

# TIMBERTREK: Exploring and Curating Sparse Decision Trees with Interactive Visualization

Zijie J. Wang<sup>1</sup> Chudi Zhong<sup>2</sup> Rui Xin<sup>2</sup> Takuya Takagi<sup>3</sup> Zhi Chen<sup>2</sup>  
Duen Horng Chau<sup>1</sup> Cynthia Rudin<sup>2</sup> Margo Seltzer<sup>4</sup>

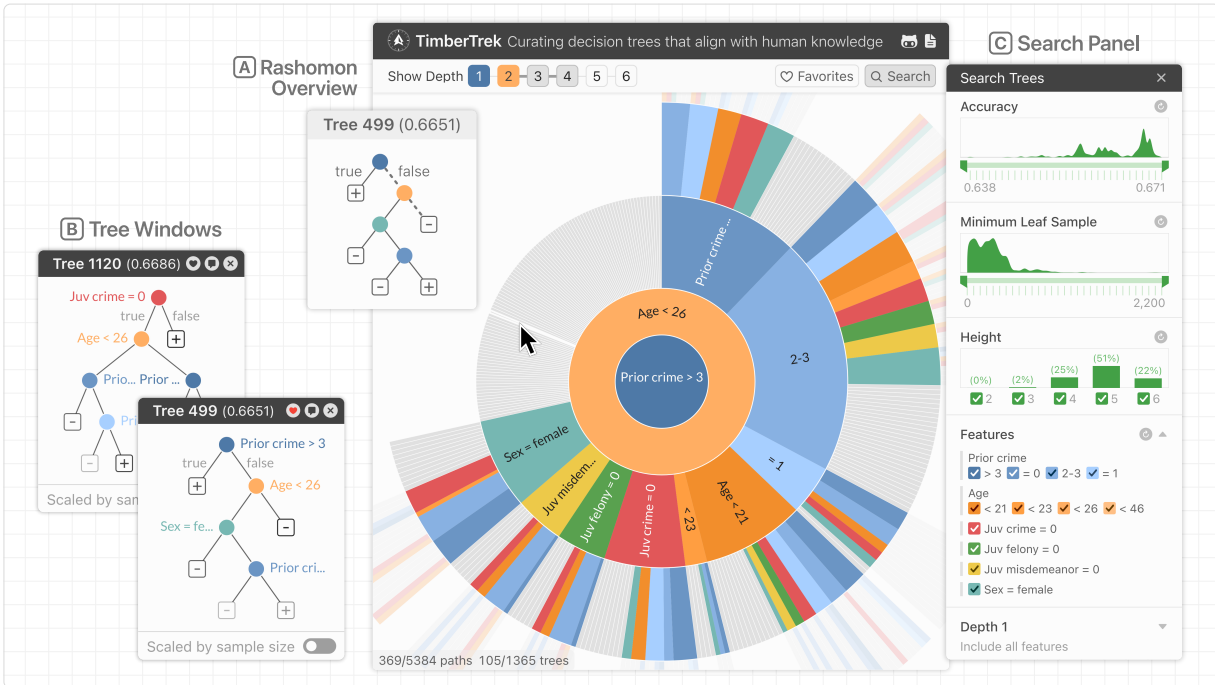


Fig. 1: TIMBERTREK empowers domain experts and data scientists to easily explore thousands of well-performing decision trees so they can find and collect those trees that best reflect their knowledge and values. Consider the task of predicting whether a criminal is likely to commit a crime in the next two years. (A) The *Rashomon Overview* visually summarizes all well-performing decision trees by organizing them based on their decision paths, enabling users to seamlessly transition across different model subsets and explore trees with similar prediction patterns. (B) Clicking a tree opens a repositionable *Tree Window* showing details of a decision tree: multiple windows allow users to compare several model candidates’ prediction patterns. (C) The *Search Panel* provides filtering tools, enabling users to quickly identify decision trees with desired properties, such as accuracy, robustness, simplicity, and used features.

## ABSTRACT

Given thousands of equally accurate machine learning (ML) models, how can users choose among them? A recent ML technique enables domain experts and data scientists to generate a complete *Rashomon set* for sparse decision trees—a huge set of almost-optimal interpretable ML models. To help ML practitioners identify models with desirable properties from this Rashomon set, we develop TIMBERTREK, the first interactive visualization system that summarizes thousands of sparse decision trees at scale. Two usage scenarios highlight how TIMBERTREK can empower users to easily explore, compare, and curate models that align with their domain knowledge and values. Our open-source tool runs directly in users’ computational notebooks and web browsers, lowering the barrier to creating more

responsible ML models. TIMBERTREK is available at the following public demo link: <https://poloclub.github.io/timbertrek>.

