

---

# Towards Robust Classification with Deep Generative Forests

---

Alvaro H. C. Correia<sup>1</sup> Robert Peharz<sup>1</sup> Cassio de Campos<sup>1</sup>

## Abstract

Decision Trees and Random Forests are among the most widely used machine learning models, and often achieve state-of-the-art performance in tabular, domain-agnostic datasets. Nonetheless, being primarily discriminative models they lack principled methods to manipulate the uncertainty of predictions. In this paper, we exploit Generative Forests (GeFs), a recent class of deep probabilistic models that addresses these issues by extending Random Forests to generative models representing the full joint distribution over the feature space. We demonstrate that GeFs are uncertainty-aware classifiers, capable of measuring the robustness of each prediction as well as detecting out-of-distribution samples.

## 1. Introduction

Decision Trees (DTs) and Random Forests (RFs) are arguably the most popular non-linear machine learning models of today. In Kaggle’s 2019 report on the *State of Data Science and Machine Learning* (Kaggle, 2019), DTs and RFs appear as second most widely used techniques, right after linear and logistic regressions. Moreover, decision trees are often considered interpretable (Freitas, 2014) and hence have enjoyed a surge in popularity with the increasing interest in explainable artificial intelligence. Nonetheless, efforts towards uncertainty-aware and reliable tree-based models are still comparatively scarce.

In this paper, we demonstrate that some of these shortcomings are addressed by Generative Forests (GeFs) (Correia et al., 2020), a class of deep probabilistic models that subsumes Random Forests. In particular, we show in a number of classification tasks that GeFs enable new principled methods to i) estimate the uncertainty of *each* of the model’s predictions and ii) monitor the input distribution to detect out-of-domain samples or distribution shifts.

---

<sup>1</sup>Eindhoven University of Technology. Correspondence to: Alvaro H. C. Correia <a.h.chaim.correia@tue.nl>.

## 2. Generative Forests

Before discussing the main ideas of the paper, we introduce Generative Forests and the required notation. As we focus on classification tasks, we denote the set of explanatory variables as  $\mathbf{X} = \{X_1, X_2, \dots, X_m\}$  and the target variable as  $Y$ . As usual, we write realisations of random variables (or collections thereof) in lowercase; for example,  $\mathbf{X} = \mathbf{x}$  or  $Y = y$ . We assume the pair  $(\mathbf{X}, Y)$  is drawn from a fixed joint distribution  $\mathbb{P}^*(\mathbf{X}, Y)$  with density  $p^*(\mathbf{X}, Y)$  and that, while the true distribution  $\mathbb{P}^*$  is unknown, we have a dataset  $\mathcal{D}_n = \{(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_n, y_n)\}$  of  $n$  i.i.d. samples from  $\mathbb{P}^*$ .

Generative Forests are in fact a class of *Probabilistic Circuits* (PCs) (Van den Broeck et al., 2019) satisfying smoothness and decomposability (Peharz et al., 2015). PCs are a family of deep density representations facilitating many exact and efficient inference routines (Darwiche, 2003; Van den Broeck et al., 2019). In short, they are computational graphs with three types of nodes: i) *distribution nodes*, ii) *sum nodes* and iii) *product nodes*. Distribution nodes compute a probability density (by an adequate choice of the underlying measure, this also subsumes probability mass functions) over some subset  $\mathbf{X}' \subseteq \mathbf{X}$ , that is, a normalised function mapping the state space of  $\mathbf{X}'$  to the non-negative real numbers. Sum nodes compute convex combinations over their children: if  $v$  is a sum node and  $\text{ch}(v)$  its children, then  $v$  computes  $v(\mathbf{x}) = \sum_{u \in \text{ch}(v)} w_{v,u} u(\mathbf{x})$ , where  $w_{v,u} \geq 0$  and  $\sum_{u \in \text{ch}(v)} w_{v,u} = 1$ . Product nodes compute the product over their children; if  $v$  is a product node, then  $v(\mathbf{x}) = \prod_{u \in \text{ch}(v)} u(\mathbf{x}_u)$ , with the collection  $\{\mathbf{x}_u\}_{u \in \text{ch}(v)}$  a partition (non-overlapping projections) of  $\mathbf{x}$ . Finally, a (smooth and decomposable) PC represents the density over all variables (here  $p(\mathbf{X}, Y)$ ) computed by its root node.

