genre classification to heart sound classification [1], [2], [3], [4], [5], [6]. While deterministic embeddings prevail, variational embeddings provide a promising Bayesian alternative. By using variational inference (VI) to infer

posterior distributions over latent features, we obtain variational embeddings which capture uncertainty and enable new analyses of transferred representations. However, the use of variational embeddings in TL remains relatively unexplored despite their ability to capture uncertainty. Uncertainty estimation is crucial for assessing model credibility and identifying unreliable predictions [7], [8], [9], [10]. Specifically, the variance of variational embeddings provides epistemic uncertainty estimates that indicate when models lack knowledge [8], [9]. This benefits building reliable artificial intelligence systems across audio domains.

In this study, we conduct a comprehensive analysis of uncertainty-aware TL with variational audio embeddings.

The associate editor coordinating the review of this manuscript and approving it for publication was Prakasam Periasamy¹.

Bayesian versions of the popular ResNet architecture [11] are first trained on the large-scale AudioSet dataset to extract robust acoustic features along with uncertainty estimates. Building off the work of [1], we call our approach variational inference pre-trained audio neural networks (VI-PANNs). VI-PANNs are then transferred to three audio classification benchmark datasets - ESC-50 [12], UrbanSound8K [13], and DCASE2013 [14]. Our goal is to rigorously evaluate the quality and utility of the resulting uncertainty from variational embeddings after transfer to sparse data scenarios.

While extensive work exists on uncertainty quantification and decomposition for classification problems [15], [16], [17], the focus of this work is on the multi-class case, where each sample contains only a single label. Unfortunately, this restriction hinders the application of existing techniques to real-world datasets where multiple labels may be present in each sample [7], [18], [19]. To address this limitation, we describe a new uncertainty decomposition method for use in multi-label classification and demonstrate its efficacy on AudioSet.

This manuscript makes several contributions:

- 1) We modify the popular ResNet-54 architecture used in [1] to create two distinct VI model variants, namely MC dropout [20] and Flipout [21]. This modification enhances the model by providing access to calibrated epistemic (model) uncertainty information, which was lacking in the original model. We pre-train these VI-PANN models on the AudioSet dataset, and share the resulting checkpoints for future research.¹
- 2) We derive a method for decomposing uncertainty information in the multi-label classification scenario.
- 3) We apply the multi-label uncertainty decomposition technique to analyze the uncertainty of VI-PANNs on the AudioSet validation set, providing insights into the model's performance in real-world scenarios.
- 4) We systematically evaluate Bayesian TL techniques with a specific emphasis on the quality of uncertainty. This evaluation is conducted using three publicly available audio datasets (ESC-50, UrbanSound8K, and DCASE2013).
- 5) In the audio pattern recognition domain, we demonstrate, for the first time, the feasibility of transferring calibrated uncertainty information alongside knowledge from upstream tasks. This transfer markedly enhances a model's capability to excel in downstream tasks, showcasing an innovative approach in leveraging uncertainty for improved model performance.

II. RELATED WORK

A. TRANSFER LEARNING IN AUDIO

Recently, with popular deep learning frameworks like PyTorch offering pre-trained initializations for modern model architectures, TL has become an integral part of modern model development workflows. Building upon this, research

efforts have leveraged the large-scale AudioSet dataset to pre-train deep neural networks for enhanced performance on downstream audio tasks [1], [5]. A common TL approach is to directly extract features from a pre-trained model fixed after the initial training. This method transfers general acoustic knowledge to new tasks without updating the model parameters. However, fine-tuning the pre-trained model by allowing parameter updates during training on the downstream data can further improve results by adapting to the task [1]. In this manuscript, we refer to these techniques as “fixed-feature” and “fine-tuned,” respectively. Leveraging large pre-trained models via either technique provides significant performance gains across various audio applications [1], [5].

B. BAYESIAN DEEP LEARNING

The inability of modern deterministic deep learning models to communicate a measure of epistemic (model) uncertainty in prediction has led to an increased interest in Bayesian deep learning (BDL), specifically in remote sensing [7], [8], [9], medical [22], [23], and safety-critical applications [24]. Although there are a number of different approaches to BDL, we focus our experiments on VI. Due to the increased speed and the ability to scale with data and models, VI is often favored over techniques like Markov Chain Monte Carlo (MCMC) [25]. In modern probabilistic machine learning libraries like BayesianTorch [26], the most common VI implementations are Flipout [21] and the Local Reparameterization Trick [27]. Due to the fact that both of these approaches represent each model weight using a Gaussian Distribution (i.e., each weight is defined using two model parameters, a mean and a variance), they effectively double the number of model parameters. In 2016, Gal et al. [20] showed that it was possible to perform VI by training a model with dropout layers preceding every weight layer and activating those dropout layers during inference. This approach, called MC dropout, does not double the number of model parameters. For this reason, along with the minimal changes required to common deep learning model architectures and training procedures, MC dropout is often favored over other VI approaches. In this work, we focus on the Flipout and MC dropout implementations of VI.

C. UNCERTAINTY QUANTIFICATION AND DECOMPOSITION IN MULTI-CLASS CLASSIFICATION

One of the primary motivations behind using BDL models is to gain access to high-quality uncertainty for predictions. The existing research literature is rich with techniques for quantifying and decomposing uncertainty. In [28], Kendall and Gal provide insight into two types of uncertainty that can be modeled. Aleatoric, or irreducible, uncertainty is the uncertainty inherent in the data. Epistemic, or reducible, uncertainty is the uncertainty about the prediction due to uncertainty about the model. In addition to providing a detailed description of these uncertainties, the authors describe a method for measuring predictive (total) uncertainty, based on output

¹<https://github.com/marko-orescanin-nps/VI-PANN>




FIGURE 1. Uncertainty calibration plots for foundation model training on AudioSet. Comparison plot of test set accuracy vs. percentage of evaluation data retained based on entropy (left), epistemic uncertainty (center), and aleatoric uncertainty (right). Shading represents a 95% confidence interval (CI).




FIGURE 2. Uncertainty box plots depicting results of Monte Carlo (MC) Dropout model (top row) and Flipout model (bottom row) trained on AudioSet. The plots compare predictive entropy (left), epistemic uncertainty (middle), and aleatoric uncertainty (right) as the models are evaluated on both the AudioSet test set and the ShipsEar dataset. Both the median (orange line) and mean (dashed green line) are presented.

variance, and decomposing the total uncertainty into its aleatoric and epistemic components using the laws of total variance. Unfortunately, the method used by Kendall and Gal requires the use of extra parameters to model the mean and variance of the model output. Kwon et al. [22] expands upon the work in [28] by proposing a method for calculating these component uncertainties without the use of additional model parameters.

Another line of research is based on the use of entropy as a measure of predictive uncertainty [16], [17]. We detail the approach of Chai [17], as we use this method for multi-class classification, and it is the basis of our multi-label classification decomposition method. In BDL multi-class classification problems, we approximate the predictive probability $p(y = c | x)$ using MC integration with M samples [29]. The average probability per class \bar{p}_c is

calculated using

$$\bar{p}_c = \frac{1}{M} \sum_{m=1}^M \hat{p}_{c_m}, \tag{1}$$

where $\hat{p}_{c_m} = p(y = c | x, \theta^m)$ and θ^m is sampled from an approximation of $p(\theta | \mathcal{D})$. Defining C as the set of all possible classes, we can then compute the entropy of a prediction with

$$\mathbb{H}[y | x, \mathcal{D}] = - \sum_{c \in C} \bar{p}_c \log \bar{p}_c. \tag{2}$$

Depeweg et al. [16] and Chai [17] use the entropy from Eq. (2) as a measure of total uncertainty and decompose that uncertainty using the following:

$$\mathbb{H}[y | x, \mathcal{D}] = \underbrace{\mathbb{I}[y, \theta | x, \mathcal{D}]}_{\text{Epistemic}} + \underbrace{\mathbb{E}_{\theta \sim p(\theta | \mathcal{D})}[\mathbb{H}[y | x, \theta]]}_{\text{Aleatoric}}, \tag{3}$$

