Individual Fairness in Optimal Decision Trees

AUTHOR

Chrysanthos Kindynis - c.kindynis@student.tudelft.nl

SUPERVISORS

Koos van der Linden
Dr. Emir Demirović

1) INTRODUCTION

- **Optimal** decisions trees guarantee maximized objective value within a given size limit.
- **Individual fairness:** ethical check on machine learning models which suggests that similar data should be classified the same [1]
- Issues with previous works: optimality not guaranteed [3] or scalability issues and not fair to the individuals [2].

2) RELATED WORK

- "**STreeD**": generalized Dynamic Programming framework [4]:
  - can construct optimal decision trees for any "separable" task
  - with a better performance than alternative approaches
- **Separability**: ability to estimate the objective value of a tree while using information found only within the current tree.

3) OBJECTIVE

> Define individual fairness as a "separable" task.
> Utilize "STreeD" framework to construct optimal decision trees for this task

Analyze the performance and scalability of this task across different data sets and depth limits.

4) RESEARCH METHOD

- Individual Fairness mathematical **formulation:**

$$\frac{\text{number of close pairs classified the same}}{\text{total number close pairs}}$$

- Define a separable approach:
  - We can gain information about Individual Fairness' **lower and upper bounds** of a tree in a separable way.
  - In a sub-tree, we find similar pairs that end up in that sub-tree. We then count the number of **similar pairs** that are **classified the same** and the number of similar pairs that are **classified differently**. The first indicates the lower bound, and the second on the upper bound of individual fairness of the sub-tree.
  - The leaf nodes provide the necessary information to the branch nodes, making the above estimation possible.
  - **Hard constraint** on I.F.: upper bound higher than threshold
  - Use these lower and upper bounds to compare solutions and argue about (pareto) **optimality**.

