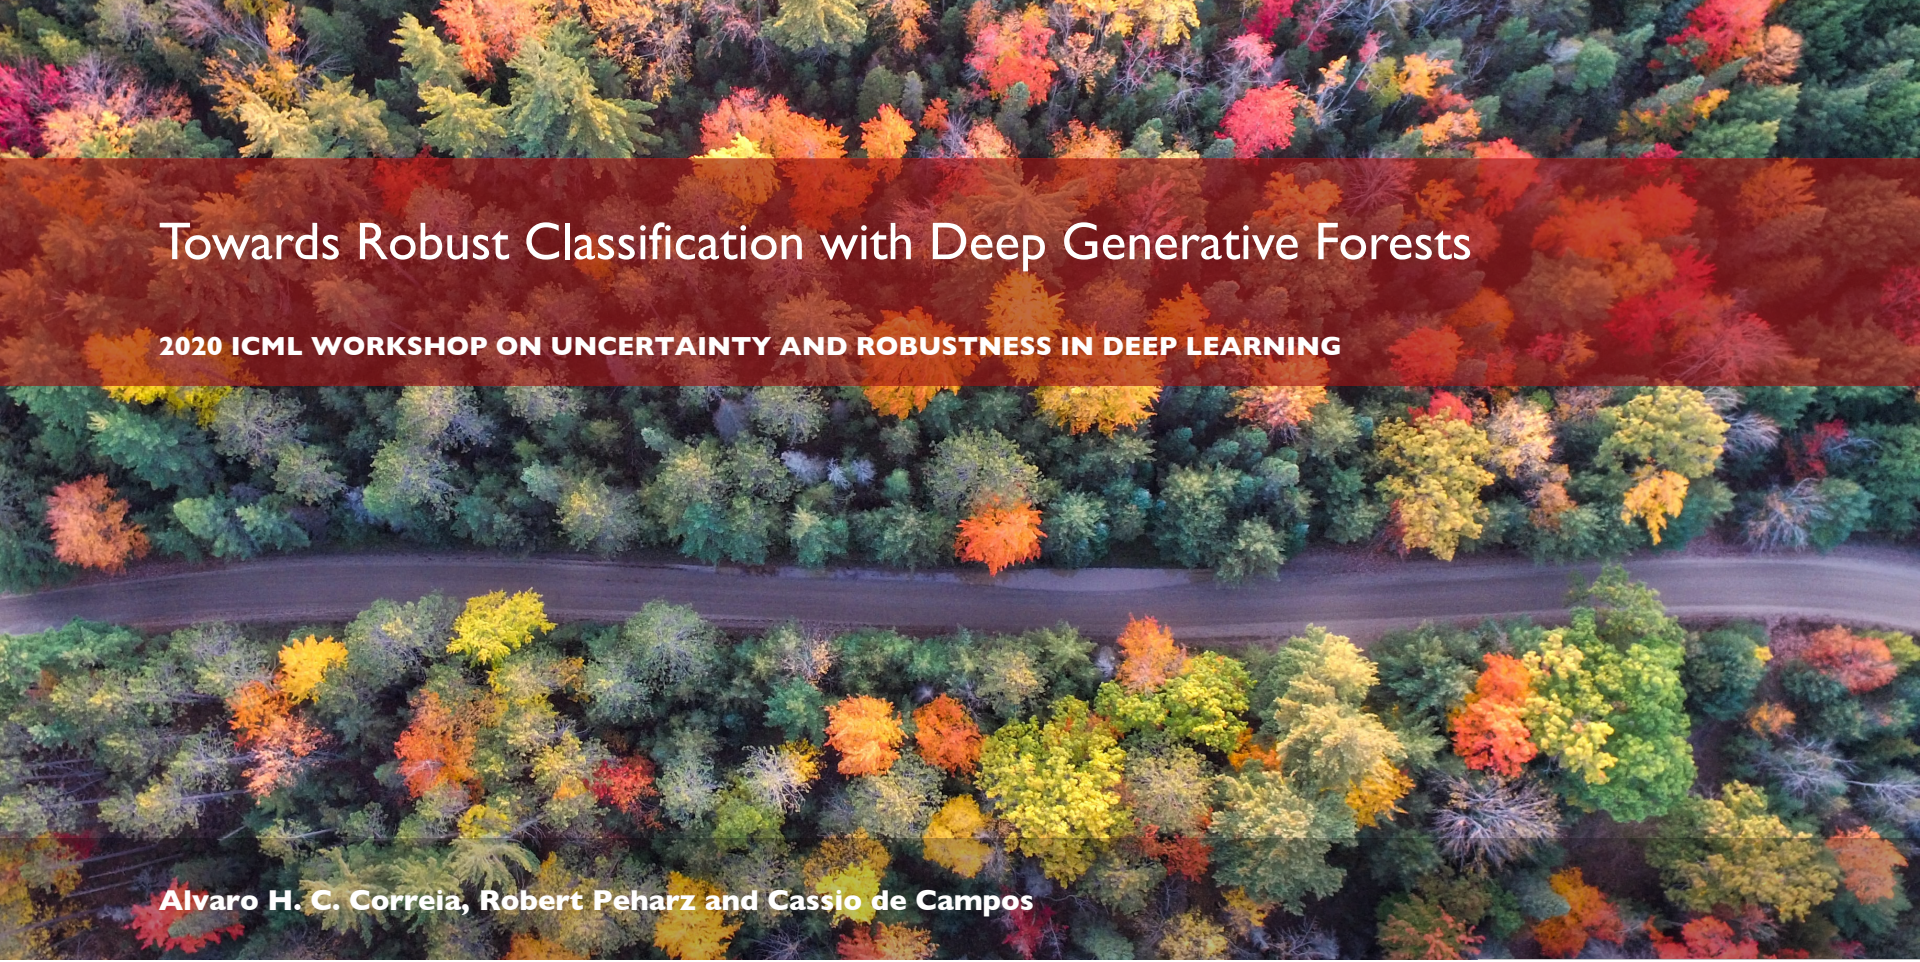

Towards Robust Classification with Deep Generative Forests

2020 ICML WORKSHOP ON UNCERTAINTY AND ROBUSTNESS IN DEEP LEARNING

Alvaro H. C. Correia, Robert Peharz and Cassio de Campos

Department of Mathematics and Computer Science
Eindhoven University of Technology

TU/e

In a nutshell

Extension of Random Forests to a full generative model.

- Discriminative structure learning.
- Model a joint distribution $p(Y, \boldsymbol{X})$ instead of a function $f: \boldsymbol{x} \mapsto y$.

Enhance Random Forests

- Outlier detection
- Robust classification
- Inference with missing values

Probabilistic Circuits (PCs)

Distribution nodes

compute a probability density over some subset of the features.

Sum nodes

compute convex combinations over their children.