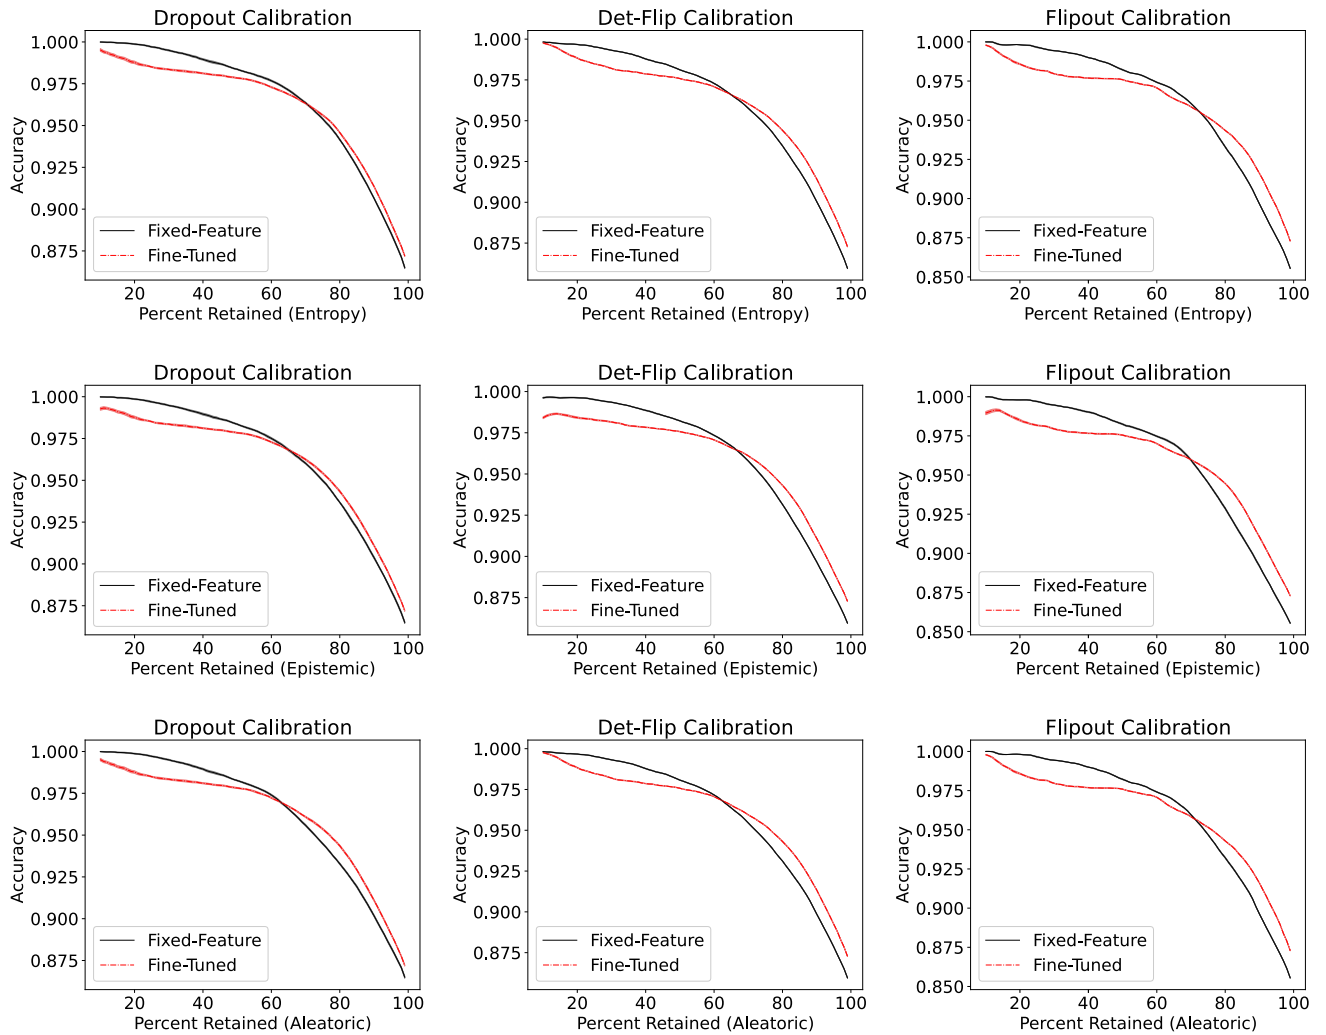


FIGURE 3. Uncertainty calibration plots comparing fixed-feature and fine-tuning TL techniques on UrbanSound8K. Comparison plots of test set accuracy vs. percentage of evaluation data retained based on Entropy (top), Epistemic Uncertainty (middle) and Aleatoric Uncertainty (bottom). Drop VI-PANN is on the left, Det-Flip VI-PANN in the center, and Flip VI-PANN on the right. Shading represents a 95% CI.

where \mathbb{E} is expected value and \mathbb{I} is mutual information. Similar to the calculation of predictive entropy in Eq. (2), we approximate the aleatoric uncertainty component using MC integration to arrive at the following estimator:

$$\mathbb{E}_{\theta \sim p(\theta|\mathcal{D})}[\mathbb{H}[y | x, \theta]] \approx -\frac{1}{M} \sum_{m=1}^M \sum_{c \in C} \hat{p}_{c_m} \log \hat{p}_{c_m}. \quad (4)$$

Finally, the epistemic uncertainty component is calculated by finding the difference between Eq. (2) and (4).

III. METHODOLOGY

A. ARCHITECTURE

As a starting point, we adopt the ResNet-54 architecture described in [1] and make use of the source code provided by the authors. In order to evaluate VI-PANN, we implement MC dropout [20] and Flipout [21] variants of the pre-trained audio neural network (PANN) architecture in [1].

For the MC dropout variant, the architecture of [1] is left unmodified during training. However, during inference, we explicitly keep dropout layers active.

In order to implement the Flipout model, we utilize the Bayesian-Torch [26] software package. Using Bayesian-Torch, we convert deterministic layers to Bayesian layers. More specifically, linear layers are converted to LinearFlipout layers and Conv2d layers are converted to Conv2dFlipout. These weight layers are initialized using the MOPED methodology described in [30]. In our case, initialization is done by calling the Bayesian-Torch `dnn_to_bnn()` function with our pre-trained deterministic model and the default `moped_delta` parameter of 0.5. We then modify the cross entropy loss function from [1] to a loss function based on the following form of the negative Evidence Lower Bound (ELBO):

$$\mathcal{L}_q = \text{KL}[q_\phi(\theta) || p(\theta)] - \mathbb{E}_q [\log p(\mathcal{D}|\theta)], \quad (5)$$

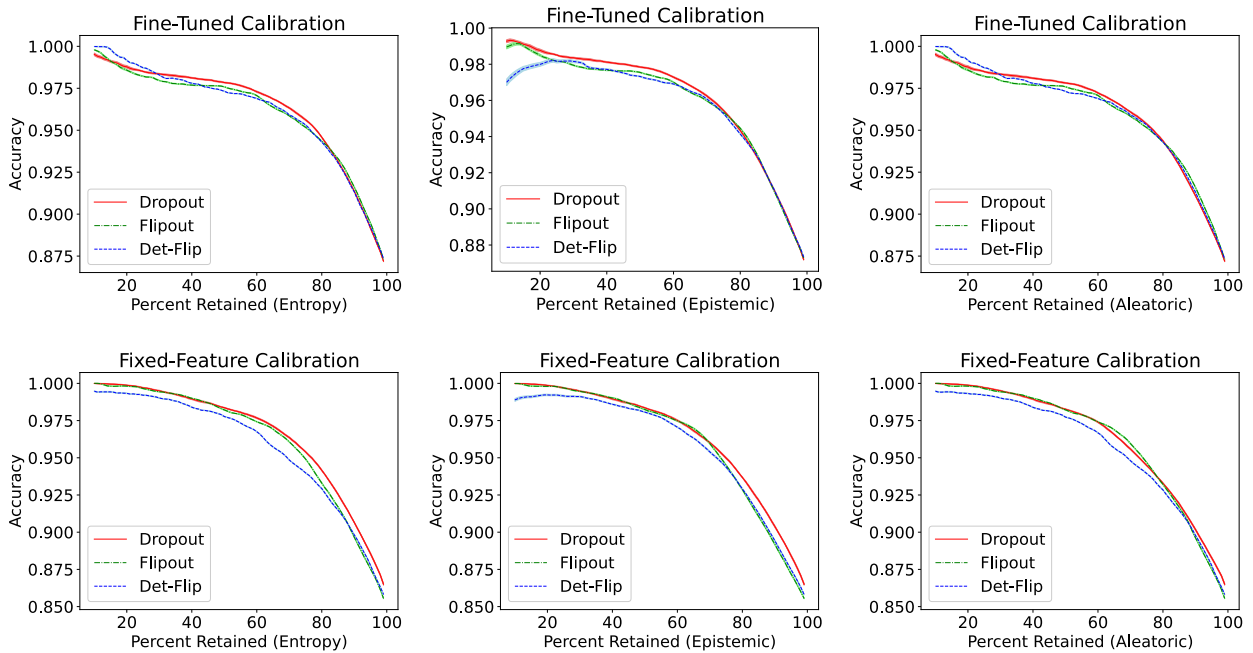


FIGURE 4. Uncertainty calibration plots comparing Drop, Flip, and Det-Flip VI-PANN variants on UrbanSound8k. Comparison plots of test set accuracy vs. percentage of evaluation data retained based on Entropy (left), Epistemic Uncertainty (center) and Aleatoric Uncertainty (right). Plots corresponding to fine-tuned models are on the top, fixed-feature model plots are on the bottom. Shading represents a 95% CI.

where KL corresponds to the Kullback-Leibler (KL) divergence, and \mathbb{E}_q represents the expected value under the probability distribution $q_\phi(\theta)$. A detailed discussion and derivation of this objective can be found in [8].

B. UNCERTAINTY QUANTIFICATION AND DECOMPOSITION IN MULTI-LABEL CLASSIFICATION

In this work, we train and evaluate BDL models on both multi-class and multi-label classification tasks. In the multi-class case (ESC-50, UrbanSound8K, and DCASE2013), we can directly apply the techniques described in Depeweg et al. [16] and Chai [17]. In the multi-label case (AudioSet), however, we must modify the multi-class uncertainty decomposition technique to account for the fact that each class is an independent binary classification problem. Following the methodology in [17], we start by calculating the predictive entropy (i.e., total uncertainty). To calculate the predictive entropy, we first calculate the entropy for each class

$$\mathbb{H}[y | x, \mathcal{D}] = -\bar{p}_c \log \bar{p}_c - (1 - \bar{p}_c) \log(1 - \bar{p}_c), \quad (6)$$

where \bar{p}_c is defined in Eq. (1). Next, to capture the total entropy of the prediction, we sum over all classes $\sum_{c \in \mathcal{C}} \mathbb{H}[y | x, \mathcal{D}]$. Borrowing the definition of \hat{p}_{c_m} from Section II-C, and modifying Eq. (2.21) from [17] for the multi-label case, we are left with the following estimator for aleatoric uncertainty

$$\begin{aligned} & \mathbb{E}_{\theta \sim p(\theta | \mathcal{D})} [\mathbb{H}[y | x, \theta]] \\ & \approx - \sum_{c \in \mathcal{C}} \frac{1}{M} \sum_{m=1}^M \hat{p}_{c_m} \log \hat{p}_{c_m} + (1 - \hat{p}_{c_m}) \log(1 - \hat{p}_{c_m}). \quad (7) \end{aligned}$$

Finally, in order to compute epistemic uncertainty, we calculate the difference between Eq. (6) and (7),

$$\mathbb{I}[y, \theta | x, \mathcal{D}] \approx \mathbb{H}[y | x, \mathcal{D}] - \mathbb{E}_{\theta \sim p(\theta | \mathcal{D})} [\mathbb{H}[y | x, \theta]]. \quad (8)$$

C. MODEL EVALUATION

In order to align with [1] and [19], we present our pre-training results using mean average precision (mAP), area under the curve (AUC), and d-prime. Similar to [1], we calculate each metric using macro-averaging (i.e., we calculate each class individually and average across classes).

