
BAR: Bayesian Activity Recognition using variational inference

Ranganath Krishnan
ranganath.krishnan@intel.com

Mahesh Subedar
mahesh.subedar@intel.com

Omesh Tickoo
omesh.tickoo@intel.com

Intel Labs
Hillsboro, OR (USA)

Abstract

Uncertainty estimation in deep neural networks is essential for designing reliable and robust AI systems. Applications such as video surveillance for identifying suspicious activities are designed with deep neural networks (DNNs), but DNNs do not provide uncertainty estimates. Capturing reliable uncertainty estimates in safety and security critical applications will help to establish trust in the AI system. Our contribution is to apply Bayesian deep learning framework to visual activity recognition application and quantify model uncertainty along with principled confidence. We utilize the stochastic variational inference technique while training the Bayesian DNNs to infer the approximate posterior distribution around model parameters and perform Monte Carlo sampling on the posterior of model parameters to obtain the predictive distribution. We show that the Bayesian inference applied to DNNs provide reliable confidence measures for visual activity recognition task as compared to conventional DNNs. We also show that our method improves the visual activity recognition precision-recall AUC by 6.2% compared to non-Bayesian baseline. We evaluate our models on Moments-In-Time (MiT) activity recognition dataset by selecting a subset of in- and out-of-distribution video samples.

1 Introduction

Activity recognition is an active area of research with multiple approaches depending on the application domain and the types of sensors [1, 2, 3]. In recent years, the DNN models applied to the visual activity recognition task outperform the earlier methods based on Gaussian Mixture Models and Hidden Markov Models [4, 5] using handcrafted features. The availability of large datasets [6, 7, 8] for visual analysis tasks has enabled the use of DNNs for activity recognition task. Vision-based activity recognition methods typically apply a combination of spatiotemporal features to classify the activities. Single frame based activity recognition methods usually apply DNN models trained on ImageNet dataset [9] to extract the spatial features. The temporal dynamics for activity recognition is typically modeled either by using a separate temporal sequence modeling such as variants of RNNs [10, 11] or by applying 3D ConvNets [12], which extend 2D ConvNets to the temporal dimension. In this work, we focus on the visual activity recognition applied to the trimmed video samples using the 3D ConvNet (C3D) architecture [13]

Probabilistic Bayesian models provide principled ways to gain insight about data and capture reliable uncertainty estimates in predictions. Bayesian deep learning has allowed bridging DNNs and probabilistic Bayesian theory to leverage the strengths of both methodologies [14, 15]. Conventional DNNs are trained to obtain the maximum likelihood estimates and they tend to disregard uncertainty around the model parameters that eventually leads to predictive uncertainty. DNNs may fail in the case of noisy or out-of-distribution data, leading to overconfident decisions that could be erroneous

as SoftMax probability does not capture overall model confidence. Instead, it represents relative probability that an input is from a particular class compared to the other classes.

Bayesian deep learning framework with Monte Carlo (MC) dropout technique has been used in visual scene understanding applications including camera relocalization [16], semantic segmentation [17], and depth regression [18]. In this work, we propose a Bayesian confidence measure applied to visual activity recognition task using the variational inference technique. To the best of our knowledge, this is the first research effort that applies Bayesian deep learning framework with variational inference for activity recognition task to capture reliable uncertainty measures.

2 Background

Bayesian neural networks (BNNs) offer a probabilistic interpretation of deep learning models by placing distributions over the model parameters. Bayesian inference can be applied to estimate the predictive distribution by propagating over the model likelihood and marginalizing over the learned posterior parameter distribution. BNNs also help in regularization by introducing distribution over network weights, capturing the posterior uncertainty around the neural network parameters. This allows transferring inherent DNN uncertainty from the parameter space to the predictive uncertainty of the unseen data.