Generative Forests are best understood by relating individual decision trees to Probabilistic Circuits. For any given DT, we can construct a corresponding PC—a *Generative Decision Tree* (GeDT)—representing a full joint density  $p(\mathbf{X}, Y)$ . In a nutshell, each decision node is converted into a sum node and each leaf into a density with support restricted to the leaf’s cell. The training samples can be figured to be routed from the root node to the leaves, following the decisions at each decision/sum node. The sum weights are given by the fraction of samples which are routed from the sum node to each of its children. The leaf densities are

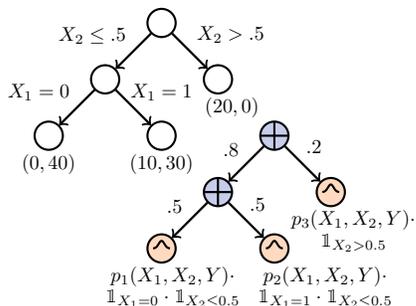


Figure 1. Illustration of a DT and its corresponding PC.

learned on the data which arrives at the respective leaves.

Note that GeDTs are proper PCs over  $(\mathbf{X}, Y)$ , albeit rather simple ones: they are tree-shaped and contain only sum nodes. Nonetheless, GeDTs are in fact a class of models, as we are free to fit arbitrarily complex functions at the leaves; say graphical models, again PCs or even advanced density estimators such as a VAEs (Kingma & Welling, 2014) or Flows (Rezende & Mohamed, 2015). In this work, however, we focus on arguably the simplest density estimator, and model the density at the leaves as  $p(\mathbf{X}, Y) = p(X_1) \dots p(X_m)p(Y)$ , with continuous and categorical variables represented by univariate normal and multinomial distributions, respectively. We show these *fully-factorised* leaves are already sufficient to equip standard RFs with effective and principled ways to detect outliers and estimate the robustness of each prediction.

The main semantic difference between DTs and GeDTs is that a DT represents a classifier, that is, a conditional distribution  $f(\mathbf{x})$ , while the corresponding GeDT encodes a full joint distribution  $p(\mathbf{X}, Y)$ —the latter naturally lends itself to classification via the conditional distribution  $p(Y | \mathbf{x}) \propto p(\mathbf{x}, Y)$ . Note that, in theory,  $p(Y | \mathbf{x})$  might differ substantially from  $f(\mathbf{x})$ , as every feature might influence classification in a GeDT, even if it never appears in any decision node of the DT. Still, it is easy to see that if the distribution at the leaves satisfy  $p(\mathbf{X}, Y) = p(\mathbf{X})p(Y)$ , then a GeDT defines the same prediction function as the original DT.

Generative Forests are ensembles of GeDTs and can also be made equivalent to the original RF by an appropriate choice of density model at the leaves. However, instead of ensuring “backwards compatibility”, in this paper we are interested in exploiting the generative properties of GeFs. To that end, we extend GeFs to model a single joint by considering a uniform mixture of GeDTs (using a sum node over the trees). That is, instead of averaging over the conditional distributions of each of the  $n_t$  trees, we define a single model that represents the joint  $p(\mathbf{X}, Y) = n_t^{-1} \sum_{j=1}^{n_t} p_j(\mathbf{X}, Y)$ , where each  $p_j$  comes from a GeDT. Since this model is essentially a mixture of the different trees, we call it GeF<sup>+</sup>.

### 3. Outlier Detection

Most of machine learning theory relies on the assumption that training and test data are sampled from the same distribution. This is a reasonable assumption—but is often violated in practice, as real-world data is constantly evolving. Reliable machine learning models should then be able to identify such violations to either suspend judgement and fail gracefully or signal the need for further data gathering and retraining.

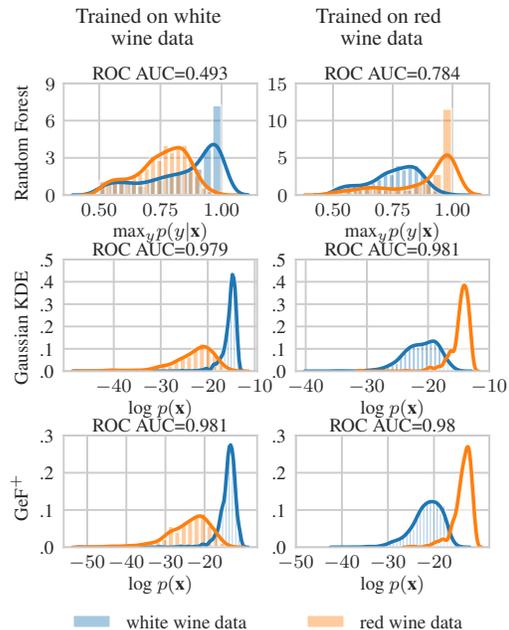


Figure 2. Normalised histograms of  $\log p(\mathbf{x})$  (KDE and GeF<sup>+</sup>) and  $\max_y p(y | \mathbf{x})$  (RF) of samples from red and white wine data.