For assessing model calibration, we draw on the insights from Filos et al. [29], who demonstrated that a well-calibrated model’s performance improves when high-uncertainty predictions are discarded. Furthermore, Ortiz et al. [9], [31] demonstrated on large scale multispectral satellite datasets (multi-year data) for both classification and regression applications that proper calibration and uncertainty quantification are critical for operational use of neural network models in geoscience applications. Consequently, we employ mAP and accuracy versus data retained curves to evaluate model calibration based on predictive entropy, aleatoric uncertainty, and epistemic uncertainty. Plot shading represents a 95% CI calculated over 20 replications.

To illustrate the practicality of calibrated model uncertainty, we assess each of our VI-PANNs on the ShipsEar dataset [32]. ShipsEar is a multi-class classification dataset comprising 90 underwater sound recordings of 11 different types of ships. This dataset was chosen because each sample is out-of-distribution (OOD), and the recordings, captured underwater with hydrophones, differ from the microphones used in the TL datasets in this study. Consequently, we can

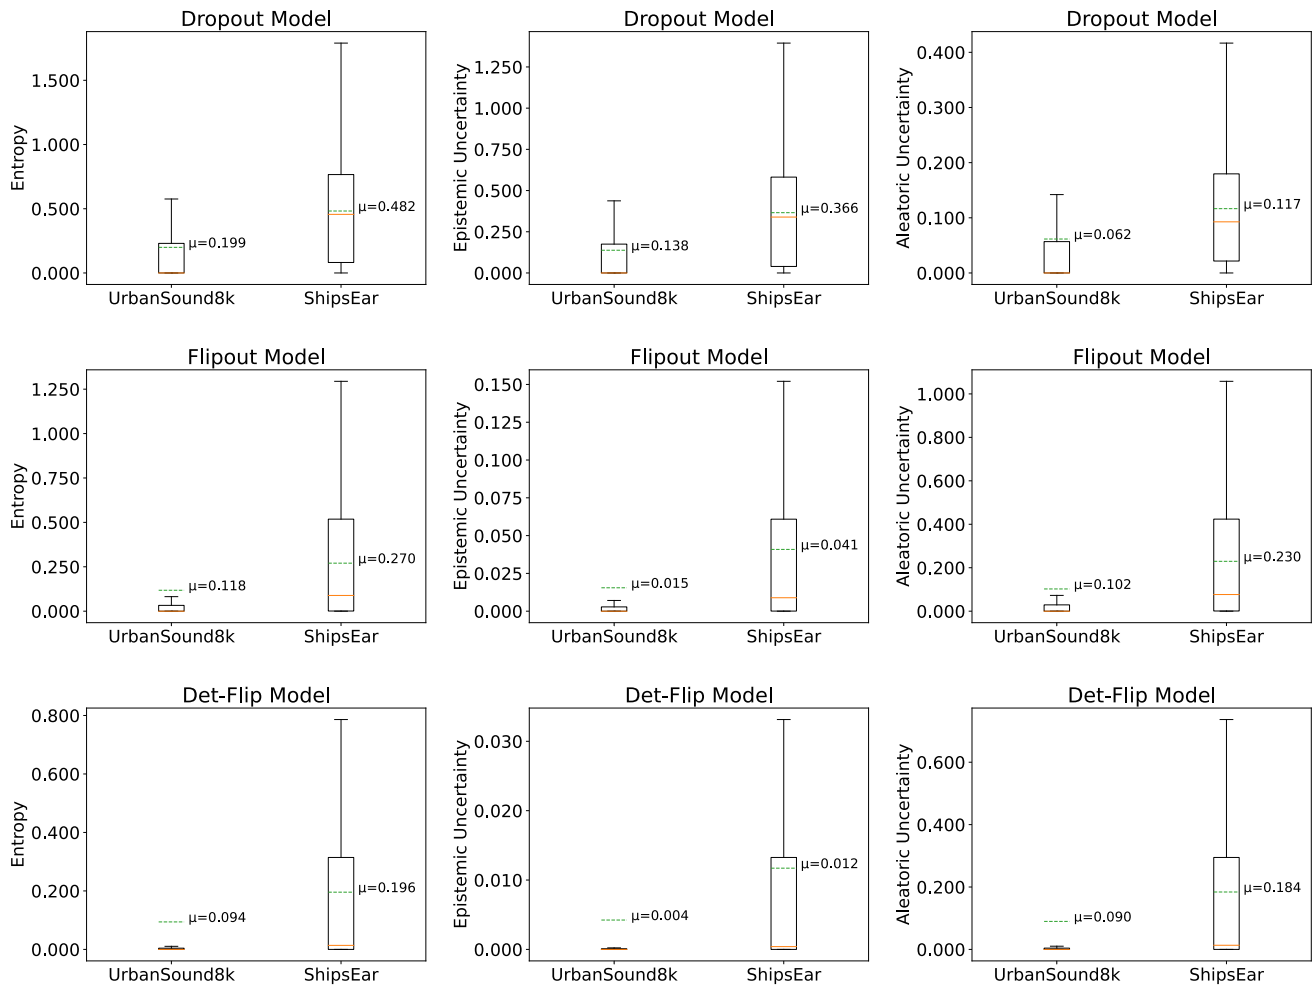


FIGURE 5. Uncertainty box plots depicting results of MC Dropout (top row), Flipout (middle row), and Det-Flip (bottom row) fine tuned on UrbanSound8k. The plots compare predictive entropy (left), epistemic uncertainty (middle), and aleatoric uncertainty (right) as the models are evaluated on both UrbanSound8k and the ShipsEar dataset. Both the median (orange line) and mean (dashed green line) are presented.

analyze the change in model uncertainty (total, aleatoric, and epistemic) when each model is evaluated on data types and distributions it hasn't been trained on.

Due to the fact that our TL datasets require cross-fold validation, we present all results averaged across folds.

D. BAYESIAN DEEP LEARNING MODEL PRE-TRAINING

In order to pre-train our models on AudioSet, we adopt the approach and hyperparameters from [1]. Specifically, to standardize and control for the acoustic pre-processing hyperparameters, enabling a direct and meaningful comparison of model performance between our VI-PANNs and the deterministic models detailed in [1], AudioSet acoustic segments are resampled to 32kHz and converted to log-mel spectrograms using a Hamming window of 1024, a hop size of 320, and 64 mel filter banks. Additionally, following the approach in [1], we remove frequencies above 14kHz and below 50Hz from the samples. For additional details on

acoustic pre-processing hyperparameter selection, we refer the interested reader to [1]. We use a batch size of 32, and an Adam optimizer with a learning rate of 0.001. For the MC dropout variants, we use a dropout rate of 0.2 for convolutional layers, and 0.5 for linear layers.

We apply this training setup to our YouTube-curated repository of approximately 1.7M 10-second, unbalanced audio clips. Similar to [1], we employ mixup [33] augmentation with $\alpha = 1.0$; however, we make no effort to balance the training dataset. As the goal of our investigation is not to match state-of-the-art performance on the AudioSet tagging task but rather construct large-parameter probabilistic versions of AudioSet pre-trained networks to investigate the benefits they confer to uncertainty analysis in the acoustic domain, we train our deterministic PANN and MC dropout VI-PANN for approximately 3M steps. In order to train our Flipout VI-PANN, we initialize the network priors and posteriors using MOPED [30] with the learned weights from our deterministic PANN. We then train the Flipout VI-PANN