Product nodes

compute the product over their children.

Probabilistic Circuits (PCs)

Wide range of tractable inference routines.

Efficient Marginalisation

Compute $p(\boldsymbol{X}')$ for any $\boldsymbol{X}' \subseteq \boldsymbol{X}$ in a single pass through the network.

Notable examples of Probabilistic Circuits

- Sum-Product Networks (Poon and Domingos 2011)
- Arithmetic Circuits (Darwich 2003)
- Cut-set Networks (Rahman et al. 2014)
- Probabilistic Sentential Decision Diagrams (Kisa et al. 2014)

Probabilistic Circuits

Example

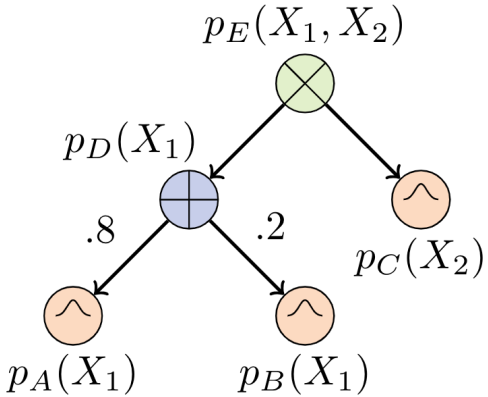

Evaluate at
$X_1 = 1$ and $X_2 = 2$

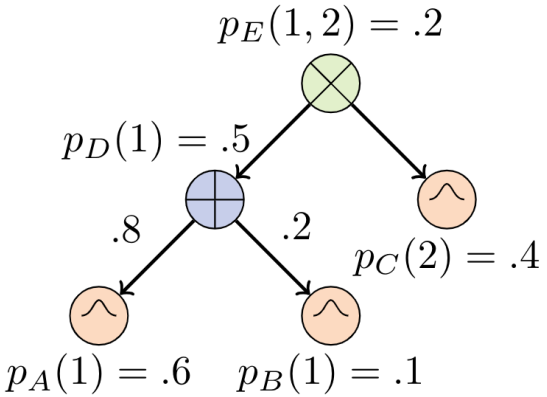

$$p_D(X_1) = 0.8p_A(X_1) + 0.2p_B(X_1)$$

$$p_E(X_1, X_2) = p_D(X_1)p_E(X_2) = \big(0.8p_A(X_1) + 0.2p_B(X_1)\big)p_C(X_2)$$

