Estimation of ML model uncertainty in particle physics event classifiers [link]

Julia Vázquez-Escobar (CIEMAT,Madrid)

Bayesian neural network approximation

Bayesian neural networks (BNN) [1] are a Bayesian approach to artificial neural networks. The calculation of posterior probabilities for the weights of the network, given some observed data, is proposed to obtain uncertainty measures, but it generally becomes intractable for a network with more than one hidden layer. An approximation resorting to dropout is presented in [2] and used in this work. Two results from the reference allow us to estimate the expected prediction of a BNN and its variance using a regular multilayer perceptron (MLP) with dropout applied.

$$\widetilde{\mathbb{E}}\left[\mathbf{y}^*\right] := \frac{1}{T}\sum_{t=1}^{T}\mathbf{f}^{\hat{\omega}_t}\left(\mathbf{x}^*\right) \underset{T\to\infty}{\longrightarrow} \mathbb{E}_{q_\theta^*(\mathbf{y}^*|\mathbf{x}^*)}\left[\mathbf{y}^*\right]$$

$$\widetilde{\mathrm{Var}}\left[\mathbf{y}^*\right] := \tau^{-1}\mathbf{I} + \frac{1}{T}\sum_{t=1}^{T}\mathbf{f}^{\hat{\omega}t}\left(\mathbf{x}^*\right)^T\mathbf{f}^{\hat{\omega}_t}\left(\mathbf{x}^*\right) - \widetilde{\mathbb{E}}\left[\mathbf{y}^*\right]^T\widetilde{\mathbb{E}}\left[\mathbf{y}^*\right] \underset{T\to\infty}{\to} \mathrm{Var}_{q_\theta^*(\mathbf{y}^*|\mathbf{x}^*)}\left[\mathbf{y}^*\right]$$

Local ensembles

Another approach to obtaining uncertainty measures when predicting with a neural network is using the Local Ensembles method [3]. This algorithm estimates the variability of a test observation's prediction when varying the model within a set that fits the training data equally well. It relies on the calculation of the prediction's gradient with respect to the parameters of the model and the Hessian of the loss function.

Proposition 1.1 *Let Δ_θ be the projection of a random perturbation with mean zero and covariance proportional to the identity $\epsilon \cdot I$ into the ensemble subspace spanned by $\{\xi_{(j)} : j > m\}$. Let P_Δ be the linearized change in prediction induced by the perturbation*

$$P_\Delta(x') := g_{\theta*}(x')^\top \Delta_\theta \approx \hat{y}(x', \theta * + \Delta_\theta) - \hat{y}(x', \theta*)$$

then $\mathcal{E}_m(x') = \varepsilon^{-1/2} \cdot SD(P_\Delta(x'))$.

Probabilistic random forest

We also consider the Probabilistic Random Forest (PRF) method [4] where a modified approach to the standard Random Forest algorithm is taken. Random Forest is an ensemble method where several decision trees (a non-parametric model) are trained to fit some data. Standard decision trees are trained to fit some data, classifying events in classes according to the values of their input features. PRF considers these features as random variables and calculates the probability of an event belonging to each class instead of obtaining a final class for the event. Since it provides output probabilities we can obtain variability measures for the prediction.