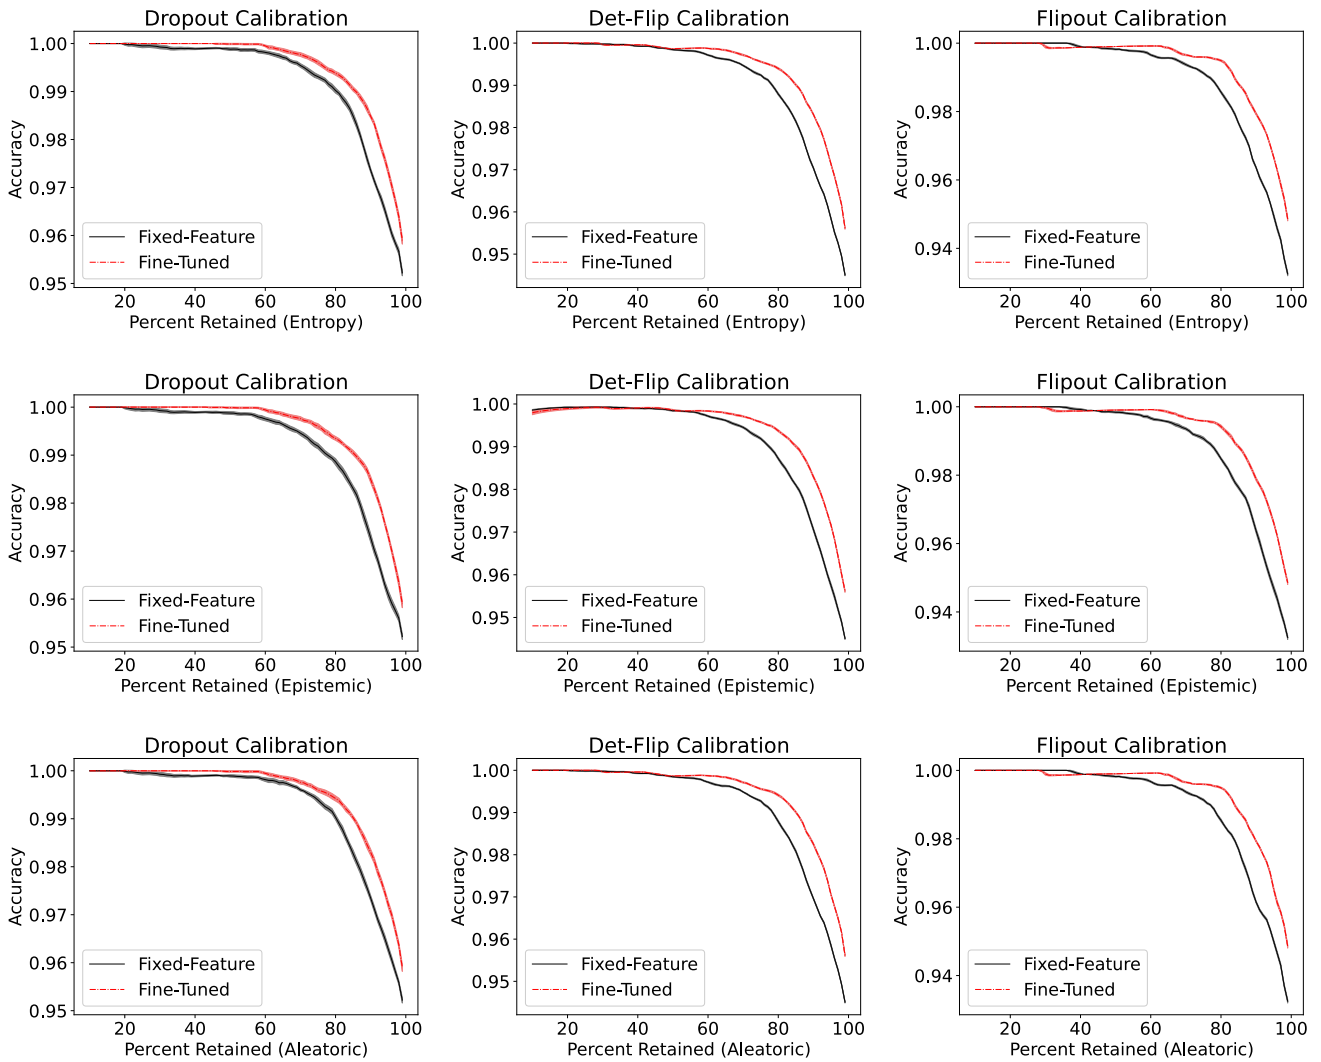


FIGURE 6. Uncertainty calibration plots comparing fixed-feature and fine-tuning TL techniques on ESC-50. Comparison plots of test set accuracy vs. percentage of evaluation data retained based on Entropy (top), Epistemic Uncertainty (middle) and Aleatoric Uncertainty (bottom). Drop VI-PANN is on the left, Det-Flip VI-PANN in the center, and Flip VI-PANN on the right. Shading represents a 95% CI.

for an additional 2M steps. The deterministic PANN and both VI-PANNs are evaluated using the AudioSet balanced evaluation split.

E. BAYESIAN TRANSFER LEARNING

In our TL experiments, we explore three distinct TL strategies.

1) Flip Strategy:

- Initialize a Flipout model with parameters from our Flipout VI-PANN.
- Replace the classification head with a new Flipout head, using the Bayesian-Torch LinearFlipout layer defaults.
- Freeze the backbone, train the classification head for 200 epochs with a learning rate of 0.001 (referred to as “fixed-feature”), then unfreeze the

backbone, reduce the learning rate by a factor of 10, and train for an additional 200 epochs (referred to as “fine-tuned”).

2) Det-Flip Strategy:

- Initialize a deterministic model with our deterministic PANN using MOPED ($moped_delta = 0.5$), as described in III-A.
- Replace the deterministic head with a Flipout head, following the Flip strategy workflow.

3) Drop Strategy:

- Initialize an MC dropout model with our MC dropout VI-PANN.
- Replace the head with an MC dropout head (dropout rate: 0.5), and follow the Flip strategy workflow.

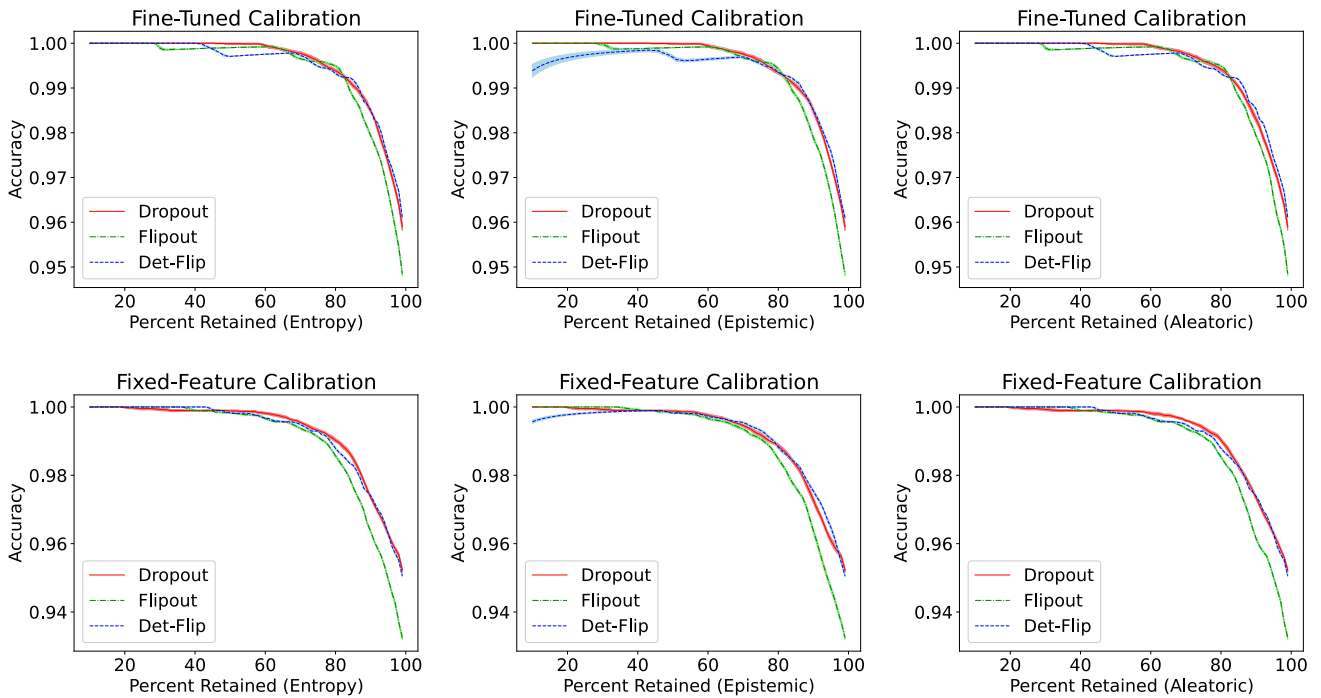


FIGURE 7. Uncertainty calibration plots comparing Drop, Flip, and Det-Flip VI-PANN variants on ESC-50. Comparison plots of test set accuracy vs. percentage of evaluation data retained based on Entropy (left), Epistemic Uncertainty (center) and Aleatoric Uncertainty (right). Plots corresponding to fine-tuned models are on the top, fixed-feature model plots are on the bottom. Shading represents a 95% CI.

TABLE 1. Model parameter counts and multiply-accumulate operations (MACs) of the three VI model variants used in the transfer learning experiments.

Dataset	Model	Params	MACs (per-sample)
ESC-50	Drop	103,340,786	13,536,497,600
	Det-Flip	107,639,588	13,545,306,248
	Flip	206,620,004	26,653,313,032
UrbanSound8k	Drop	103,258,826	13,536,333,760
	Det-Flip	107,475,668	13,544,814,568
	Flip	206,456,084	26,652,821,352
DCASE2013	Drop	103,258,826	81,862,378,176
	Det-Flip	107,475,668	81,870,858,984
	Flip	206,456,084	160,695,176,040

For comparison, we include results from a deterministic baseline (**Det**):

- Initialize a deterministic model with our deterministic PANN.
- Replace the head with a deterministic head and follow the Flip strategy training workflow.

These diverse strategies allow us to assess the impact of different transfer learning approaches on model performance.

For comparison, Table 1 contains a summary of TL model variant, number of parameters, and the number of multiply-accumulate operations (MACs). The presentation is segmented by dataset, as the input feature shape and the number of classes have distinct impacts on the MACs and number of model parameters, respectively.

