Learning Decision Trees When You Cannot Trust Your Labels

Łukasz Sztukiewicz[1,2] Jack H. Good[2] Artur W. Dubrawski[2]

[1] Faculty of Computing and Telecommunications, Poznan University of Technology, Poznan, Poland
[2] Auton Lab, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA, USA

Motivation

Context

- Real-world datasets are noisy
- Corrupted labels can hinder performance, increase size and prolong training times of models
- Relabeling is complex, costly and time-consuming

Objective

- To robustly learn a single decision tree without making assumptions about label noise

Methods

Kernel Density Decision Tree[1]

- Fuzzification natively represents uncertainty in the tree structure

- Kernels smooth and increase margin of decision boundaries

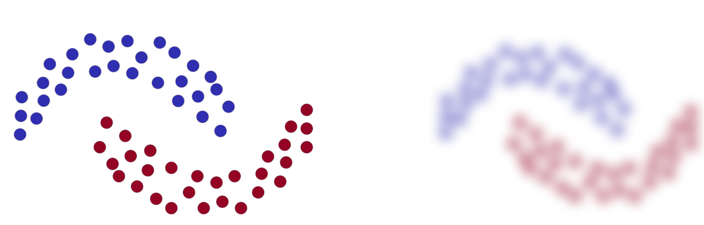

Robust Splitting Criterion[2]

- Takes into account unreliability of the data during decision tree induction

Feature Transformation Learning[3]

- Gradient-based optimization helps in obtaining small, performant trees

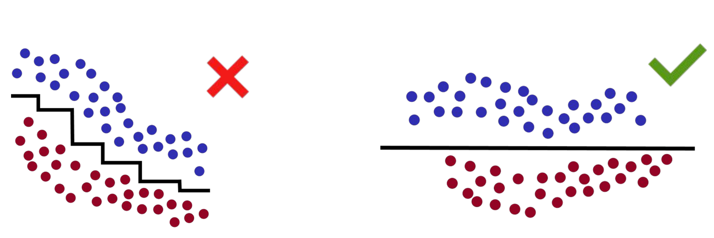

Demonstration

We compare the decision boundaries of selected methods trained on noisy data with 20% corrupted labels.