**Index Terms:** Human-centered computing—Visual Analytics

## 1 INTRODUCTION

It is essential to understand how machine learning (ML) models make predictions in high-stakes settings such as healthcare, finance, and criminal justice. Researchers have made great strides in developing interpretable models [e.g., 11, 30, 43] that perform competitively with state-of-the-art black-box models yet have transparent and simple structures [13, 44]. Some recent research focuses on *operationalizing* interpretability—leveraging an understanding of the domain to create more responsible and trustworthy ML systems [36, 46].

To help ML practitioners build trustworthy models, researchers have recently developed a technique to generate the full set of almost-optimal sparse decision trees [47]. This set of high-performing models is called the *Rashomon set* [18, 21], named after the *Rashomon effect* in statistics [8]. A Rashomon set of sparse decision trees can have thousands of inherently interpretable and almost-equally accurate models [47], providing opportunities for users to choose ones that best align with their knowledge and needs (e.g., fairness, monotonicity, simplicity) [18, 20]. However, the large size of the Rashomon set and diversity of models within it pose challenges

<sup>1</sup>Georgia Institute of Technology. {jayw|polo}@gatech.edu

<sup>2</sup>Duke University. {chudi.zhong|rui.xin926|zhi.chen1}@duke.edu, cynthia@cs.duke.edu

<sup>3</sup>Fujitsu Laboratories. takagi.takuya@fujitsu.com

<sup>4</sup>University of British Columbia. mseltzer@cs.ubc.ca

for users wishing to effectively explore the set and compare these accurate models [40]. To tackle this critical challenge, we **contribute**:

- **TIMBERTREK, the first interactive visualization tool** that empowers domain experts and data scientists to easily explore the Rashomon set of sparse decision trees and curate models with desired properties. Fig. 1 shows an example of TIMBERTREK in action, helping users explore 5384 decision paths from 1365 trees for recidivism risk assessment [28]. Advancing over prior visual analytics tools designed for interpretable ML models [e.g., 24, 46], our tool overcomes unique design challenges identified from a literature review of recent work in ML interpretability (§ 3).
- **Novel interactive system design** that leverages Sunburst [41] to summarize the entire Rashomon set at scale by organizing decision trees based on their decision paths (Fig. 1A). Through animation and *focus+context* [10] techniques (§ 4), our tool enables users to seamlessly traverse the full spectrum of abstraction levels: from the highest-level Sunburst overview (Fig. 2A), to intermediate levels of model subsets with similar prediction patterns (Fig. 2B), to the lowest-level node-link representation of individual trees (Fig. 2C). Two usage scenarios highlight how our tool can help users curate models with desired properties (§ 5).
- **An open-source<sup>1</sup> and web-based implementation** that broadens people’s access to trustworthy ML techniques (§ 4.5). We develop TIMBERTREK with modern web technologies so that anyone can access our tool directly in their web browsers and computational notebooks. For a demo video of TIMBERTREK, visit <https://youtu.be/3eGqTmsStJM>.

We hope our work helps democratize cutting-edge responsible ML techniques as well as inspires and informs future work in human-AI interaction and visual analytics for interpretable ML.

## 2 BACKGROUND & RELATED WORK

**Decision trees and Rashomon set.** Decision trees have been popular for more than half a century due to their accuracy and flexibility [31, 32]. These predictive models have a tree structure, where each branch node assesses a condition, and each leaf makes a prediction. With modern optimization techniques [25, 30], sparse decision trees—decision trees using a small set of features—have been gaining popularity in health care and criminal justice [39], as these models are not only accurate but also simple enough to be memorized by humans [40]. To help users explore diverse and accurate models and eventually find ones that they can trust, recent researchers have developed an algorithm to generate the whole Rashomon set of sparse decision trees [47]. Given a dataset with binary features and binary labels, two hyperparameters (sparsity penalty  $\lambda$  and loss tolerance  $\epsilon$ ), this algorithm will find all binary decision trees with at most  $\epsilon \times \ell$  loss on the training data, where  $\ell$  is the loss of the optimal tree. To characterize a Rashomon set, researchers have proposed visualization techniques to study the model construction process [26] and feature importance [20]. Different from these techniques, TIMBERTREK is the first interactive tool that summarizes a Rashomon set with varying levels of abstraction and enables users to curate good models.