IV. RESULTS AND DISCUSSION

A. DATASETS

To comprehensively evaluate the uncertainty-aware transfer learning approach, we select a single foundation

TABLE 2. Characteristics of the datasets used to evaluate VI-PANN embeddings in transfer learning.

Dataset	Num Clips	Num Classes	Duration	Folds
ESC-50	2,000	50	5 sec	5
UrbanSound8K	8,732	10	≤ 4 sec	10
DCASE2013	100	10	30 sec	5

dataset, AudioSet, and three diverse audio classification datasets - ESC-50, UrbanSound8K, and DCASE2013. These datasets offer various sound recognition tasks to assess the generalization of variational embeddings. A summary of the TL dataset details is presented in Table 2.

AudioSet [19]: A large-scale audio event recognition dataset consisting of 2.1M 10-second audio samples. Each of these samples were extracted from videos on YouTube and hand-annotated. Of the approximately 2.1M videos listed in the original AudioSet paper [19], we were only able to obtain approximately 1.7M (many videos from the published dataset are no longer available via the links from [19]). AudioSet has an ontology of 527 classes and is a heavily imbalanced, multi-label dataset (i.e., one or more labels can be present in a given sample). It is essential to note that the label quality varies significantly across classes, with some having noisy labels and others exhibiting high-quality annotations. Within the AudioSet data, there are three splits: a balanced evaluation split, a balanced training split, and an unbalanced training split.

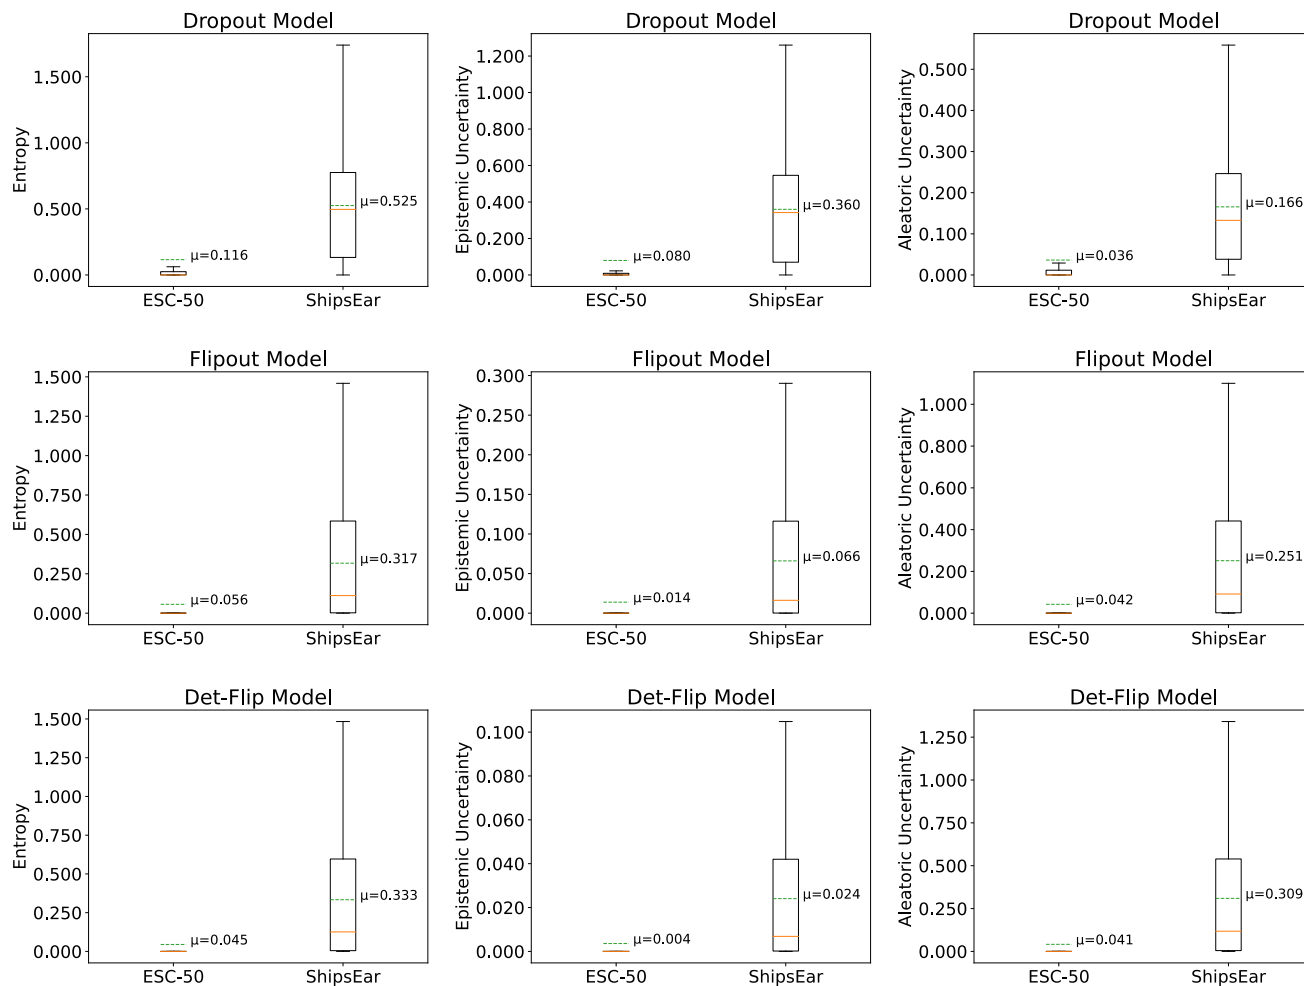


FIGURE 8. Uncertainty box plots depicting results of MC Dropout (top row), Flipout (middle row), and Det-Flip (bottom row) fine tuned on ESC-50. The plots compare predictive entropy (left), epistemic uncertainty (middle), and aleatoric uncertainty (right) as the models are evaluated on both the ESC-50 and the ShipsEar dataset. Both the median (orange line) and mean (dashed green line) are presented.

ESC-50 [12]: A multi-class classification dataset which consists of 2000 five-second recordings organized into 50 classes. Split into 5-folds for cross-validation, it covers a variety of environmental sound events like gunshots, dogs barking, and applause. The ESC-50 dataset is suitable for evaluating fine-grained event recognition capabilities.

UrbanSound8K [13]: A multi-class classification dataset containing 8732 urban sound excerpts up to 4 seconds, categorized into 10 classes. Split into 10-folds for cross-validation, this dataset is commonly used to evaluate a model’s ability to identify ambient urban noises such as air conditioner, car horn, and children playing.

DCASE2013 [14]: A multi-class dataset with 10 classes representing various acoustic scenes and events. It consists of 100 audio samples, each 30 seconds in duration, and is split into 5-folds for cross-validation. DCASE2013 is commonly used to evaluate acoustic scene/event classification in medium duration recordings.

Together, these datasets enable a rigorous evaluation of our approach on diverse audio classification tasks with labelled data far more scarce than the large foundation dataset. Furthermore, the variety of sounds and context shifts across datasets allows us to evaluate the ability of VI-PANNs to transfer and generalize their learned variational embeddings. Analyzing uncertainties on these datasets will reveal how embedding distributions capture model credibility across different acoustic environments and events.

B. FOUNDATION MODEL (AUDIOSET)

The results of our AudioSet pre-training are summarized in Table 3. Each of our models exhibits comparable performance, as measured by mAP, AUC, and d-prime, to the ResNet-54 PANN presented in [1]. Alongside performance metrics, we provide predictive entropy (total uncertainty), epistemic uncertainty, and aleatoric uncertainty calibration plots for our VI-PANNs in Fig. 1.