Generative Decision Trees (GeDTs) and Forests (GeFs)

Representation of Decision Trees as Probabilistic Circuit

- Convert each internal node to a sum node

  Weights are given by the probability mass of each children

- Convert each leaf into a distribution node

  Fit a density over the instances in each leaf

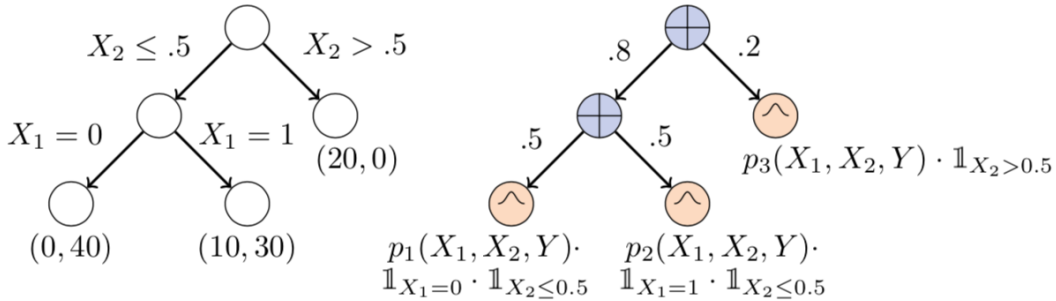

TU/e

Outlier detection

Generative Models are Natural Outlier Detectors

We detect out-of-domain samples simply by monitoring $p(\boldsymbol{X})$.

$$p(\boldsymbol{x}) = \sum_{y} p(y, \boldsymbol{x})$$

This comes at little extra cost, since Generative Forests perform classification over the joint, and all terms $p(y, \boldsymbol{x})$ are already computed in the context of classification.

$$\hat{y} = \text{argmax}_y \sum_{y} p(y, \boldsymbol{x})$$

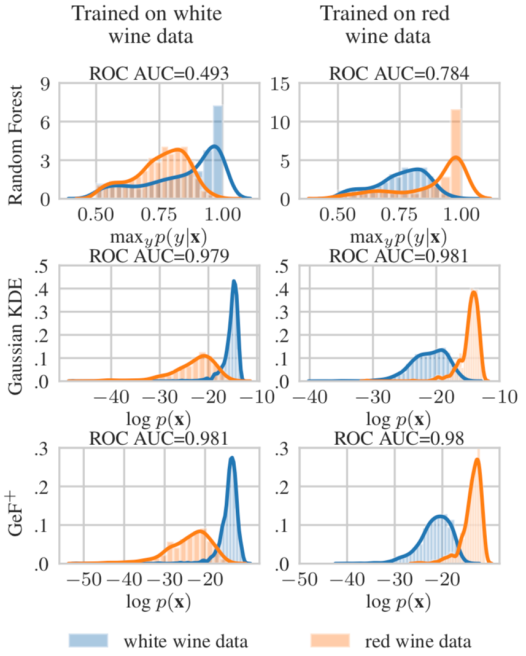

Robust Classification

Sensitivity analysis

Perturb the model parameters until the predicted class changes.

ϵ-contamination of a vector of parameters $\boldsymbol{w}$

$$C_{\boldsymbol{w},\epsilon} = \{(1-\epsilon)\boldsymbol{w} + \epsilon\boldsymbol{v} : v_j \geq 0, \sum v_j = 1\}$$

ϵ-robustness

The largest ϵ for which all parameters in $C_{w,\epsilon}$ yield the same classification.

$$\forall y' \neq y : \max_{\boldsymbol{w} \in C_{\boldsymbol{w},\epsilon}} \mathbb{E}_{\boldsymbol{w}} [\mathbb{I}(Y = y') - \mathbb{I}(Y = y) \mid \boldsymbol{x}] < 0$$

Robust Classification

ϵ-robustness correlates to accuracy

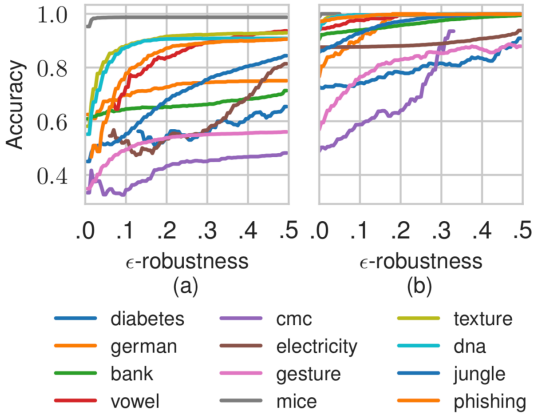

Accuracy of predictions with ϵ-robustness (a) below and (b) above different thresholds for 12 OpenML datasets.