Given training dataset $D = \{x, y\}$ with inputs $x = x_1, \dots, x_N$ and their corresponding outputs $y = y_1, \dots, y_N$, in parametric Bayesian settings we would like to infer a distribution over weights w as a function $y = f_w(x)$ that represents structure of the DNN model. With the posterior for model parameters inferred during Bayesian neural network training, we can predict the output for a new data point by propagating over the model likelihood $p(y|x, w)$ while drawing samples from the learned parameter posterior $p(w|D)$. Equation 1 below shows predictive distribution of output y^* given new input x^* :

$$p(y^* | x^*, D) = \int p(y^* | x^*, w) p(w | D) dw \quad (1)$$

In Bayesian neural networks, some of the techniques to achieve tractable inference include: (i) Markov Chain Monte Carlo (MCMC) sampling based probabilistic inference [14, 19] (ii) Variational inference techniques to infer the tractable approximate posterior distribution around model parameters [20, 21, 22] and (iii) Monte Carlo (MC) dropout approximate inference [23]. In our work, we use variational inference approach to infer the approximate posterior distribution around the model parameters.

Variational inference [24] is an active area of research in Bayesian deep learning, which uses gradient based optimization. This technique approximates complex probability distribution $p(w|D)$ with a simpler distribution $q_\theta(w)$, parameterized by variational parameters θ while minimizing the Kullback-Leibler (KL) divergence. Minimizing the KL divergence is equivalent to maximizing the log evidence lower bound [25, 23].

$$\mathcal{L}_{V\mathcal{I}} := \int q_\theta(w) \log p(y | x, w) dw - KL[q_\theta(w) || p(w)] \quad (2)$$

We evaluate the model uncertainty using Bayesian active learning by disagreement (BALD) [26] for the visual activity recognition task. BALD quantifies mutual information between parameter posterior distribution and predictive distribution, which captures model uncertainty, as shown in Equation 3.

$$BALD := H(y^* | x^*, D) - \mathbb{E}_{q_\theta(w)}[H(y^* | x^*, w)] \quad (3)$$

where, $H(y^* | x^*, D)$ is the predictive entropy given by:

$$H(y^* | x^*, D) = - \sum_{i=0}^{K-1} p_{i\mu} * \log p_{i\mu} \quad (4)$$

and $p_{i\mu}$ is predictive mean probability of i^{th} class from T Monte Carlo samples.




Figure 1: Vision-based activity recognition using ResNet101 C3D architecture. Bayesian variational inference is applied to the final three fully connected layers. Monte Carlo sampling on the posterior of model parameters provides the predictive distribution.

3 Bayesian DNN architecture

We use ResNet-101 C3D [13] architecture and replace the final layer with three fully connected variational layers followed by categorical distribution, as shown in Figure 1. The weights and bias parameters in the fully connected variational layers are modeled through mean-field normal distribution, and the network is trained using Bayesian variational inference based on KL divergence [21, 22]. We use Flipout [27], which is an efficient method that decorrelates the gradients within a mini-batch by implicitly sampling pseudo-independent weight perturbations for each input. We perform 40 stochastic forward passes on the final three fully connected variational layers with Monte Carlo sampling on the weight and bias posterior distributions, while the remaining layers of the network are considered to be deterministic. We compare the proposed architecture with MC dropout approximate Bayesian inference [23] method. We analyze the confidence measure and model uncertainty estimates for true (correct) and false (incorrect) predictions obtained from Bayesian DNN models.

For the comparison with the non-Bayesian baseline, we maintain the same model depth as the Bayesian DNN model and use three deterministic fully connected final layers for the non-Bayesian DNN model. The dropout layer is used after every fully connected layer to avoid over-fitting of the model. In the rest of the document, we refer the non-Bayesian DNN model as simply the DNN model.

4 Results

We evaluate our models on the MiT video activity recognition dataset [7] and train our model on a subset of 54 classes, which are used as in-distribution data. We use a different subset of 54 classes as out-of-distribution data. The subset of 54 classes for each category are selected after subjective




Figure 2: Comparison of DNN and Bayesian DNN visual activity recognition results using Precision-Recall and ROC plots.

Table 1: Comparison of accuracies for DNN, Bayesian DNN MC Dropout and Stochastic Variational Inference (Stochastic VI) models applied to the subset of MiT dataset (in-distribution classes).