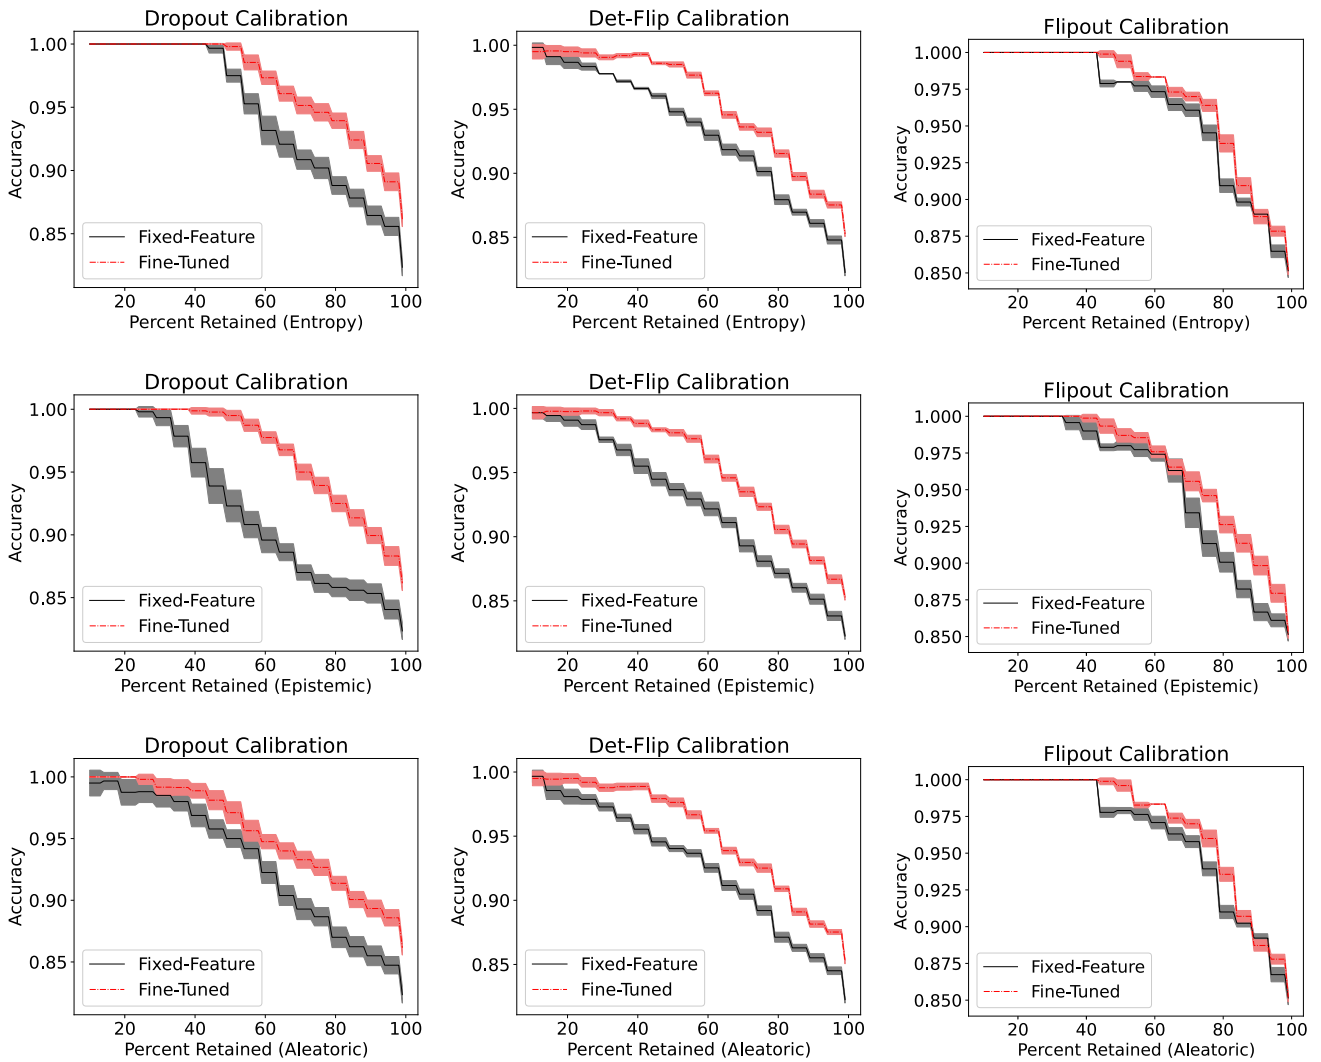


FIGURE 9. Uncertainty calibration plots comparing fixed-feature and fine-tuning TL techniques on DCASE2013. Comparison plots of test set accuracy vs. percentage of evaluation data retained based on Entropy (top), Epistemic Uncertainty (middle) and Aleatoric Uncertainty (bottom). Drop VI-PANN is on the left, Det-Flip VI-PANN in the center, and Flip VI-PANN on the right. Shading represents a 95% CI.

Our Flipout VI-PANN demonstrates well-calibrated uncertainty across all three measures. In contrast, the MC dropout VI-PANN exhibits poor calibration. A well-calibrated model typically shows improved performance as high-uncertainty predictions are discarded [29]. The observed poor calibration is likely attributed to the complexity and significant class imbalance within the AudioSet dataset. Additionally, the lack of explicit learning of the dropout parameter during the training process may contribute to this issue [34].

In Fig. 2, we present box plots that compare the model uncertainty on both the AudioSet test set and the ShipsEar dataset. Although results from both the Flipout and MC Dropout models are included, our primary focus is on the Flipout model due to its demonstrated calibration across all three types of uncertainty. Upon analyzing the model’s response to samples from the ShipsEar dataset, we observe a subtle increase in both average entropy and

aleatoric uncertainty when compared to the AudioSet test set. Conversely, the average epistemic uncertainty remains consistent across both datasets. Notably, the plots reveal a tighter distribution of uncertainty on the ShipsEar dataset in contrast to the AudioSet dataset. This phenomenon is likely attributed to the diversity of the input data, extreme class imbalance, and unsatisfactory label quality observed in many underrepresented classes within the AudioSet dataset.

C. TRANSFER LEARNING

The transfer learning (TL) results for ESC-50, UrbanSound8K, and DCASE2013 are detailed in Table 4. For context, the results of training models from scratch (i.e., without TL) are also provided in Table 5. Each model variant (Det, Det-Flip, Drop, Flip) exhibits comparable performance across the three datasets. Notably, fine-tuned models demonstrate a substantial increase in accuracy compared to

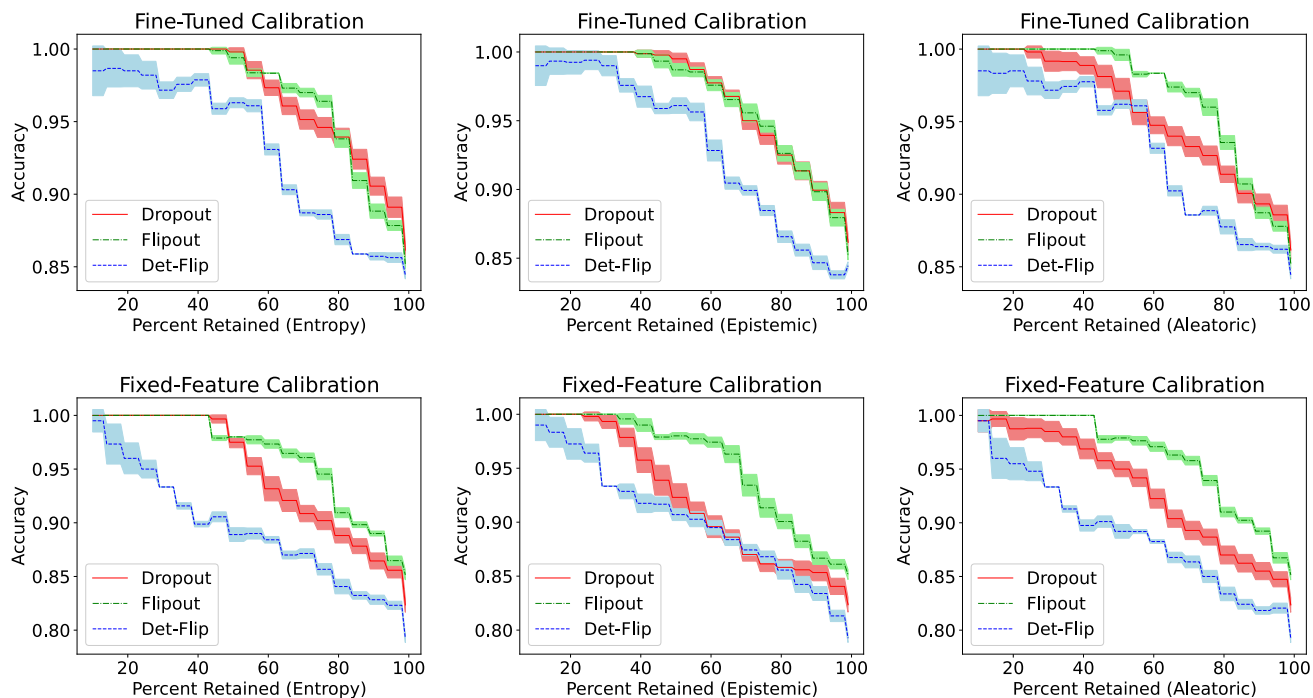


FIGURE 10. Uncertainty calibration plots comparing Drop, Flip, and Det-Flip VI-PANN variants on DCASE2013. Comparison plots of test set accuracy vs. percentage of evaluation data retained based on Entropy (left), Epistemic Uncertainty (center) and Aleatoric Uncertainty (right). Plots corresponding to fine-tuned models are on the top, fixed-feature model plots are on the bottom. Shading represents a 95% CI.

TABLE 3. Model mean average precision (mAP), area under the receiver operating characteristic curve (AUC), and d-prime after pre-training on the AudioSet dataset.

Model	mAP	AUC	d-prime
Det	0.418	0.973	2.732
Drop	0.426	0.975	2.762
Flip	0.415	0.972	2.708
ResNet-54 PANN [1]	0.429	0.971	2.675

fixed-feature (fixed) models and both TL techniques provide significant performance increases over the models trained from scratch on the TL datasets. As expected, when trained from scratch on the relatively small TL datasets, the high-capacity ResNet-54 models perform relatively poorly and suffer from overfitting. The Flip variant which contains 2x the learnable parameters, when compared to the other variants, performs particularly poorly. Although the Det PANN slightly outperforms others in accuracy on ESC-50 and DCASE2013, it lacks the capability to provide access to epistemic uncertainty in predictions.

For reference, we present the results alongside the state-of-the-art (SOTA) results for each dataset. It is essential to clarify that the primary aim of this study was not to achieve SOTA performance on these datasets. Instead, our goal was to demonstrate performance comparable to existing methods while also offering calibrated epistemic uncertainty information. Nonetheless, our VI-PANNs demonstrate comparable

performance (within 2 to 3% accuracy) to these SOTA approaches. It’s worth noting that many of the model architectures employed to achieve SOTA performance on these datasets are Transformer-based. In contrast to our approach, these Transformer-based architectures are deterministic and do not provide access to calibrated epistemic uncertainty information.

Furthermore, we provide calibration plots for UrbanSound8K (Figs. 3 and 4), ESC-50 (Figs. 6 and 7), and DCASE2013 (Figs. 9 and 10). These plots reveal that, following TL, when fixing the features of the base model and after fine-tuning, all three variants of our VI-PANNs result in well-calibrated models. The stairstep pattern evident in the DCASE2013 plots is attributed to the comparatively small size of the DCASE2013 dataset.