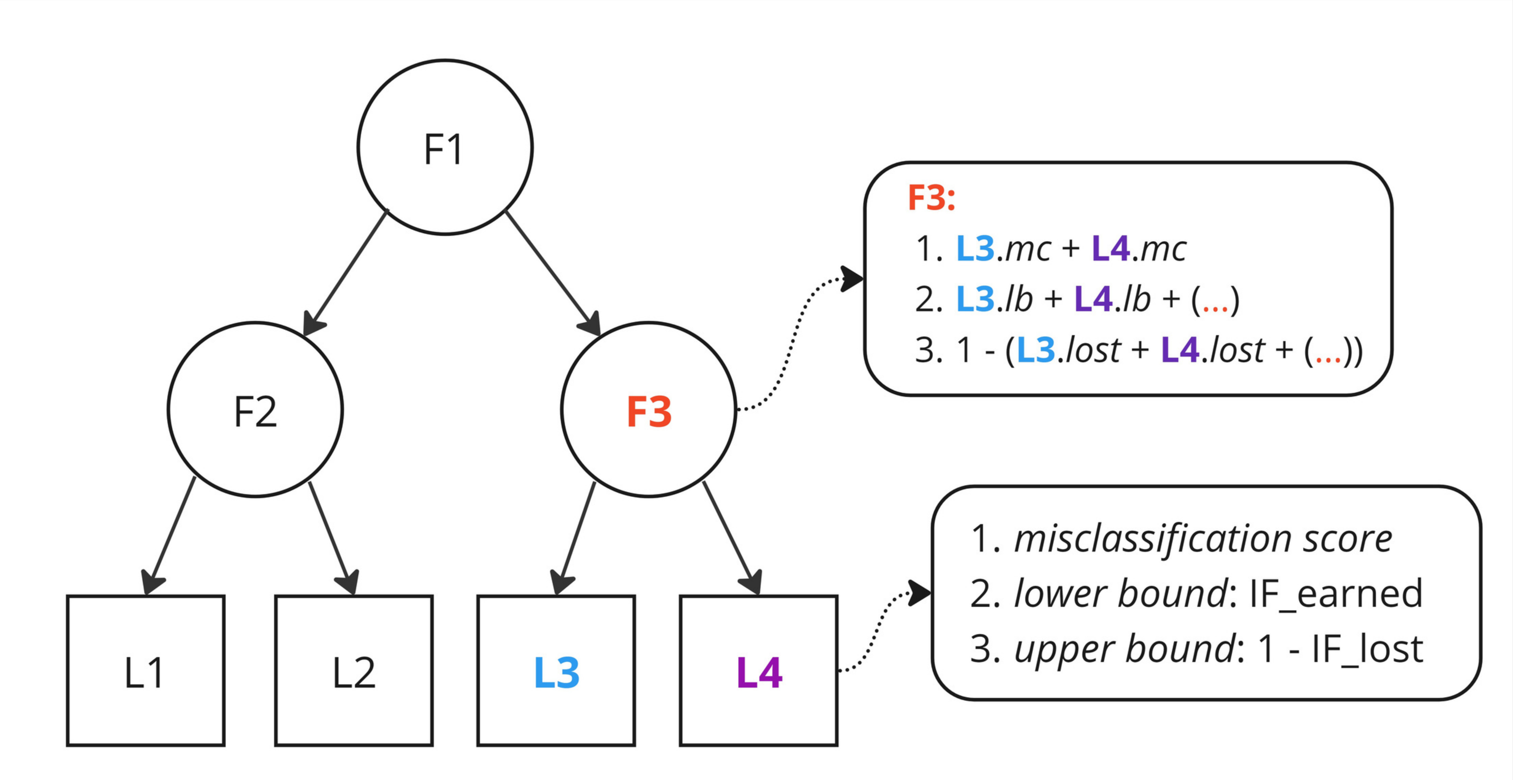

Figure 1: Example of a decision tree - illustrating the relation between a branching node and its two subtees

5) EXPERIMENTS ANALYSIS

We evaluate the runtime performance of our algorithm, across synthetic data sets. To present the effect of each of the three parameters, by varying the values of each parameter while keeping the other two fixed. We found that:
Depth 2: solved in milliseconds
Depth 3: reachable in many cases
Depth 4: reachable if few number of similar individuals in data set