**Visual analytics for model selection.** Iterating and selecting good models is a critical part of ML workflows [1, 2]. Visual analytics tools have shown great success in facilitating model selection [e.g., 6, 12, 16], as they enable the integration of domain knowledge in model development [14, 42]. For example, BEAMES [19] is an interactive tool that helps domain experts impose model constraints and search for linear models that meet specified constraints. Researchers have also developed visual analytics tools for interpreting tree-based models [e.g., 15, 29, 34, 50]. However, these tools focus on understanding single tree ensembles (i.e., random forest [7] and gradient-boosted trees [22]) instead of choosing from

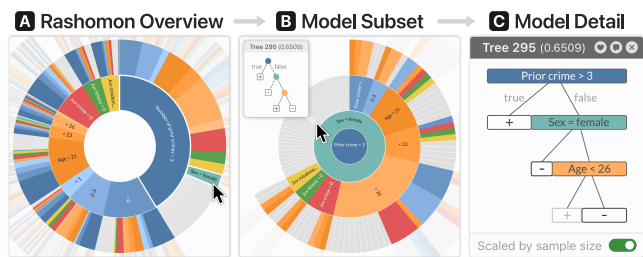


Fig. 2: TIMBERTREK’s tightly integrated views enable users to characterize and curate model candidates by seamlessly traversing across abstraction levels. (A) The *Rashomon Overview* summarizes the entire Rashomon set; (B) zoomable Sunburst enables users to focus on a subset of models with similar prediction patterns; and (C) the *Tree Window* presents the details of a selected decision tree.

a collection of standalone decision trees. Possibly closest in spirit to TIMBERTREK are TREEPOD [33] and a system designed by Padua et al. [37]. Both tools guide users to select satisfactory decision trees by tuning the parameters of decision tree algorithms. In contrast, our tool visualizes the complete Rashomon set of sparse decision trees with different levels of abstraction. For TIMBERTREK, every tree produced has both performance and interpretability guarantees.

## 3 DESIGN GOALS

Through synthesizing recent work in interpretable ML and ML workflows, we identify four design goals (G1–G4) for TIMBERTREK.

- G1. Visual Summary of the whole Rashomon set.** Depending on hyperparameters  $\lambda$  and  $\epsilon$ , the Rashomon set’s size can vary from hundreds to tens of thousands [47], posing challenges for users to explore models in this set [40]. Therefore, we aim to design scalable visualizations to summarize a large number of sparse decision trees to help users gain a better understanding of the landscape of the Rashomon set (§ 4.1).
- G2. Fluid transition between different levels of abstraction.** To curate models, users need to characterize all model candidates [40], identify important features [20, 21], and compare individual models [2]. Therefore, we would like to design a focus + context display to help users easily connect the Rashomon set landscape to individual models (§ 4.1). In addition, we would like to design query mechanisms to enable users to quickly pinpoint models with desirable properties (§ 4.2, § 4.4).
- G3. Model comparison.** Model selection often requires model comparison [2, 18]. In our case, all models are interpretable and have similar accuracies; thus, we would like TIMBERTREK to help users compare decision tree structures and prediction patterns (§ 4.3). As each model in the Rashomon set provides a *different* and *incomplete* explanation of the real-world phenomena [8], comparison can also help users gain insights into patterns within the full set of reasonable possibilities.
- G4. Fit into model development workflows.** Computational notebooks, such as Jupyter Notebook [27], have revolutionized how ML practitioners develop models [38]. To make model curation accessible and fit into the current workflows, we would like TIMBERTREK to work in both web browsers and computational notebooks. Finally, we open-source our implementation to support future design, research, and development of visual analytics tools for interpretable ML (§ 4.5).

## 4 VISUALIZATION INTERFACE OF TIMBERTREK

Following the design goals, TIMBERTREK (Fig. 1) tightly integrates four components: the *Rashomon Overview* providing a hierarchical overview of all decision trees in a Rashomon set (§ 4.1), the *Search Panel* enabling users to find trees with desirable properties (§ 4.2),

<sup>1</sup>TIMBERTREK code: <https://github.com/poloclub/timbertrek>

*Tree Windows* showing details of selected decision trees (§ 4.3), and the *Favorite Panel* documenting curated trees (§ 4.4).

#### 4.1 Summarizing the Whole Rashomon Set

A Rashomon set can contain thousands of sparse binary decision trees. Consider criminal recidivism prediction as an example [28]; each branch node in a decision tree assesses a feature condition (e.g., **juvenile crime = 0**), and a leaf node makes a prediction (e.g., the subject is likely to reoffend in two years). Decision rules [13, 42] can be represented as a path from the root to a leaf. For example, the left-most path of Tree 1071 (Fig. 3A) represents “IF (**juvenile crime = 0**) AND (**prior crime > 3**) THEN (the subject is likely to re-offend).” To help users characterize the Rashomon set, we can organize decision trees based on their decision rules.