In each of the UrbanSound8k calibration plots, the curve of the TL model learned from fixed-features starts at a lower accuracy and crosses over that of the TL model fine-tuned, achieving a higher accuracy at the same percentage of data retained. These results suggest that fine-tuning may have an adverse effect on model calibration if care is not taken in the fine-tuning process. Intuitively, one might extend the hyperparameter evaluation of fine-tuning and model selection to include calibration curves.

In Figs. 5, 8, and 11, we depict box plots that compare model uncertainty for each of the TL datasets (UrbanSound8k, ESC-50, and DCASE2013) and the ShipsEar dataset. Since the fine-tuned models demonstrated superior

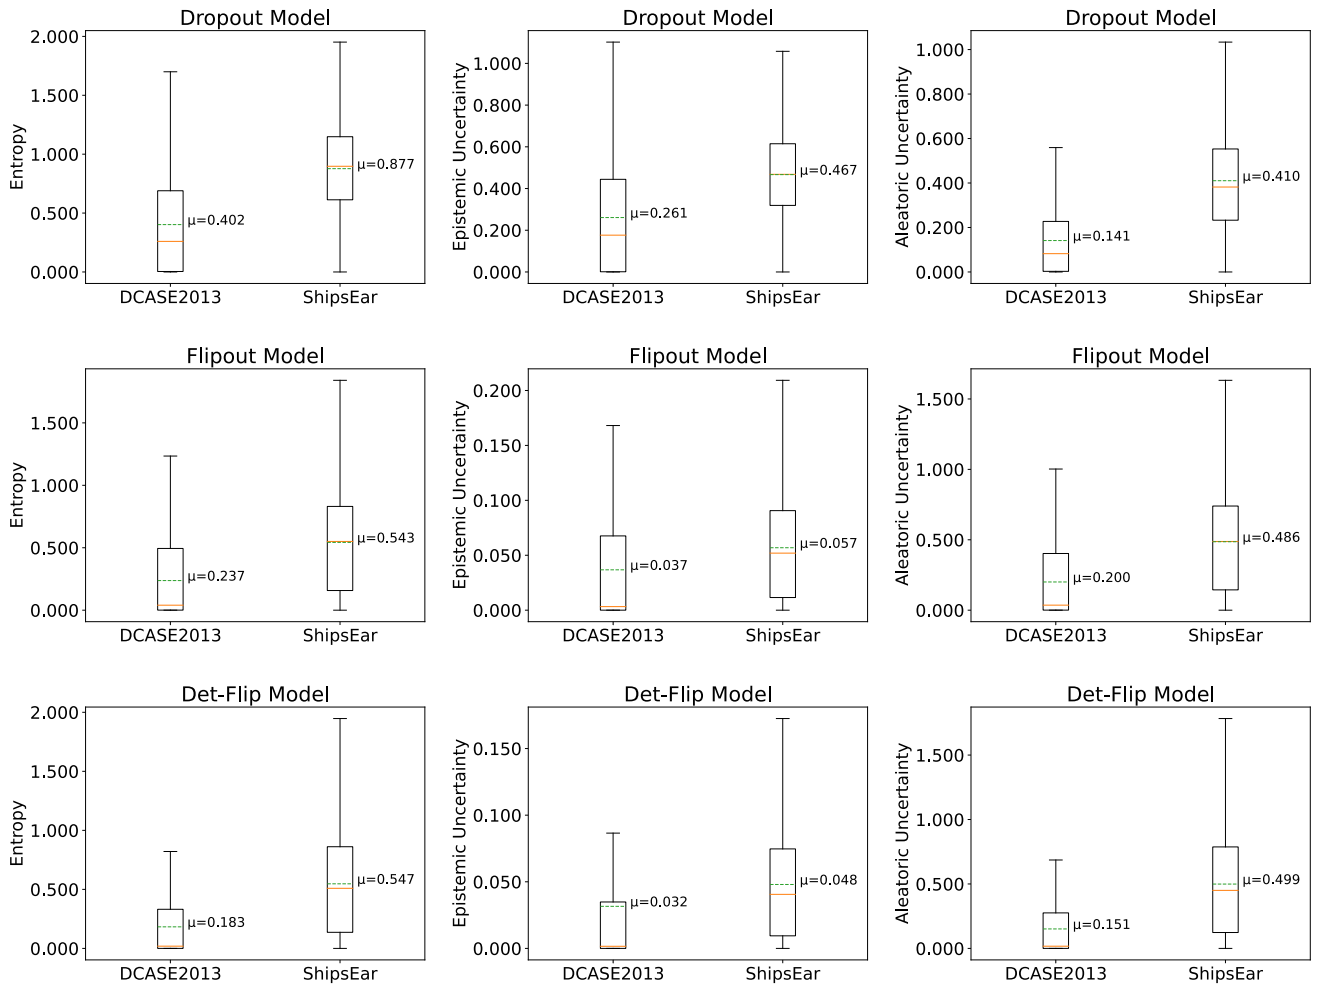


FIGURE 11. Uncertainty box plots depicting results of MC Dropout (top row), Flipout (middle row), and Det-Flip (bottom row) fine tuned on DCASE2013. The plots compare predictive entropy (left), epistemic uncertainty (middle), and aleatoric uncertainty (right) as the models are evaluated on both DCASE2013 and the ShipsEar dataset. Both the median (orange line) and mean (dashed green line) are presented.

TABLE 4. Transfer learning experiment model accuracies on ESC-50, UrbanSound8k, and DCASE2013.

Model	ESC-50		UrbanSound8K		DCASE2013	
	Fixed	Fine-Tuned	Fixed	Fine-Tuned	Fixed	Fine-Tuned
Flip	0.934	0.950	0.857	0.873	0.860	0.870
Det-Flip	0.951	0.962	0.858	0.874	0.790	0.850
Drop	0.952	0.958	0.864	0.873	0.830	0.860
Det	0.951	0.964	0.858	0.870	0.870	0.880
SOTA	0.984 [35]		0.908 [36]		0.916 [5]	

TABLE 5. Baseline model accuracies after training on ESC-50, UrbanSound8k, and DCASE2013 without transfer learning.

Model	ESC-50	UrbanSound8k	DCASE2013
Flip	0.422	0.794	0.400
Det-Flip	0.834	0.815	0.610
Drop	0.861	0.827	0.680
Det	0.858	0.827	0.730

performance compared to the fixed-feature models, we showcase the results of the fine-tuned Dropout, Flipout, and Det-Flip models. In contrast to the AudioSet results, when assessed on the ShipsEar dataset, all three model variants exhibit a notable increase in average entropy, epistemic uncertainty, and aleatoric uncertainty compared to the TL dataset. Moreover, the plots illustrate a considerably broader distribution of uncertainty compared to the TL datasets. This outcome is anticipated given that these models perform exceptionally well with low uncertainty on the TL datasets.

V. CONCLUSION

In this study, we introduce VI-PANNs as a Bayesian alternative to widely adopted deterministic audio embedding methods. Trained on AudioSet, our VI-PANNs exhibit calibrated models through the use of the Flipout approach, underscoring the significance of variational audio embeddings. By adapting uncertainty decomposition techniques for multi-label classification, we enable a nuanced analysis of uncertainty estimates not only on AudioSet but also on other multi-label datasets. Notably, our work represents the first

adaptation of the uncertainty decomposition from [16] and [17] for application in multi-label problems.

Our transfer learning (TL) experiments on well-established datasets demonstrate comparable or improved performance compared to previous state-of-the-art methods, leveraging a similar model architecture. Importantly, our Det-Flip VI-PANN, constructed with a deterministic PANN and a Flipout classification head, achieves high performance at a relatively low cost compared to pre-training a full Flipout model. This establishes robust baselines for uncertainty-aware audio transfer learning in scenarios with limited labeled data, offering valuable insights for practitioners.

Crucially, the presented methodology for TL with variational audio embeddings is universal and applicable to diverse audio tasks. The insights gained emphasize the intrinsic value of Bayesian neural networks in facilitating reliable and transparent transfer learning within the audio domain.