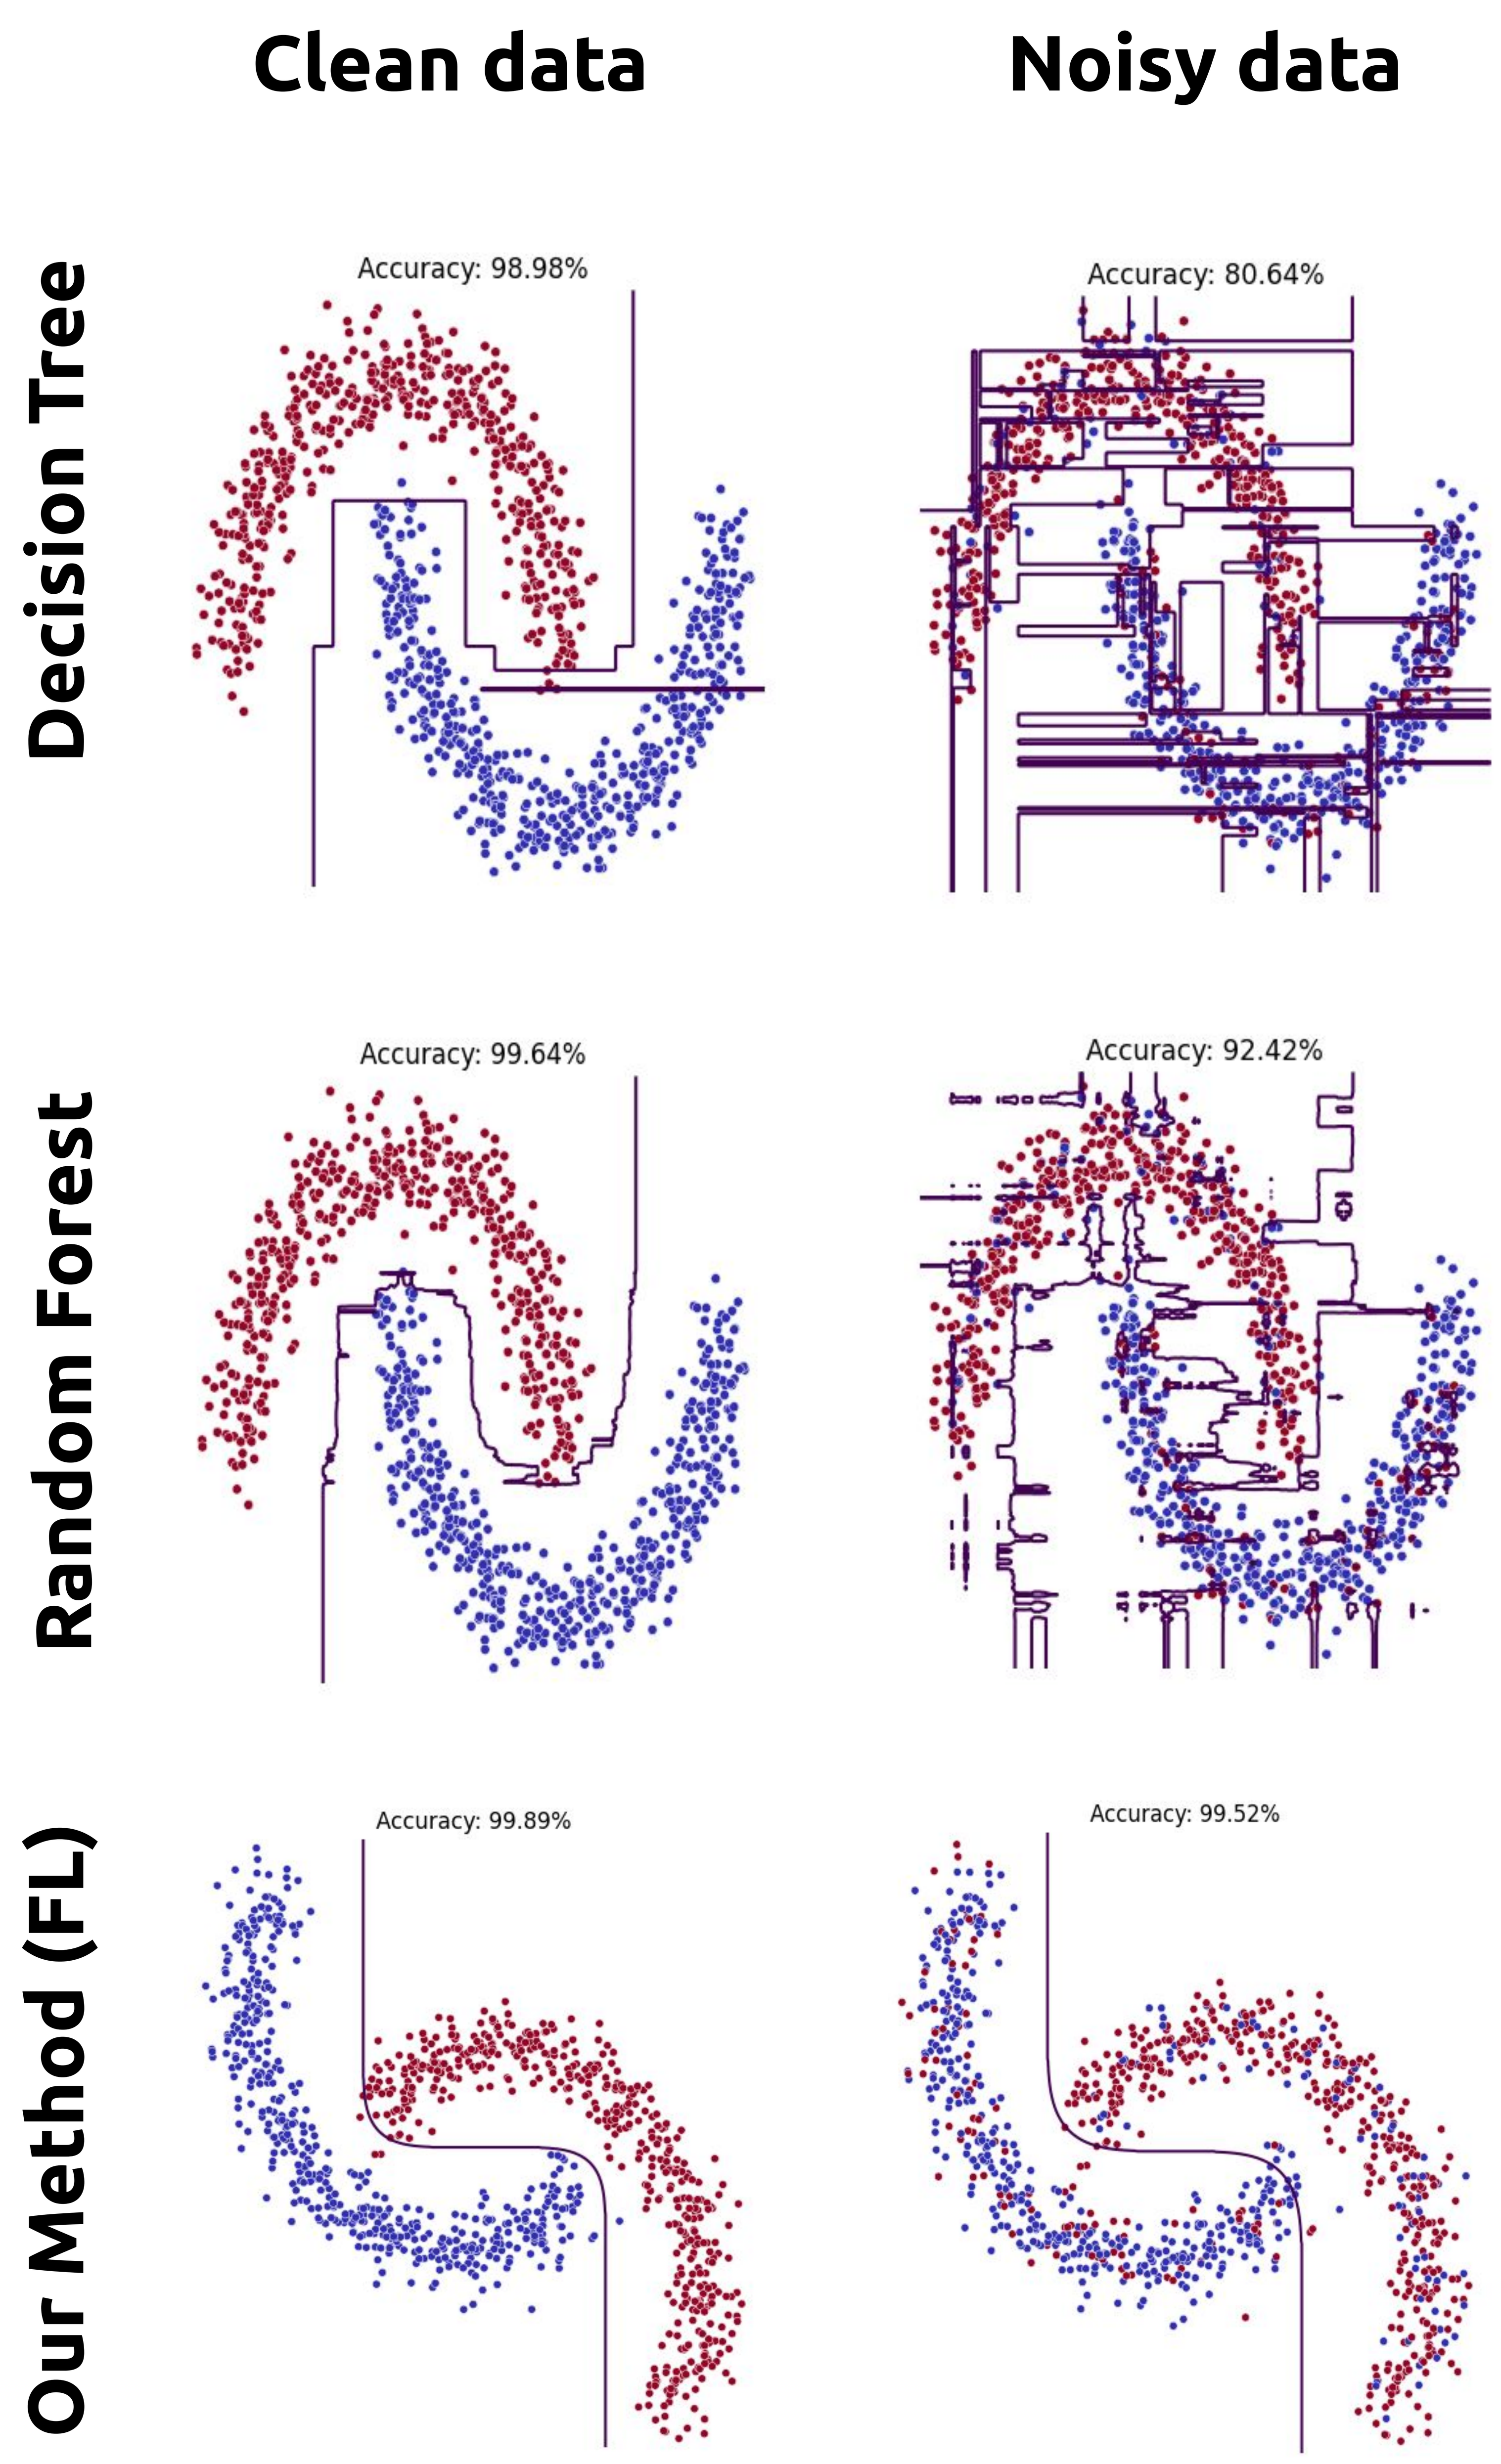

Robustness

We report Expected Loss of Accuracy (ELA) with respect to the clean dataset. The score is averaged across 10%, 20%, 30%, and 40% noise ratios scaled by a factor of 100 to improve readability. The best result is bolded, second best is underlined.