Model	Top1 (%)	Top5 (%)
DNN Model	53.90	81.25
Bayesian DNN (MC Dropout) Model	53.40	80.81
Bayesian DNN (Stochastic VI) Model	54.10	81.40




Figure 3: Density histogram of confidence measures for true (correct) and false (incorrect) predictions using DNN and Bayesian DNN models: A distribution skewed towards right (near 1.0 on x-axis) indicates the model has higher confidence in predictions than the distribution skewed towards left. [The density histogram is a histogram with area normalized to one. Plots are overlaid with kernel density curves for better readability.]

evaluation to confirm the selected classes are not ambiguous and they fall into two distinct distribution of classes.

In Figure 2, comparison of precision-recall (left) and ROC (right) plots are presented for the three models. The Bayesian DNN stochastic variational inference (Stochastic VI) results achieve an improvement of 6.2% in precision-recall AUC over the non-Bayesian DNN model and an improvement of 3.63% over the MC dropout approximate Bayesian inference model. The ROC AUC curve for Bayesian DNN stochastic variational inference shows an improvement of 1.3% over the DNN results. The improvement in AUC for the stochastic variational inference model applied to activity recognition task is attributed to better measure of the confidence estimates obtained from the predictive distribution through Monte Carlo sampling. The classification accuracy for MiT in-distribution validation data is presented in Table 1. The accuracies for proposed stochastic variational inference model is similar to the DNN model.




Figure 4: Density histogram of confidence measures for in- and out-of-distribution samples using DNN and Bayesian DNN models. [The density histogram is a histogram with area normalized to one. Plots are overlaid with kernel density curves for better readability.]

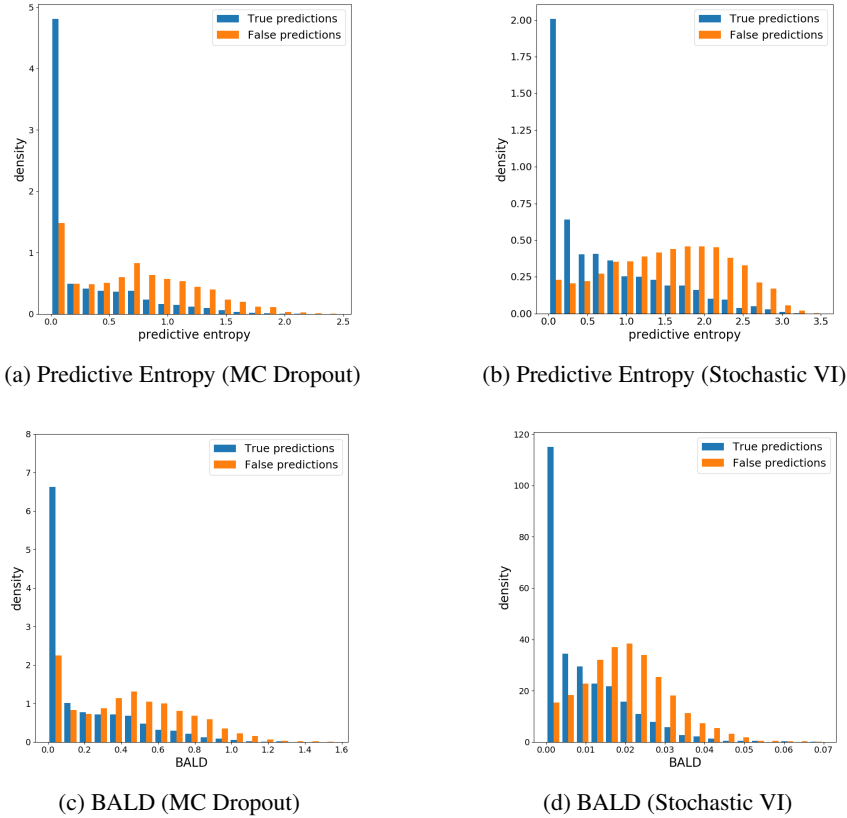


Figure 5: Density histogram of uncertainty measures from Bayesian DNN MC dropout and stochastic variational inference (Stochastic VI) models for true (correct) and false (incorrect) predictions from in-distribution samples. [The density histogram is a histogram with area normalized to one.]