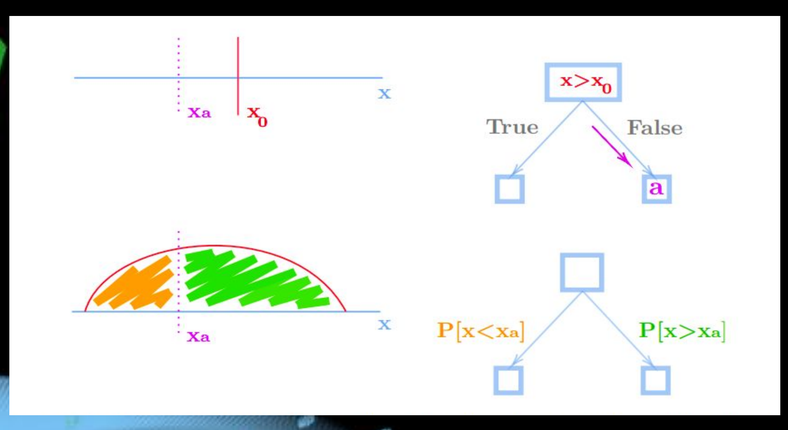

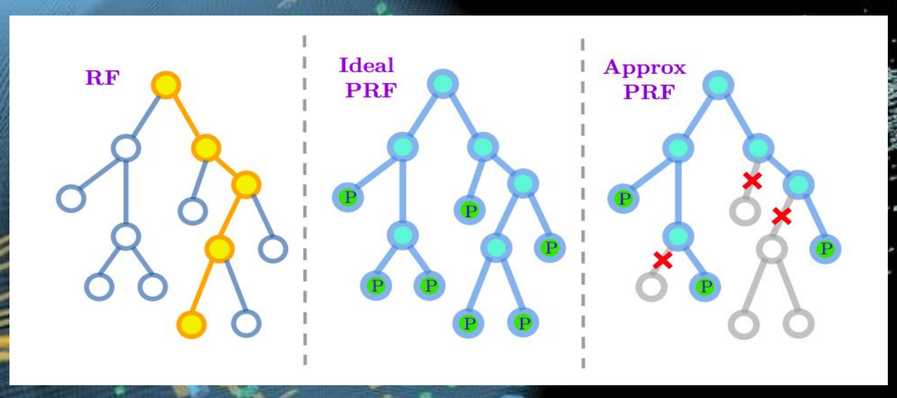

Results

The three approaches we propose make use of Machine Learning algorithms widely used for classification tasks; the BNN approximation and the Local Ensembles method are applied to a MLP whereas PRF is a modification of the standard Random Forest algorithm. They present the novelty of providing predictions together with their uncertainties related to the modelling, assessing the quality of the model.