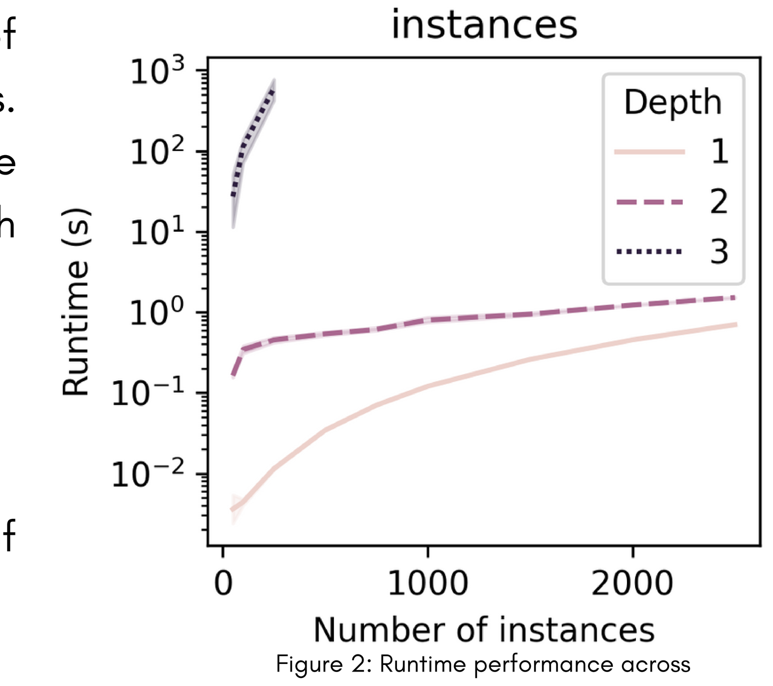

Figure 2: Runtime performance across

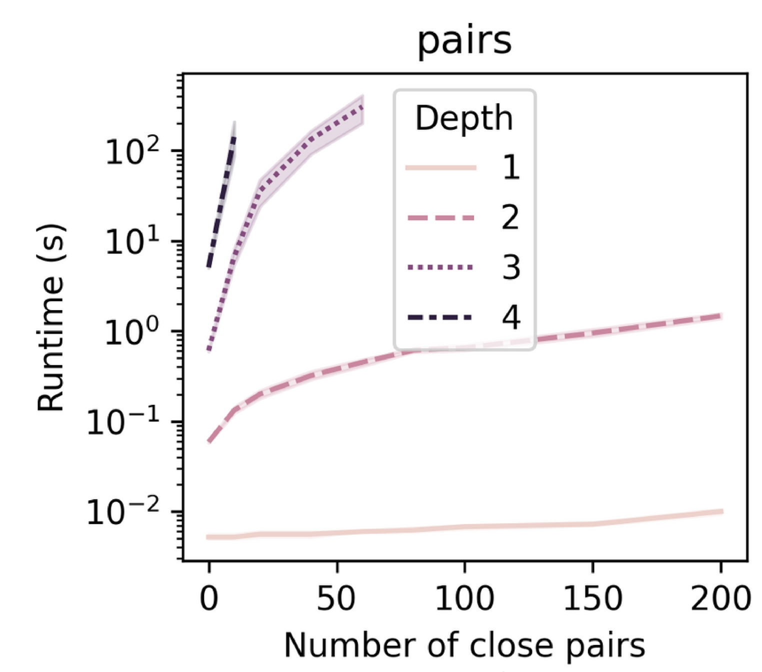

Figure 3: Runtime performance across different number of close pairs

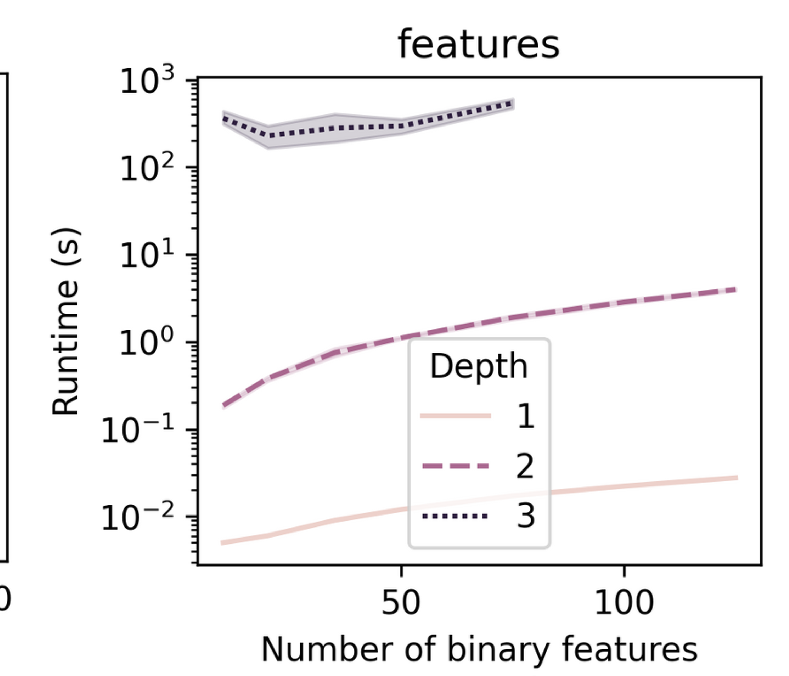

Figure 4: Runtime performance across different number of features

6) CONCLUSION

- The "STreeD" framework was successfully utilized to find optimal decision trees with the lowest misclassification score, and an individual fairness value about a threshold.
- Scalability results show a promising and competent performance against common approaches to optimal decision trees.

RELATED LITERATURE

1. Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. Fairness through awareness. In Proceedings of the 3rd Innovations in Theoretical Computer Science Conference, ITCS '12, page 214-226, New York, NY, USA, 2012
2. Sina Aghaei, Mohammad Javad Azizi, and Phebe Vayanos. Learning optimal and fair decision trees for non-discriminative decision-making. In Proceedings of the AAAI conference on artificial intelligence, volume 33, pages 1418-1426, 2019.
3. Francesco Ranzato, Caterina Urban, and Marco Zanella. Fairness-aware training of decision trees by abstract interpretation. In Proceedings of the 30th ACM International Conference on Information Knowledge Management, CIKM '21, page 1508-1517, New York, NY, USA, 2021. Association for Computing Machinery
4. Jacobus G. M. van der Linden, Mathijs M. de Weerdt, and Emir Demirovi´c. Optimal decision trees for separable objectives: Pushing the limits of dynamic programming, 2023.

TU Delft