Generative models offer a natural and principled way to detect outliers or distribution shifts. As they innately fit the joint distribution of the training data, they are capable of estimating the likelihood of every new sample, flagging unlikely ones as potential anomalies. In a GeF<sup>+</sup> this is done by monitoring the marginal  $p(\mathbf{X})$ , which comes at no extra cost; classification is performed over the joint  $p(Y, \mathbf{X})$  and computing  $p(\mathbf{x})$  only requires summing over all possible classes,  $p(\mathbf{x}) = \sum_y p(y, \mathbf{x})$ .

We illustrate outlier detection in GeF<sup>+</sup>s using the wine quality dataset (Cortez et al., 2009) (where the class is a scale of quality of wine) with a variant of transfer testing (Bradshaw et al., 2017). We learn two different GeF<sup>+</sup>s, each with only one type of wine (red or white), and compute the log-density of unseen data (70/30 train-test split) for the two wine types. As we see in the histograms of Figure 2, the marginal distribution over the joints does provide a strong signal to identify out-of-domain instances. We compare GeF<sup>+</sup>s to a Gaussian Kernel Density Estimator (KDE) and to a common baseline

for deep models (Hendrycks & Gimpel, 2016), whereby the probability of the predicted class,  $\max_y p(y|\mathbf{x})$ , is used as a signal to detect outliers. We see from the histograms and the ROC (receiver operating characteristic curve) scores, that our models largely outperform the baseline while being comparable to KDEs, even though the structure of a GeF<sup>+</sup> is learned in a discriminative manner.

We note that previous works have already proposed using Random Forests for outlier detection (Liu et al., 2008). However, these models are typically directly trained to identify anomalies and have that as their sole purpose, while GeF<sup>+</sup>s are unique in that, while being primarily classifiers (or regressors) they also effectively detect out-of-domain samples.

#### 4. Robust Classification

Outlier detection is related to the concept of *vagueness* or *epistemic uncertainty*, that is, the lack of sufficient statistical support for issuing a prediction. However, machine learning models are often confronted with cases where the data supports the thesis that a given instance is associated with more than a single class with high probability. That is commonly referred to as *aleatory uncertainty*. Effectively quantifying both types of uncertainty is indispensable in critical applications, where overconfident predictions may lead to catastrophic failures. One common approach to estimate the model’s confidence in classification tasks is to manipulate the reported probability  $p(Y|\mathbf{x})$  (Guo et al., 2017; Liang et al., 2017). Still, this is not only overly simplistic but also fails to distinguish the types of uncertainty.

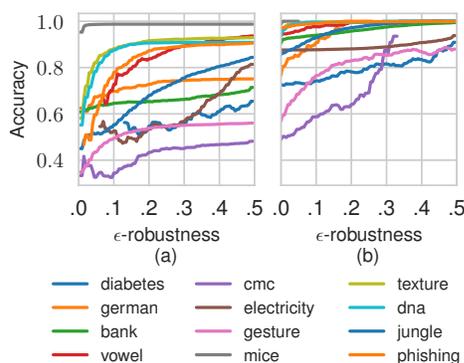


Figure 3. Accuracy of predictions with  $\epsilon$ -robustness (a) below and (b) above different thresholds for 12 OpenML datasets. Some curves end abruptly because we only computed the accuracy when 30 or more data points were available for a given threshold.

GeF<sup>+</sup>s offer an arguably more principled approach rooted in the notion of *robustness* (Dietterich, 2017) as obtained with credal sum-product networks (Mauá et al., 2017; Mauá et al., 2018). In a nutshell, we evaluate how much we can perturb all parameters of the model without changing its prediction on a given instance. Formally, we quantify this

perturbation with the concept of  $\epsilon$ -contamination for each of the sum nodes in a PC. If  $w$  is the vector of weights of a given sum node, then its  $\epsilon$ -contamination is given by the set

$$\mathcal{C}_{w,\epsilon} = \{(1 - \epsilon)w + \epsilon v : v_j \geq 0, \sum_j v_j = 1\}.$$