	AUC
PRF	0.969 ± 0.005
LE	0.951 ± 0.006
BA	0.9904 ± 0.0012

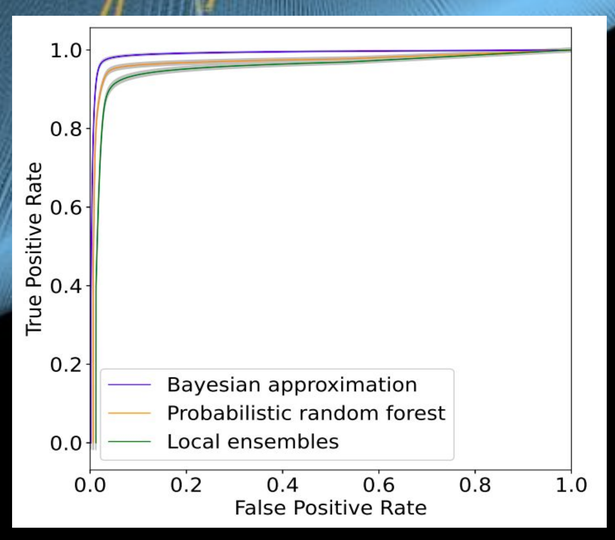

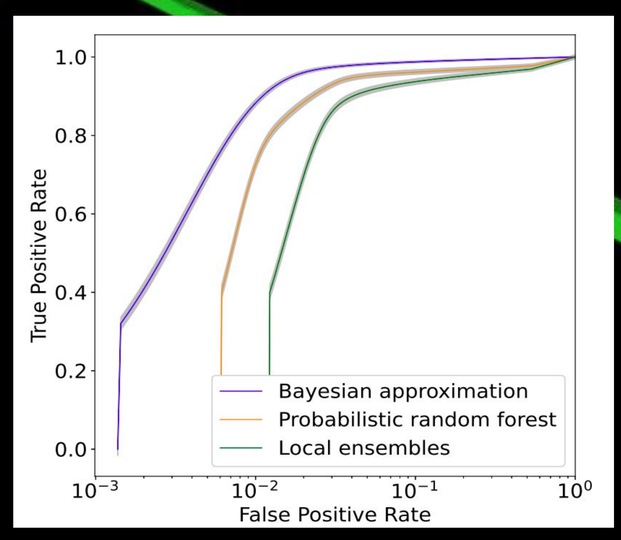

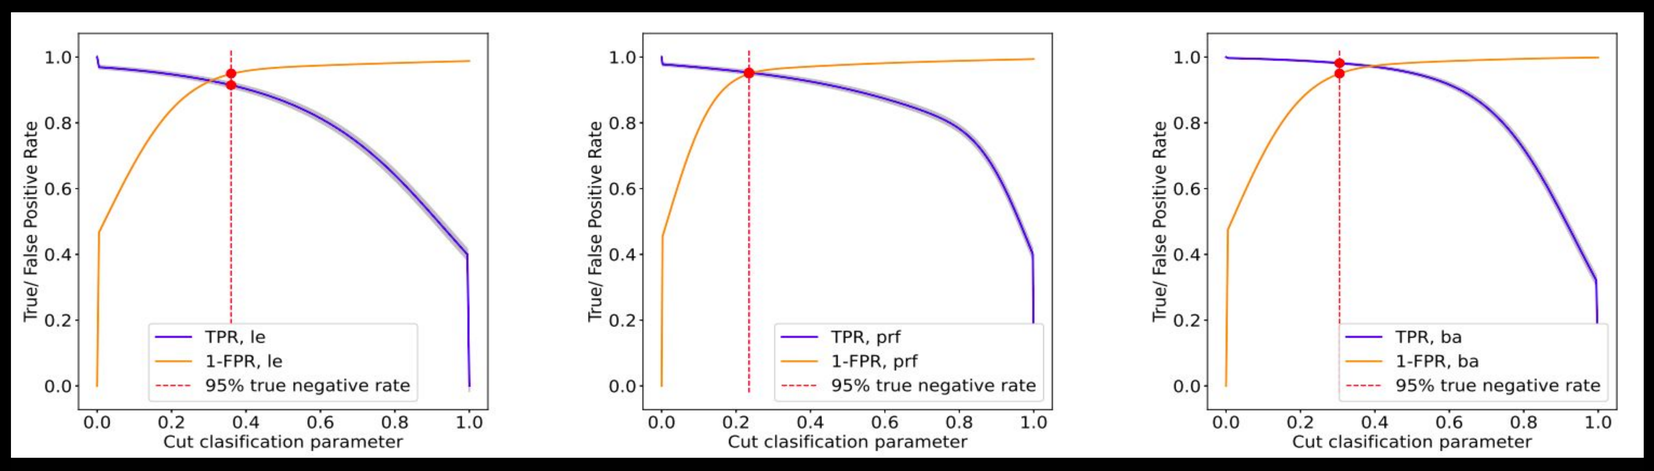

(a) Local Ensembles (b) Probabilistic random forest (c) ANN with Dropout

	TPR when FNR=95%
PFR	0.953 ± 0.010
LE	0.915 ± 0.008
BA	0.982 ± 0.003

[1] D.J.C. MacKay, Neural Comput. 4 (3) (1992) 448–472.; [2] Y. Gal, in: Uncertainty in Deep Learning, 2016, pp. 29–61, ch. 3.
[3] D. Madras, J. Atwood, A. D'Amour, in: ICLR 2020 International Conference on Learning Representations, 2019.
[4] I. Reis, D. Baron, S. Shahaf, Probabilistic random forest: a machine learning algorithm for noisy datasets, arXiv:1811.05994, 2018, p. 17.