In Figure 3, the density histograms of confidence measure for true (correct) and false (incorrect) predictions is presented. The density histogram is a histogram with area normalized to one. The confidence measure for the conventional DNN is the SoftMax probabilities used for the predictions. The mean of categorical predictive distribution obtained from Monte Carlo sampling provides the confidence measure for Bayesian DNNs. The DNN model density histogram shows a peak near higher confidence values for the true predictions, similar to the Bayesian DNN models. The DNN SoftMax probability measure (Figure 3 (a)) is skewed towards higher probability values for the false predictions. On the contrary, predictive mean probabilities obtained from the Bayesian DNN Monte Carlo dropout (Figure 3 (b)) and stochastic variational inference (Figure 3 (c)) models provide more reliable confidence measure indicating lower confidence values for the false predictions.

In Figure 4, density histogram of confidence measures for in- and out-of-distribution samples is presented. All the models show higher confidence values for in-distribution data. For out-of-distribution data, DNN SoftMax probability values (Figure 4 (a)) show overconfident measure for the erroneous predictions, whereas Bayesian DNNs (Figure 4 (b) & (c)) models show lower confidence for out-of-distribution data. Stochastic variational inference approach gives better results than MC dropout approach showing a more pronounced peak towards lower value for out-of-distribution samples.

Bayesian DNN models quantify uncertainty which is beneficial to identify false predictions and out-of-distribution samples. We compare predictive entropy and BALD [26] uncertainty measures for Bayesian Monte Carlo dropout and stochastic variational inference models. In Figure 5, the density histogram of uncertainty estimates for true and false predictions using the MiT dataset is presented. The uncertainty estimates for both the models are skewed towards higher values for the false predictions, but are skewed towards lower uncertainty for the true predictions. This indicates the models reliably capture the uncertainty associated with the predictions.