**Sunburst overview.** We first construct a trie of decision rules by extracting decision paths from all leaves of all trees in the Rashomon set (G1). A trie leaf is a decision tree that contains a decision rule matching the leaf’s ancestors. We use Sunburst [41] to visualize the trie in the *Rashomon Overview* (Fig. 1A). The Sunburst consists of  $h$  concentric rings, where  $h$  is the height of the decision rule trie (i.e., the maximal tree height in the Rashomon set). A ring corresponds to a level in the trie, and it is segmented into annular sectors, where each sector represents a split condition used in decision rules. Each sector cumulatively starts a subtree, where the children are conditions used in the next level of the trie. A sector’s size is proportional to the number of its descendant decision trees.

**Color encoding.** We use the Hue-chroma-luminance (HCL) [49] color space to represent the sector’s color. Decision tree methods often binarize continuous and categorical features into multiple split conditions covering various ranges [25, 30]. Therefore, we use different hue values to represent different features (e.g., **prior crime**, **age**, **sex**), and luminance values to represent different ranges of the same feature (e.g., **prior > 3**, **prior = 0**, **prior = 1**). We use gray to encode **leaf sectors**: a leaf of the trie indicates the end of a decision rule, and each leaf links to a decision tree that contains this decision rule. Finally, we group sectors using the same feature and sort sectors based on the number of trees within the Rashomon set.

**Model exploration.** The *Rashomon Overview* not only summarizes all decision trees in the Rashomon set (G1) but also highlights *feature importance* through sectors’ size, as important features are often used by many accurate models thus yield larger sectors [20]. To help users further explore decision trees with similar prediction patterns, the *Rashomon Overview* provides smooth transitions between different levels of Rashomon set abstraction (G2). When users click a sector, the Sunburst’s root switches to the selected split condition and only displays its descendants (Fig. 2B)—helping users focus on an interesting subtree. When users hover over a **leaf sector**, our tool shows the corresponding decision tree as a link-node diagram [35] (shown on the right), where the selected decision rule is highlighted as animated dash lines and binary outputs as  $\oplus$  and  $\ominus$ . Users can control how many levels to display in Sunburst via the depth panel, whose colors match clicked sectors (Fig. 1A-top left).

#### 4.2 Searching Models with Desirable Properties

The *Rashomon Overview* provides a bird’s-eye view, enabling users to follow different decision paths to explore different subsets of decision trees in the Rashomon Set. To allow users to quickly pinpoint trees with desirable properties (G2), the *Search Panel* (Fig. 1C) offers a suite of filtering panels to control what trees to display in the *Rashomon Overview*. For example, with the accuracy and minimum sample leaf size sliders, users can focus on trees with desired accuracy and robustness. Similarly, users can use checkboxes to filter models by tree height and the use of specific features (Fig. 4A).




Fig. 3: TIBBERTREK helps users inspect and curate decision trees. (A1) The *Tree Window* visualizes a decision tree’s prediction pattern via a node-link diagram; (A2) a funnel-like complementary view scales tree nodes by their sample sizes. (B) The *Favorite Panel* allows users to keep track of bookmarked models with curation documentation.

#### 4.3 Comparing Individual Decision Trees

The combination of *Rashomon Overview* and *Search Panel* provides users with a searchable directory of decision trees. After identifying interesting model candidates, users often need to compare them to identify ones suitable for practical use (G3). When a user clicks a **leaf sector**, a *Tree Window* appears, visualizing the corresponding decision tree as a node-link diagram (Fig. 3-A1). We use opacity to encode a leaf node’s accuracy, allowing users to quickly inspect a tree’s purity and prediction confidence [9]. This window is repositionable through dragging, and there can be multiple windows open at once—a user can create *Tree Windows* for all interesting model candidates and easily compare their structures and prediction patterns side by side. When comparing decision trees, users are also interested in the sample sizes of nodes in addition to the tree structure [42]. Therefore, when a user toggles the sample-size switch, the *Tree Window* transitions the node width to represent the percentage of training samples that fall into each split condition (branch node) or prediction (leaf node) (Fig. 3-A2). This novel funnel-like node-link diagram can help users quickly identify important nodes and evaluate model robustness via the node sample sizes [9].