This definition naturally leads to the idea of  $\epsilon$ -robustness: the largest  $\epsilon$  for which all parameter configurations in  $\mathcal{C}_{w,\epsilon}$  yield the same classification. We run such analysis for all of the nodes at once: let  $\mathcal{C}_\epsilon$  represent the collection  $\{\mathcal{C}_{w,\epsilon}\}$  for all sum nodes in the PC and let  $w$  be one possible choice of a  $w$  in each of the nodes.<sup>1</sup> We compute whether there is a label  $y$  of the class such that

$$\forall y' \neq y : \max_{w \in \mathcal{C}_\epsilon} \mathbb{E}_w[\mathbb{1}(Y = y') - \mathbb{1}(Y = y)|\mathbf{x}] < 0,$$

and if so, we declare  $y$  robust for threshold  $\epsilon$ . The maximum  $\epsilon$  such that this is true (which we find by binary search) is what we call the  $\epsilon$ -robustness of a given prediction (Mauá et al., 2018). Note that, since GeF<sup>+</sup>s have a tree structure and sum nodes with out-degree bounded by a constant, the time for computing  $\epsilon$ -robustness in GeF<sup>+</sup>s is linear in the input size (Correia & de Campos, 2019; Mauá et al., 2018).

We experiment with  $\epsilon$ -robustness in a selection of 12 datasets from the OpenML-CC18 benchmark<sup>2</sup> (Vanschoren et al., 2013). Once more, we use GeF<sup>+</sup>s composed of 30 trees with fully-factorised leaves and a 70/30 train-test split. In Figure 3, we defined a number of robustness thresholds and, for each of them, we computed the accuracy of the models over instances for which  $\epsilon$  was above and below the threshold. We clearly see there is a positive correlation; the higher the  $\epsilon$ -robustness of a prediction, the more likely it is to be correct. Obviously, the computation of robustness does not require knowing the true labels.

This concept of robustness has a clear interpretation. Given that for  $\epsilon = 0$  we have  $\mathcal{C}_{w,\epsilon} = \{w\}$  and for  $\epsilon = 1$  we have the whole simplex, we can interpret the value of  $\epsilon$  as the “percentage” of variation that we allow in the parameters for *each prediction*. That is in contrast to typical uncertainty measures where individual uncertainty values are hard to interpret in isolation (Mentch & Hooker, 2016; Shaker & Hüllermeier, 2020).

#### 5. Discussion and Further Experiments

We also trained GeF<sup>+</sup>s on the Mnist (LeCun et al., 2010) and Fashion-Mnist (Xiao et al., 2017) datasets to visually evaluate the samples with different  $\epsilon$ -robustness values. In Figure 4, we report test instances with lowest and highest

<sup>1</sup>Multinomial leaf nodes can be contaminated in the very same manner as sum nodes, while normal leaf nodes are contaminated in their means while keeping variance fixed (de Wit et al., 2019).

<sup>2</sup><https://www.openml.org/s/99/data>

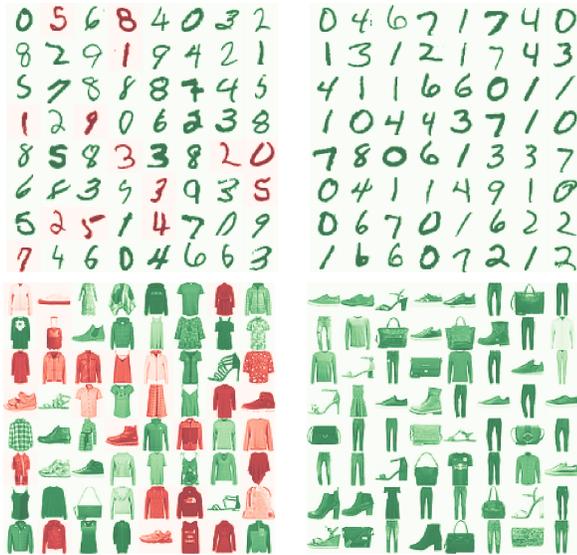


Figure 4. Samples from (Fashion-)Mnist datasets with lowest (left) and highest (right)  $\epsilon$ -robustness in the test set. Correctly and incorrectly classified examples are shown in green and red, respectively.

$\epsilon$ -robustness for each dataset. We see that samples for which the prediction is less robust are not only less likely to be correctly classified but also often contain irregular shapes and patterns, justifying the model’s uncertainty.