REFERENCES

- Q. Kong, Y. Cao, T. Iqbal, Y. Wang, W. Wang, and M. D. Plumbley, "PANNs: Large-scale pretrained audio neural networks for audio pattern recognition," *IEEE/ACM Trans. Audio, Speech, Language Process.*, vol. 28, pp. 2880–2894, 2020.
- E. Tsalera, A. Papadakis, and M. Samarakou, "Comparison of pre-trained CNNs for audio classification using transfer learning," *J. Sensor Actuator Netw.*, vol. 10, no. 4, p. 72, Dec. 2021.
- K. Choi, G. Fazekas, M. Sandler, and K. Cho, "Transfer learning for music classification and regression tasks," in *Proc. 18th Int. Soc. Music Inf. Retr. (ISMIR) Conf.*, Suzhou, China, 2017, pp. 141–149.
- T. Koike, K. Qian, Q. Kong, M. D. Plumbley, B. W. Schuller, and Y. Yamamoto, "Audio for audio is better? An investigation on transfer learning models for heart sound classification," in *Proc. 42nd Annu. Int. Conf. IEEE Eng. Med. Biol. Soc. (EMBC)*, Jul. 2020, pp. 74–77.
- P. Lopez-Meyer, J. A. del Hoyo Ontiveros, H. Lu, and G. Stemmer, "Efficient end-to-end audio embeddings generation for audio classification on target applications," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process. (ICASSP)*, Jun. 2021, pp. 601–605.
- A. Diment and T. Virtanen, "Transfer learning of weakly labelled audio," in *Proc. IEEE Workshop Appl. Signal Process. to Audio Acoust. (WASPAA)*, Oct. 2017, pp. 6–10.
- B. Beckler, A. Pfau, M. Orescanin, S. Atchley, N. Villemez, J. E. Joseph, C. W. Miller, and T. Margolina, "Multilabel classification of heterogeneous underwater soundscapes with Bayesian deep learning," *IEEE J. Ocean. Eng.*, vol. 47, no. 4, pp. 1143–1154, Oct. 2022.
- J. Fischer, M. Orescanin, P. Leary, and K. B. Smith, "Active Bayesian deep learning with vector sensor for passive sonar sensing of the ocean," *IEEE J. Ocean. Eng.*, vol. 48, no. 3, pp. 837–852, Jul. 2023.
- P. Ortiz, M. Orescanin, V. Petkovic, S. W. Powell, and B. Marsh, "Decomposing satellite-based classification uncertainties in large Earth science datasets," *IEEE Trans. Geosci. Remote Sens.*, vol. 60, 2022, Art. no. 4106211.
- M. Orescanin, V. Petkovic, S. W. Powell, B. R. Marsh, and S. C. Heslin, "Bayesian deep learning for passive microwave precipitation type detection," *IEEE Geosci. Remote Sens. Lett.*, vol. 19, pp. 1–5, 2022.
- K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 770–778.
- K. J. Piczak, "ESC: Dataset for environmental sound classification," in *Proc. 23rd ACM Int. Conf. Multimedia*. New York, NY, USA: ACM Press, Oct. 2015, pp. 1015–1018.
- J. Salamon, C. Jacoby, and J. P. Bello, "A dataset and taxonomy for urban sound research," in *Proc. 22nd ACM Int. Conf. Multimedia*. New York, NY, USA: Association for Computing Machinery, Nov. 2014, pp. 1041–1044.
- D. Stowell, D. Giannoulis, E. Benetos, M. Lagrange, and M. D. Plumbley, "Detection and classification of acoustic scenes and events," *IEEE Trans. Multimedia*, vol. 17, no. 10, pp. 1733–1746, Oct. 2015.
- Y. Gal, "Uncertainty in deep learning," Ph.D. thesis, Dept. Eng., Univ. Cambridge, Cambridge, U.K., 2016.
- S. Depeweg, J. M. Hernández-Lobato, F. Doshi-Velez, and S. Udluft, "Decomposition of uncertainty in Bayesian deep learning for efficient and risk-sensitive learning," in *Proc. 35th Int. Conf. Mach. Learn.*, vol. 80, J. Dy and A. Krause, Eds., Jul. 2018, pp. 1184–1193.
- L. R. Chai, "Uncertainty estimation Bayesian neural networks and links to interpretability," M.S. thesis, Dept. Eng., Univ. Cambridge, Cambridge, U.K., 2018.
- A. M. Pfau, "Multi-label classification of underwater soundscapes using deep convolutional neural networks," M.S. thesis, Dept. Comput. Sci., Nav. Postgraduate School, Monterey, CA, USA, 2020.
- J. F. Gemmeke, D. P. W. Ellis, D. Freedman, A. Jansen, W. Lawrence, R. C. Moore, M. Plakal, and M. Ritter, "Audio set: An ontology and human-labeled dataset for audio events," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process. (ICASSP)*, Mar. 2017, pp. 776–780.
- Y. Gal and Z. Ghahramani, "Dropout as a Bayesian approximation: Representing model uncertainty in deep learning," in *Proc. Int. Conf. Mach. Learn.*, 2016, pp. 1050–1059.
- Y. Wen, P. Vicol, J. Ba, D. Tran, and R. B. Grosse, "Flipout: Efficient pseudo-independent weight perturbations on mini-batches," in *Proc. 6th Int. Conf. Learn. Represent. (ICLR)*, Vancouver, BC, Canada, Apr. 2018. [Online]. Available: <https://openreview.net/forum?id=rJNpifWab>
- Y. Kwon, J.-H. Won, B. J. Kim, and M. C. Paik, "Uncertainty quantification using Bayesian neural networks in classification: Application to biomedical image segmentation," *Comput. Statist. Data Anal.*, vol. 142, Feb. 2020, Art. no. 106816.
- P. McClure, C. Y. Zheng, J. Kaczmarzyk, J. Rogers-Lee, S. S. Ghosh, D. Nielson, P. A. Bandettini, and F. Pereira, "Distributed weight consolidation: A brain segmentation case study," in *Proc. Adv. Neural Inf. Process. Syst., Annu. Conf. Neural Inf. Process. Syst.*, Montréal, QC, Canada, S. Bengio, H. M. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, Eds., Dec. 2018, pp. 4097–4107.
- F. Arnez, H. Espinoza, A. Radermacher, and F. Terrier, "Towards dependable autonomous systems based on Bayesian deep learning components," in *Proc. 18th Eur. Dependable Comput. Conf. (EDCC)*, Sep. 2022, pp. 65–72.
- M. D. Hoffman, D. M. Blei, C. Wang, and J. Paisley, "Stochastic variational inference," *J. Mach. Learn. Res.*, vol. 14, pp. 1303–1347, May 2013.
- R. Krishnan, P. Esposito, and M. Subedar. (Jan. 2022). *Bayesian-Torch: Bayesian Neural Network Layers for Uncertainty Estimation*. [Online]. Available: <https://github.com/IntellLabs/bayesian-torch>
- D. P. Kingma, T. Salimans, and M. Welling, "Variational dropout and the local reparameterization trick," in *Advances in Neural Information Processing Systems*, vol. 28, C. Cortes, N. Lawrence, D. Lee, M. Sugiyama, and R. Garnett, Eds. Red Hook, NY, USA: Curran Associates, 2015.
- A. Kendall and Y. Gal, "What uncertainties do we need in Bayesian deep learning for computer vision?" in *Advances in Neural Information Processing Systems*, vol. 30, I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, Eds. Red Hook, NY, USA: Curran Associates, 2017.
- A. Filos, S. Farquhar, A. N. Gomez, T. G. J. Rudner, Z. Kenton, L. Smith, M. Alizadeh, A. de Kroon, and Y. Gal, "A systematic comparison of Bayesian deep learning robustness in diabetic retinopathy tasks," 2019, *arXiv:1912.10481*.
- R. Krishnan, M. Subedar, and O. Tickoo, "Specifying weight priors in Bayesian deep neural networks with empirical Bayes," in *Proc. AAAI Conf. Artif. Intell.*, vol. 34, 2020, pp. 4477–4484.
- P. Ortiz, E. Casas, M. Orescanin, S. W. Powell, V. Petkovic, and M. Hall, "Uncertainty calibration of passive microwave brightness temperatures predicted by Bayesian deep learning models," *Artif. Intell. for Earth Syst.*, vol. 2, no. 4, Oct. 2023, Art. no. e220056.
- D. Santos-Domínguez, S. Torres-Guijarro, A. Cardenal-López, and A. Pena-Gimenez, "ShipsEar: An underwater vessel noise database," *Appl. Acoust.*, vol. 113, pp. 64–69, Dec. 2016.
- H. Zhang, M. Cisse, Y. N. Dauphin, and D. Lopez-Paz, "mixup: Beyond empirical risk minimization," in *Proc. Int. Conf. Learn. Represent.*, 2018. [Online]. Available: <https://openreview.net/forum?id=r1Ddp1-Rb>
- Y. Gal, J. Hron, and A. Kendall, "Concrete dropout," in *Proc. Adv. Neural Inf. Process. Syst.*, 2017, pp. 3584–3593.
- S. Srivastava and G. Sharma, "OmniVec: Learning robust representations with cross modal sharing," in *Proc. IEEE/CVF Winter Conf. Appl. Comput. Vis.*, Jan. 2024, pp. 1236–1248.
- S. Verbitskiy, V. Berikov, and V. Vyshegorodtsev, "ERANNs: Efficient residual audio neural networks for audio pattern recognition," *Pattern Recognit. Lett.*, vol. 161, pp. 38–44, Sep. 2022.



JOHN FISCHER received the B.S. degree in information technology from the United States Naval Academy, in 2009, and the master's degree in computer and information technology from the University of Pennsylvania, in 2010. He is currently pursuing the Ph.D. degree with the Computer Science Department, Naval Postgraduate School. He is also an Active-Duty AH-1Z (Cobra) Pilot with United States Marine Corps. His research interests include machine learning, artificial intelligence, Bayesian deep learning, active learning, federated learning, passive sonar, and remote sensing.



ERIC ECKSTRAND received the B.S. degree in computer science from the United States Naval Academy and the joint master's degree in computer science and systems engineering from the University of Delaware and the Naval Postgraduate School, Monterey, CA, USA. He is currently a Machine Learning Researcher with the Data Science and Analytics Group, Naval Postgraduate School.

...



MARKO ORESCANIN (Member, IEEE) received the Ph.D. degree in electrical and computer engineering from the University of Illinois at Urbana-Champaign, Champaign, IL, USA, in 2010. Since 2019, he has been an Assistant Professor with the Computer Science Department, Naval Postgraduate School, Monterey, CA, USA. From 2011 to 2019, he was with Bose Corporation, MA, USA, where he primarily worked on research and advanced development of signal processing and machine learning algorithms for audio and speech enhancement in consumer electronics. He left Bose as the Senior Manager of the AI and Data Group, with a focus on Consumer Electronics Business Unit. His research interests include signal processing, machine learning, artificial intelligence, Bayesian deep learning, acoustics, passive and active sonar, unmanned vehicles, and remote sensing.