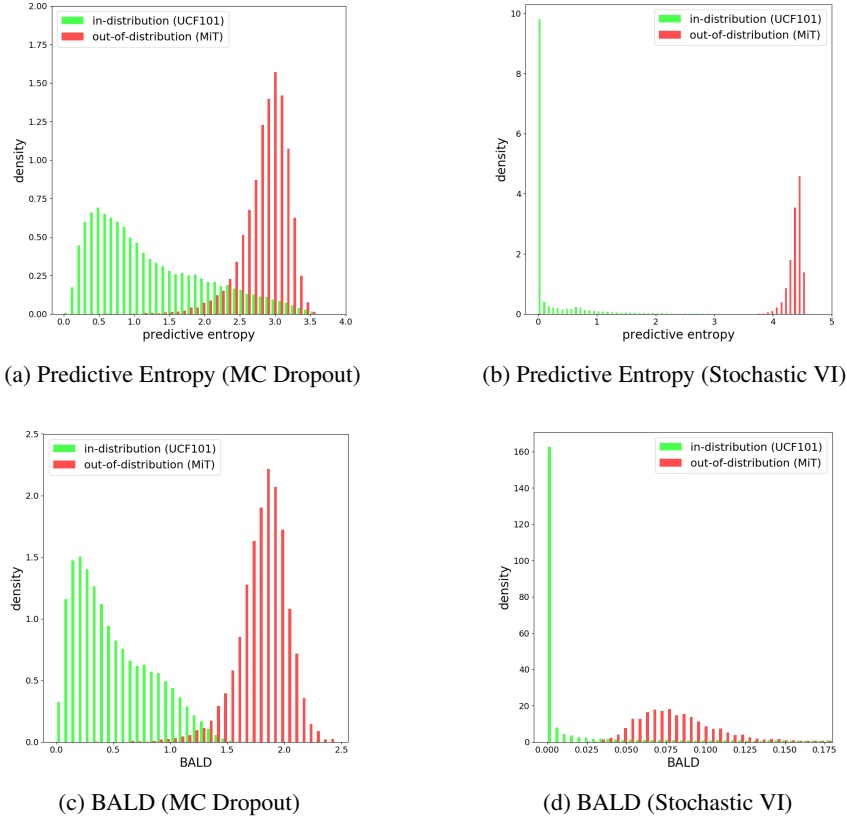


Figure 6: Density histogram of uncertainty measures from Bayesian DNN MC dropout and stochastic variational inference (Stochastic VI) models for UCF101 dataset as in-distribution samples and MiT datasets as out-of-distribution samples. [The density histogram is a histogram with area normalized to one.]

We also compare the uncertainty estimates using UCF101 [28] visual action recognition dataset consisting of 101 activity classes. The comparison of uncertainty measures for UCF101 dataset as in-distribution samples and the MiT dataset as out-of-distribution samples is shown in Figure 6. Predictive entropy and BALD uncertainty measures demonstrate a clear separation of in- and out-of-distribution samples. The stochastic variational inference approach shows better separation of the peaks for in- and out-of-distribution samples (Figure 6 (b) & (d)) compared to the Monte Carlo dropout method (Figure 6 (a) & (c)).

These results indicate the Bayesian DNNs for activity recognition task provide better confidence measure compared to the non-Bayesian DNN model, along with uncertainty estimates to reliably identify out-of-distribution data.

5 Conclusions

We presented a novel Bayesian DNN model for visual activity recognition which uses stochastic variational inference technique during training to infer the approximate posterior distribution around model parameters and perform Monte Carlo sampling on the posteriors. We evaluated our model using a subset of classes from the Moments-in-Time (MiT) dataset. The presented results demonstrate Bayesian DNNs can provide more reliable confidence measures as compared to the conventional DNNs. The uncertainty estimates obtained for wrong predictions and out-of-distribution samples indicate Bayesian DNNs can be more reliable for the visual activity recognition. We envision to extend this framework to multimodal inputs (e.g., audio and vision) where the confidence estimates obtained from Bayesian DNNs can help in fusing information based on uncertainty estimates to achieve better predictions.