#### 4.4 Curating Trustworthy Models

The goal of TIBBERTREK is to empower users to explore and curate decision trees for practical use. Once a user has identified a satisfactory model, they can click the heart button  $\heartsuit$  in the *Tree Window* to bookmark a decision tree. Bookmarked trees appear in the *Favorite Panel* (Fig. 3B). To help users track their reasons and contexts for choosing a particular decision tree, the *Favorite Panel* allows users to attach a comment to each bookmarked model. Alternatively, users can also click the comment button  $\text{comment}$  to add comments directly in the *Tree Window* (Fig. 4C). These comments allow users to continue their model curation in the future (G4), help ML auditors audit models before deployment [3, 46], and help improve ML transparency regarding the model development process [23]. Finally, users can click the save button  $\text{save}$  to export bookmarked decision trees with curation comments; TIBBERTREK’s companion package provides an API to load and deploy saved models.

#### 4.5 Accessible, Open-source Implementation

TIBBERTREK is a web-based interactive visualization tool built with *D3.js* [5]: users can access our tool with any web browser or directly in computational notebooks. To promote the accessibility of our tool, and strongly align with our VIS community’s open science practice, we have released TIBBERTREK on the Python Package Index (PyPI),<sup>2</sup> so that users can easily install our tool and integrate it into their ML development workflows (G4). We have also open sourced our implementation: future researchers can quickly adapt our design to other forms of model curation.

<sup>2</sup>PyPI repository: <https://pypi.org/project/timbertrek/>

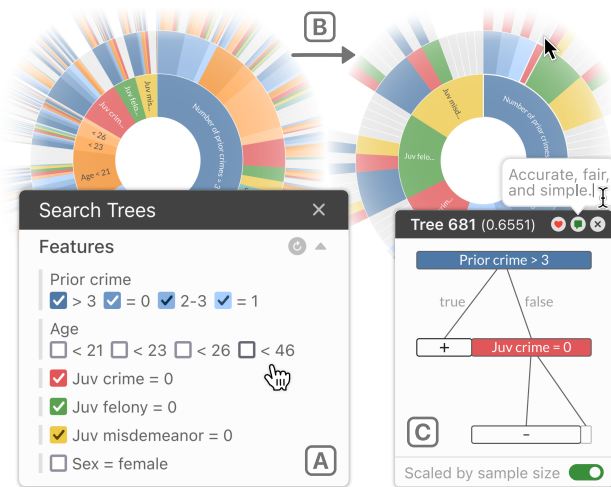


Fig. 4: With TIMBERTREK, users can easily search for models with desirable properties. Here, in the example of criminal recidivism assessment, (A) the *Search Panel* enables users to query models that do not use protected attributes (e.g., age and sex); (B) the *Rashomon Overview* animates to only display models meeting the query criteria. (C) The *Tree Window* provides details of selected model query result.

## 5 USAGE SCENARIOS

We present two hypothetical scenarios with real datasets to demonstrate how TIMBERTREK can potentially help data scientists and domain experts gain a better understanding of the model-world relationship and curate more trustworthy models. We generate a Rashomon set with 1,365 trees ( $\lambda = 0.01$ ,  $\epsilon = 1.05$ ) for the COMPAS recidivism dataset [28] (§ 5.1), and a Rashomon set with 911 trees ( $\lambda = 0.15$ ,  $\epsilon = 1.015$ ) for the Car Evaluation dataset [4] (§ 5.2).

### 5.1 Discovering Fair Rules for Recidivism Assessment

Mei is a data scientist who develops transparent and fair ML models to help inform judicial bail decisions. To explore the relationship between diverse variables and the risk of criminal recidivism, she has generated a Rashomon set of sparse decision trees on past criminal recidivism data (we use COMPAS [28] to illustrate this scenario). This dataset includes defendants’ demographic information and criminal history; the outcome variable is binary—indicating whether a defendant is likely to reoffend in the next two years.

To understand the similarities and differences between all 1,365 about-equally accurate models from the Rashomon set, Mei loads all models into TIMBERTREK. The *Rashomon Overview* visualizes all decision trees (Fig. 2A). Inspecting the Sunburst’s first ring, Mei quickly realizes **prior crime** may be the most important feature to assess recidivism risk. This is because the root is the most powerful node in a decision tree [9], and *more than half* of models in the Rashomon set choose **prior crime** as their roots. Mei opens the *Search Panel* (Fig. 1C) and searches for trees that do not use **prior crime** at any depth—there is no tree meeting this criterion. Similarly, Mei hypothesizes **sex** is the least important feature due to the small sizes of its sectors (Fig. 2A). Mei’s hypotheses regarding feature importance match previous study results on COMPAS [20].