dataset	DecisionTree	GradientBoosting	RandomForest	Ours (Base)	Ours (FL)
balance-scale	53.8 ± 3.0	24.3 ± 2.2	31.8 ± 2.6	<u>19.5 ± 1.4</u>	**15.9 ± 0.9**
btsc	51.4 ± 2.8	37.0 ± 2.2	46.7 ± 1.8	<u>30.5 ± 2.0</u>	**29.6 ± 2.2**
iris	32.5 ± 4.2	24.0 ± 3.6	17.8 ± 3.3	**7.7 ± 2.5**	<u>8.9 ± 3.0</u>
kc2	37.5 ± 2.9	29.3 ± 2.8	31.8 ± 3.0	**22.2 ± 2.5**	<u>24.3 ± 2.9</u>
wdbc	31.7 ± 2.5	17.6 ± 2.7	<u>13.7 ± 2.8</u>	**13.3 ± 2.9**	16.2 ± 3.2

Performance

We demonstrate the models' relative performance to a standard Decision Tree algorithm measured by the Relative Weighted F1 score.

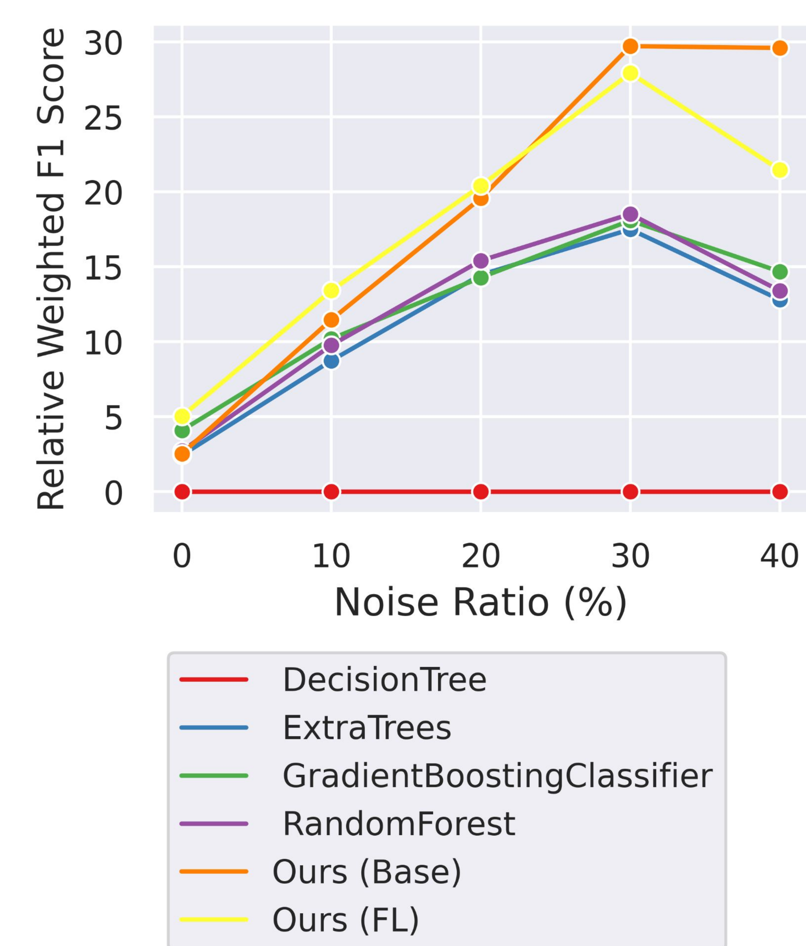

Acknowledgement

The authors would like to thank Rachel Burcin and John Dolan for making the RISS program possible and members of AutonLab for their continuous guidance and support. This work was partially funded by the Defense Advanced Research Projects Agency under award FA8750-17-2-0130.

References

[1] J. H. Good, K. Miller, and A. Dubrawski, "Kernel Density Decision Trees" *Proceedings of the AAAI Spring Symposium on AI Engineering*, 2022

[2] J. Abellán and S. Moral, "Building classification trees using the total uncertainty criterion" *International Journal of Intelligent Systems*, 2003

[3] J. H. Good, T. Kovach, K. Miller and A. Dubrawski, "Feature Learning for Interpretable, Performant Decision Trees" *Advances in Neural Information Processing Systems*, 2023