References

- [1] Andreas Bulling, Ulf Blanke, and Bernt Schiele. A tutorial on human activity recognition using body-worn inertial sensors. *ACM Computing Surveys (CSUR)*, 46(3):33, 2014.
- [2] Shawn Hershey, Sourish Chaudhuri, Daniel PW Ellis, Jort F Gemmeke, Aren Jansen, R Channing Moore, Manoj Plakal, Devin Platt, Rif A Saurous, Bryan Seybold, et al. Cnn architectures for large-scale audio classification. In *Acoustics, Speech and Signal Processing (ICASSP), 2017 IEEE International Conference on*, pages 131–135. IEEE, 2017.
- [3] Fabian Caba Heilbron, Victor Escorcia, Bernard Ghanem, and Juan Carlos Niebles. Activitynet: A large-scale video benchmark for human activity understanding. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 961–970, 2015.
- [4] Lasitha Piyathilaka and Sarath Kodagoda. Gaussian mixture based hmm for human daily activity recognition using 3d skeleton features. In *Industrial Electronics and Applications (ICIEA), 2013 8th IEEE Conference on*, pages 567–572. IEEE, 2013.
- [5] Eunju Kim, Sumi Helal, and Diane Cook. Human activity recognition and pattern discovery. *IEEE Pervasive Computing*, 9(1):48, 2010.
- [6] Will Kay, Joao Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, et al. The kinetics human action video dataset. *arXiv preprint arXiv:1705.06950*, 2017.
- [7] Mathew Monfort, Bolei Zhou, Sarah Adel Bargal, Alex Andonian, Tom Yan, Kandan Ramakrishnan, Lisa Brown, Quanfu Fan, Dan Gutfruehd, Carl Vondrick, et al. Moments in time dataset: one million videos for event understanding. *arXiv preprint arXiv:1801.03150*, 2018.
- [8] Sami Abu-El-Haija, Nisarg Kothari, Joonseok Lee, Paul Natsev, George Toderici, Balakrishnan Varadarajan, and Sudheendra Vijayanarasimhan. Youtube-8m: A large-scale video classification benchmark. *arXiv preprint arXiv:1609.08675*, 2016.
- [9] Christian Szegedy, Sergey Ioffe, Vincent Vanhoucke, and Alexander A Alemi. Inception-v4, inception-resnet and the impact of residual connections on learning. In *AAAI*, volume 4, page 12, 2017.
- [10] Limin Wang, Yuanjun Xiong, Zhe Wang, Yu Qiao, Dahua Lin, Xiaoou Tang, and Luc Van Gool. Temporal segment networks: Towards good practices for deep action recognition. In *European Conference on Computer Vision*, pages 20–36. Springer, 2016.
- [11] Bolei Zhou, Alex Andonian, and Antonio Torralba. Temporal relational reasoning in videos. *arXiv preprint arXiv:1711.08496*, 2017.
- [12] Du Tran, Lubomir Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. Learning spatiotemporal features with 3d convolutional networks. In *Proceedings of the IEEE international conference on computer vision*, pages 4489–4497, 2015.
- [13] Kensho Hara, Hirokatsu Kataoka, and Yutaka Satoh. Can spatiotemporal 3d cnns retrace the history of 2d cnns and imagenet? In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6546–6555, 2018.
- [14] Radford M Neal. *Bayesian learning for neural networks*, volume 118. Springer Science & Business Media, 2012.
- [15] Yarin Gal. *Uncertainty in deep learning*. University of Cambridge, 2016.
- [16] Alex Kendall and Roberto Cipolla. Modelling uncertainty in deep learning for camera relocalization. In *2016 IEEE international conference on Robotics and Automation (ICRA)*, pages 4762–4769. IEEE, 2016.
- [17] Alex Kendall, Vijay Badrinarayanan, and Roberto Cipolla. Bayesian segnet: Model uncertainty in deep convolutional encoder-decoder architectures for scene understanding. *arXiv preprint arXiv:1511.02680*, 2015.

- [18] Alex Kendall, Yarin Gal, and Roberto Cipolla. Multi-task learning using uncertainty to weigh losses for scene geometry and semantics. *arXiv preprint arXiv:1705.07115*, 3, 2017.
- [19] Max Welling and Yee W Teh. Bayesian learning via stochastic gradient langevin dynamics. In *Proceedings of the 28th International Conference on Machine Learning (ICML-11)*, pages 681–688, 2011.
- [20] Alex Graves. Practical variational inference for neural networks. In *Advances in neural information processing systems*, pages 2348–2356, 2011.
- [21] Rajesh Ranganath, Sean Gerrish, and David M Blei. Black box variational inference. *arXiv preprint arXiv:1401.0118*, 2013.
- [22] Charles Blundell, Julien Cornebise, Koray Kavukcuoglu, and Daan Wierstra. Weight uncertainty in neural networks. *arXiv preprint arXiv:1505.05424*, 2015.
- [23] Yarin Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning*, pages 1050–1059, 2016.
- [24] David M Blei, Alp Kucukelbir, and Jon D McAuliffe. Variational inference: A review for statisticians. *Journal of the American Statistical Association*, 112(518):859–877, 2017.
- [25] Christopher M Bishop. Pattern recognition and machine learning (information science and statistics) springer-verlag new york. Inc. Secaucus, NJ, USA, 2006.
- [26] Neil Houlsby, Ferenc Huszár, Zoubin Ghahramani, and Máté Lengyel. Bayesian active learning for classification and preference learning. *arXiv preprint arXiv:1112.5745*, 2011.
- [27] Yeming Wen, Paul Vicol, Jimmy Ba, Dustin Tran, and Roger Grosse. Flipout: Efficient pseudo-independent weight perturbations on mini-batches. *arXiv preprint arXiv:1803.04386*, 2018.
- [28] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. Ucf101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*, 2012.