Mei believes **prior crime** is indeed an informative feature for recidivism prediction, but she notices many accurate models use sensitive features such as **age** and **sex**. Making bail decisions with these models could be problematic [17]. Therefore, Mei decides to find accurate models that do not use any sensitive features. To do that, she uses the *Search Panel* to query decision trees without any **age** or **sex** nodes (Fig. 4A). TIMBERTREK finds 33 trees that meet this criterion (Fig. 4B). After inspecting these trees by hovering over the *leaf sectors*, Mei finds her favorite model—Tree 681 that

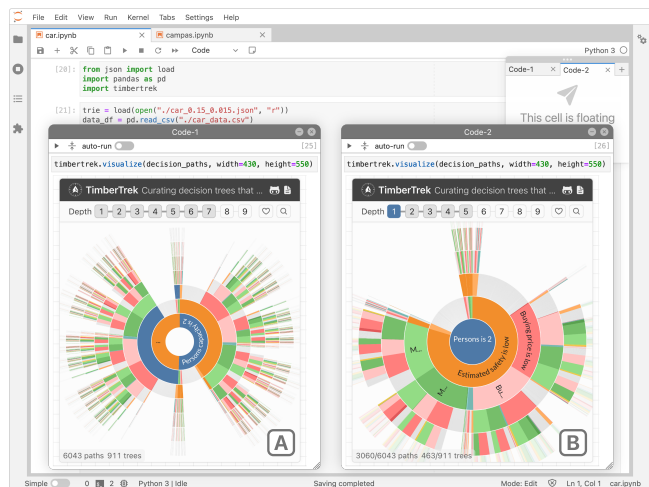


Fig. 5: Our tool works in computational notebooks—commonly used in ML development workflows. With sticky cells, users can create multiple TIMBERTREK instances and compare different Rashomon subsets side by side. Take Car Evaluation as an example: (A) reveals all trees in the Rashomon set use either **person** or **safety** as root, and (B) provides an additional view to focus on trees with a **person** root.

does not use sensitive features, yields a high accuracy, and is simple enough to be memorized (Fig. 4C). Mei adds it to her curation list and writes a short comment to document why she chooses this tree.

### 5.2 Curating Trustworthy Models for Insurance Quote

Robaire is an ML developer building models to help his company make automobile insurance quotes. For accountability to end-users, his company requires all models to be transparent and easy to explain. Therefore, Robaire decides to curate trustworthy models from the Rashomon set of sparse decision trees. We use the Car Evaluation dataset [4] to illustrate this scenario; the task is to predict whether a vehicle’s value is acceptable with typical vehicle features. After loading TIMBERTREK directly in JupyterLab (Fig. 5A), Robaire quickly notices all 911 accurate trees (median accuracy is 0.92)—with an even distribution—use either **person capacity** or **safety score** as their first split. Curious, Robaire decides to compare model structures between two subsets of trees with different roots. He opens a second TIMBERTREK instance in a notebook sticky cell [45], where he clicks the **blue sector** in the first ring to focus on trees with a **person capacity** root (Fig. 5B). He repeats the same process on the other cell to choose a different root. Comparing two subsets side by side, Robaire finds almost all trees use these two features in their first two splits and use diverse combinations of other features in further levels. Therefore, to avoid potential overfitting and keep the model simple enough for end-users to understand, Robaire eventually chooses two trees that use only **person capacity** and **safety score** nodes.

## 6 DISCUSSION & CONCLUSION

We present TIMBERTREK, the first visualization system that summarizes the entire Rashomon set and empowers ML practitioners to explore, compare, and curate models with desired properties. Our current prototype does not scale to multi-class classification trees, decision trees with many features (limitation of our color encoding), or trees with many levels. However, interpretability requires decision trees to limit the number of features and levels [40, 48]. In addition, our design principles such as focus + context, model comparison and query are generalizable to other model types. Future researchers can use our tool as a research instrument to probe how users would select ML models when many models are approximately equally accurate. We hope our work will inspire future research and development of tools that can empower users to interpret and trust ML technologies.

## ACKNOWLEDGMENTS

We thank anonymous reviewers for their valuable feedback. This work was supported in part by a J.P. Morgan PhD Fellowship, NSF grants IIS-1563816, CNS-1704701, NIH/NIDA grant DA054994-01, DARPA GARD, Fujitsu, gifts from Intel, NVIDIA, Bosch, Google.