We emphasise outlier detection and robustness estimation are related but different notions, and GeF<sup>+</sup> effectively distinguishes them. Figure 5 shows a few of the most likely and unlikely (Fashion-)Mnist samples under the training data distribution. While samples are ordered by their marginal density  $p(\mathbf{x})$ , the background light is proportional to their  $\epsilon$ -robustness, with darker colours for larger  $\epsilon$ . We can clearly see how these measures differ as, for example, although the model deems 1s highly likely,  $\epsilon$ -robustness seems to vary with the shape/orientation of the trace.

Moreover, these two measures complement each other and allow us to better understand the underlying cause of the model’s uncertainty. Notably, for a consistent model—one that fits the true data generating distribution if given sufficient data—and a sample  $\mathbf{x}$  with high  $p(\mathbf{x})$  and low  $\epsilon$ -robustness, one may infer there is high aleatory uncertainty. A number 9 with an incomplete circle at the top is a good example of a pattern in handwritten digits that, albeit likely, is still hard to tell apart from a number 4. Conversely, an instance might be misshaped and hence unlikely, but still be associated to high robustness values. In that case, epistemic uncertainty is dominant, that is, the model has not been trained on similar examples and its high confidence estimate should not be trusted. Distinguishing the two types of uncertainty is not only fundamental to better understand the task at hand but also to establish the correct course of ac-

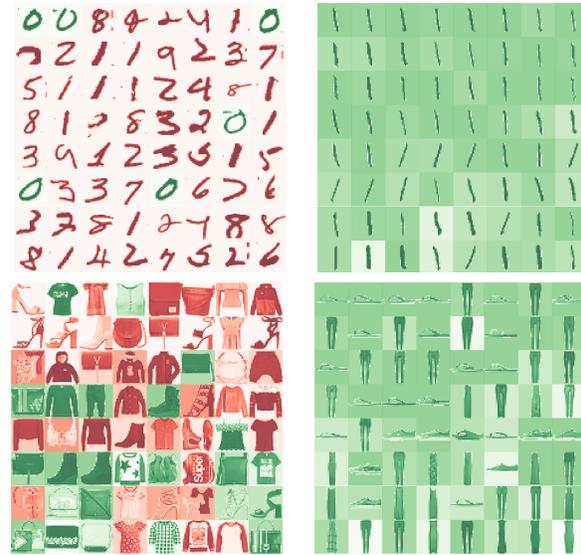


Figure 5. Samples from (Fashion-)Mnist datasets with lowest (left) and highest (right)  $p(\mathbf{x})$  in the test set. The background light is proportional to the  $\epsilon$ -robustness.

tion; namely, suspend judgement when faced with aleatory uncertainty or collect more data and possibly retrain the model in cases of epistemic uncertainty.

In all experiments<sup>3</sup>, the trees are made “deep”, that is, we keep splitting the feature space until each leaf cell contains either only samples of one class or a single sample. That means the average depth of our models is  $\Theta(\log n)$ , with  $n$  the number of samples in the training data (Louppe, 2014). Such deep trees make GeF<sup>+</sup>’s highly expressive, while the overall ensemble, by and large, avoids overfitting. It is also worth noticing that our models are learned as regular Random Forests, with bootstrapping and gini-impurity criterion, and afterwards converted to GeFs. Moreover, we use fully-factorised leaves,  $p(\mathbf{X}, Y) = p(X_1) \dots p(X_m)p(Y)$ , which are trivial to learn and compute but also achieve similar results as the original RF (identical predictions in each tree of the RF (Correia et al., 2020)).

## 6. Conclusion

While more experimentation is still needed, initial results indicate Generative Forests are a promising extension of Random Forests that effectively leverages the properties of Probabilistic Circuits to detect out-of-domain samples and estimate the robustness of its own predictions. We believe this is not only an important step towards more reliable machine learning but also a promising avenue for future research on deep hybrid discriminative-generative models.

<sup>3</sup>The source code will be available at the authors’ web-page.