Samples from (Fashion-)Mnist datasets with lowest (left) and highest (right) ϵ-robustness in the test set.

Robust Classification

ϵ-robustness differs substantially from $p(\boldsymbol{x})$

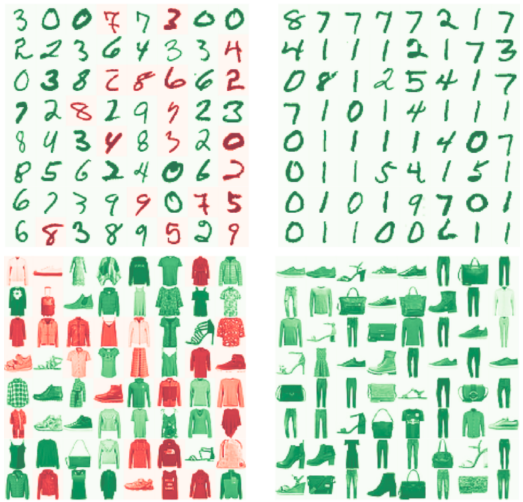

Samples from (Fashion-)Mnist datasets with lowest (left) and highest (right) ϵ-robustness in the test set.

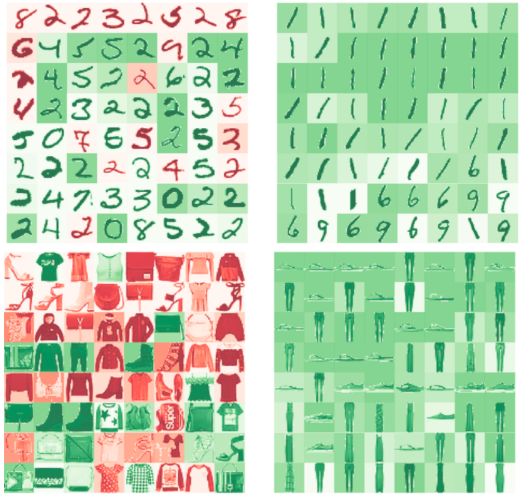

Samples from (Fashion-)Mnist datasets with lowest (left) and highest (right) $p(\boldsymbol{x})$ in the test set.

Inference with Missing Values

Marginalise the non-observed features $X_{\neg o}$

$$P(Y|X_o) = \frac{\int_{x_{\neg o}} P(Y, X_o, x_{\neg o}) dx_{\neg o}}{\sum_y \int_{x_{\neg o}} P(Y, X_o, x_{\neg o}) dx_{\neg o}}$$

Marginalisation with Generative Forests is tractable!
Perform marginalisation at the leaves and evaluate the trees as usual.

We can show this classifier is Bayes-consistent for any pattern of missing values!

Inference with Missing Values

Some Experimental Results

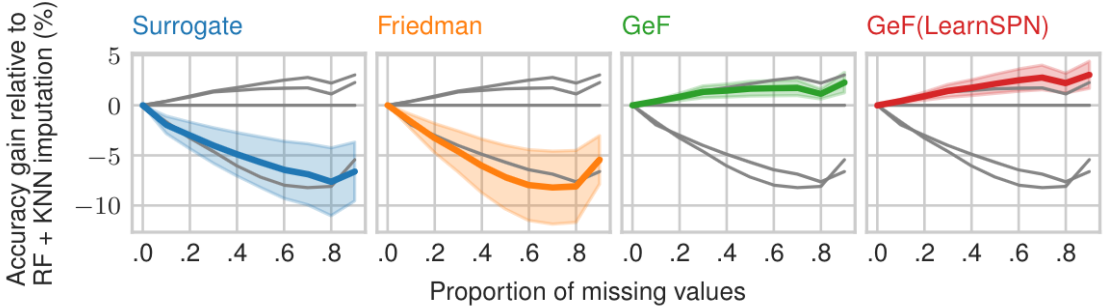

Average (across 21 datasets) accuracy gain relative to RFs (100 trees) plus KNN imputation against percent of missing values. Confidence intervals (95%) are also computed across the datasets.

Conclusion

Generative Forests

- Recent class of Probabilistic Circuits
- Hybrid discriminative-generative models
- Preserve the prediction function and overall characteristics of Random Forests
- Equip Random Forests with cool generative properties:
  - Outlier detection
  - Robust classification
  - Handling missing values

Thanks for your time!