## REFERENCES

- [1] S. Amershi, A. Begel, C. Bird, R. DeLine, H. Gall, E. Kamar, N. Nagappan, B. Nushi, and T. Zimmermann. Software Engineering for Machine Learning: A Case Study. *ICSE*, May 2019.
- [2] S. Amershi, M. Cakmak, W. B. Knox, and T. Kulesza. Power to the People: The Role of Humans in Interactive Machine Learning. *AI Magazine*, 35, Dec. 2014.
- [3] S. Amershi, K. Inkpen, J. Teevan, R. Kikin-Gil, E. Horvitz, D. Weld, M. Vorvoreanu, A. Fournay, B. Nushi, P. Collisson, J. Suh, S. Iqbal, and P. N. Bennett. Guidelines for Human-AI Interaction. *CHI*, 2019.
- [4] M. Bohanec and V. Rajkovic. Knowledge Acquisition and Explanation for Multi-Attribute Decision Making. 1988.
- [5] M. Bostock, V. Ogievetsky, and J. Heer. D<sup>3</sup> Data-Driven Documents. *IEEE TVCG*, 17, Dec. 2011.
- [6] L. Bradel, C. North, L. House, and S. Leman. Multi-Model Semantic Interaction for Text Analytics. In *2014 IEEE Conference on Visual Analytics Science and Technology (VAST)*, Oct. 2014.
- [7] L. Breiman. Random forests. *Machine learning*, 45, 2001.
- [8] L. Breiman. Statistical Modeling: The Two Cultures (with comments and a rejoinder by the author). *Statistical Science*, 16, Aug. 2001.
- [9] L. Breiman, J. H. Friedman, R. A. Olshen, and C. J. Stone. *Classification and Regression Trees*. First edition, 1984.
- [10] S. K. Card, J. D. Mackinlay, and B. Shneiderman. *Readings in Information Visualization: Using Vision to Think*. The Morgan Kaufmann Series in Interactive Technologies. 1999.
- [11] R. Caruana, Y. Lou, J. Gehrke, P. Koch, M. Sturm, and N. Elhadad. Intelligible Models for HealthCare: Predicting Pneumonia Risk and Hospital 30-day Readmission. *KDD*, Aug. 2015.
- [12] D. Cashman, A. Perer, R. Chang, and H. Strobel. Ablate, Variate, and Contemplate: Visual Analytics for Discovering Neural Architectures. *IEEE TVCG*, 26, Jan. 2020.
- [13] C.-H. Chang, S. Tan, B. Lengerich, A. Goldenberg, and R. Caruana. How Interpretable and Trustworthy are GAMs? *KDD*, Aug. 2021.
- [14] A. Chatzimpampas, R. M. Martins, I. Jusufi, K. Kucher, F. Rossi, and A. Kerren. The State of the Art in Enhancing Trust in Machine Learning Models with the Use of Visualizations. *Computer Graphics Forum*, 39, June 2020.
- [15] A. Chatzimpampas, R. M. Martins, and A. Kerren. VisRuler: Visual Analytics for Extracting Decision Rules from Bagged and Boosted Decision Trees. *arXiv:2112.00334*, Apr. 2022.
- [16] A. Chatzimpampas, R. M. Martins, K. Kucher, and A. Kerren. StackGenVis: Alignment of Data, Algorithms, and Models for Stacking Ensemble Learning Using Performance Metrics. *IEEE TVCG*, 2021.
- [17] A. Chouldechova. Fair Prediction with Disparate Impact: A Study of Bias in Recidivism Prediction Instruments. *Big Data*, 5, 2017.
- [18] A. Coston, A. Rambachan, and A. Chouldechova. Characterizing Fairness Over the Set of Good Models Under Selective Labels. 2021.
- [19] S. Das, D. Cashman, R. Chang, and A. Endert. BEAMES: Interactive Multimodel Steering, Selection, and Inspection for Regression Tasks. *IEEE Computer Graphics and Applications*, 39, Sept. 2019.
- [20] J. Dong and C. Rudin. Exploring the Cloud of Variable Importance for the Set of All Good Models. *Nature Machine Intelligence*, 2, 2020.
- [21] A. Fisher, C. Rudin, and F. Dominici. All models are wrong, but many are useful: Learning a variable’s importance by studying an entire class of prediction models simultaneously. *Journal of Machine Learning Research*, 20, 2019.
- [22] J. H. Friedman. Greedy Function Approximation: A Gradient Boosting Machine. *The Annals of Statistics*, 29, 2001.
- [23] L. Hancox-Li. Robustness in Machine Learning Explanations: Does It Matter? *ACM FAccT*, Jan. 2020.