## References

- Bradshaw, J., Matthews, A. G. d. G., and Ghahramani, Z. Adversarial examples, uncertainty, and transfer testing robustness in gaussian process hybrid deep networks. *arXiv preprint arXiv:1707.02476*, 2017.
- Correia, A. H. C. and de Campos, C. P. Towards scalable and robust sum-product networks. In *Scalable Uncertainty Management - 13th International Conference, SUM 2019, Proceedings*, Lecture Notes in Computer Science, pp. 409–422, Germany, 2019. Springer.
- Correia, A. H. C., Peharz, R., and de Campos, C. P. Joints in Random Forests. *arXiv e-prints*, art. arXiv:2006.14937, June 2020. URL <https://arxiv.org/abs/2006.14937>.
- Cortez, P., Cerdeira, A., Almeida, F., Matos, T., and Reis, J. Modeling wine preferences by data mining from physicochemical properties. *Decision Support Systems*, 47(4): 547–553, nov 2009.
- Darwiche, A. A differential approach to inference in Bayesian networks. *Journal of the ACM*, 50(3):280–305, 2003.
- de Wit, R., de Campos, C. P., Conaty, D., and del Rincon, J. M. Robustness in sum-product networks with continuous and categorical data. In De Bock, J., de Campos, C. P., de Cooman, G., Quaeghebeur, E., and Wheeler, G. (eds.), *Proceedings of the Eleventh International Symposium on Imprecise Probabilities: Theories and Applications*, volume 103 of *Proceedings of Machine Learning Research*, pp. 156–158. PMLR, 03–06 Jul 2019.
- Dietterich, T. G. Steps toward robust artificial intelligence. *AI Magazine*, 38(3):3–24, 2017.
- Freitas, A. A. Comprehensible classification models: a position paper. *ACM SIGKDD explorations newsletter*, 15(1):1–10, 2014.
- Guo, C., Pleiss, G., Sun, Y., and Weinberger, K. Q. On calibration of modern neural networks. In *Proceedings of the 34th International Conference on Machine Learning - Volume 70*, pp. 1321–1330. JMLR. org, 2017.
- Hendrycks, D. and Gimpel, K. A baseline for detecting misclassified and out-of-distribution examples in neural networks. *arXiv preprint arXiv:1610.02136*, 2016.
- Kaggle. Kaggle’s State of Data Science and Machine Learning 2019. Technical report, 2019. URL <https://www.kaggle.com/kaggle-survey-2019>.
- Kingma, D. P. and Welling, M. Auto-encoding variational Bayes. In *ICLR*, 2014. arXiv:1312.6114.
- LeCun, Y., Cortes, C., and Burges, C. Mnist handwritten digit database. 2010.
- Liang, S., Li, Y., and Srikant, R. Enhancing the reliability of out-of-distribution image detection in neural networks. *arXiv preprint arXiv:1706.02690*, 2017.
- Liu, F. T., Ting, K. M., and Zhou, Z.-H. Isolation forest. In *2008 Eighth IEEE International Conference on Data Mining*, pp. 413–422. IEEE, 2008.
- Louppe, G. *Understanding Random Forests: From Theory to Practice*. PhD thesis, University of Liege, 2014.
- Mauá, D. D., Conaty, D., Cozman, F. G., Poppenhaeger, K., and de Campos, C. P. Robustifying sum-product networks. *International Journal of Approximate Reasoning*, 101: 163–180, 2018.
- Mauá, D. D., Cozman, F. G., Conaty, D., and Campos, C. P. Credal sum-product networks. In Antonucci, A., Corani, G., Couso, I., and Destercke, S. (eds.), *Proceedings of the Tenth International Symposium on Imprecise Probability: Theories and Applications*, volume 62 of *Proceedings of Machine Learning Research*, pp. 205–216. PMLR, 10–14 Jul 2017.
- Mentch, L. and Hooker, G. Quantifying uncertainty in random forests via confidence intervals and hypothesis tests. *The Journal of Machine Learning Research*, 17(1): 841–881, 2016.
- Peharz, R., Tschitschek, S., Pernkopf, F., and Domingos, P. On theoretical properties of sum-product networks. In *Artificial Intelligence and Statistics*, pp. 744–752, 2015.
- Rezende, D. J. and Mohamed, S. Variational inference with normalizing flows. In *Proceedings of ICML*, pp. 1530–1538, 2015.
- Shaker, M. H. and Hüllermeier, E. Aleatoric and epistemic uncertainty with random forests. In *International Symposium on Intelligent Data Analysis*, pp. 444–456. Springer, 2020.
- Van den Broeck, G., Di Mauro, N., and Vergari, A. Tractable probabilistic models: Representations, algorithms, learning, and applications. <http://web.cs.ucla.edu/~guyvdb/slides/TPMTutorialUAI19.pdf>, 2019. Tutorial at UAI 2019.
- Vanschoren, J., van Rijn, J. N., Bischl, B., and Torgo, L. Openml: Networked science in machine learning. *SIGKDD Explorations*, 15(2):49–60, 2013. doi: 10.1145/2641190.2641198. URL <http://doi.acm.org/10.1145/2641190.2641198>.
- Xiao, H., Rasul, K., and Vollgraf, R. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*, 2017.