- [24] F. Hohman, A. Head, R. Caruana, R. DeLine, and S. M. Drucker. Gamut: A Design Probe to Understand How Data Scientists Understand Machine Learning Models. *CHI*, May 2019.
- [25] X. Hu, C. Rudin, and M. Seltzer. Optimal Sparse Decision Trees. In *Neural Information Processing Systems*, volume 32, 2019.
- [26] N. Kissel and L. Mentch. Forward Stability and Model Path Selection. *arXiv:2103.03462*, Mar. 2021.
- [27] T. Kluyver and others. Jupyter Notebooks - a Publishing Format for Reproducible Computational Workflows. *ELPUB*, 2016.
- [28] J. Larson, S. Mattu, L. Kirchner, and J. Angwin. How We Analyzed the COMPAS Recidivism Algorithm. *ProPublica*, 9, 2016.
- [29] Y. Li, T. Fujiwara, Y. K. Choi, K. K. Kim, and K.-L. Ma. A Visual Analytics System for Multi-Model Comparison on Clinical Data Predictions. *Visual Informatics*, 4, June 2020.
- [30] J. Lin, C. Zhong, D. Hu, C. Rudin, and M. Seltzer. Generalized and Scalable Optimal Sparse Decision Trees. In *International Conference on Machine Learning*, July 2020.
- [31] W.-Y. Loh. Fifty Years of Classification and Regression Trees: Fifty Years of Classification and Regression Trees. *International Statistical Review*, 82, Dec. 2014.
- [32] J. N. Morgan and J. A. Sonquist. Problems in the Analysis of Survey Data, and a Proposal. *J Am Stat Assoc*, 58, June 1963.
- [33] T. Muhlbacher, L. Linhardt, T. Moller, and H. Piringer. TreePOD: Sensitivity-Aware Selection of Pareto-Optimal Decision Trees. *IEEE TVCG*, 24, Jan. 2018.
- [34] M. P. Neto and F. V. Paulovich. Explainable Matrix - Visualization for Global and Local Interpretability of Random Forest Classification Ensembles. *IEEE TVCG*, 27, Feb. 2021.
- [35] T. Nguyen, T. Ho, and H. Shimodaira. A Visualization Tool for Interactive Learning of Large Decision Trees. *IEEE ICTAI*, 2000.
- [36] H. Nori, R. Caruana, Z. Bu, J. H. Shen, and J. Kulkarni. Accuracy, interpretability, and differential privacy via explainable boosting. In *International Conference on Machine Learning*, 2021.
- [37] L. Padua, H. Schulze, K. Matković, and C. Delrieux. Interactive Exploration of Parameter Space in Data Mining: Comprehending the Predictive Quality of Large Decision Tree Collections. *Computers & Graphics*, 41, June 2014.
- [38] J. M. Perkel. Reactive, Reproducible, Collaborative: Computational Notebooks Evolve. *Nature*, 593, May 2021.
- [39] C. Rudin. Stop Explaining Black Box Machine Learning Models for High Stakes Decisions and Use Interpretable Models Instead. *Nature Machine Intelligence*, 1, May 2019.
- [40] C. Rudin, C. Chen, Z. Chen, H. Huang, L. Semenova, and C. Zhong. Interpretable Machine Learning: Fundamental Principles and 10 Grand Challenges. *Statistics Surveys*, 16, Jan. 2022.
- [41] J. Stasko, R. Catrambone, M. Guzdial, and K. McDonald. An Evaluation of Space-Filling Information Visualizations for Depicting Hierarchical Structures. *International Journal of Human-Computer Studies*, 2000.
- [42] D. Streeb, Y. Metz, U. Schlegel, B. Schneider, M. El-Assady, H. Neth, M. Chen, and D. Keim. Task-Based Visual Interactive Modeling: Decision Trees and Rule-Based Classifiers. *IEEE TVCG*, 2021.
- [43] B. Ustun and C. Rudin. Learning Optimized Risk Scores. *Journal of Machine Learning Research*, 20, 2019.
- [44] C. Wang, B. Han, B. Patel, and C. Rudin. In Pursuit of Interpretable, Fair and Accurate Machine Learning for Criminal Recidivism Prediction. *Journal of Quantitative Criminology*, Mar. 2022.
- [45] Z. J. Wang, K. Dai, and W. K. Edwards. StickyLand: Breaking the Linear Presentation of Computational Notebooks. *CHI EA*, 2022.
- [46] Z. J. Wang, A. Kale, H. Nori, P. Stella, M. Nunnally, D. H. Chau, M. Vorvoreanu, J. W. Vaughan, and R. Caruana. GAM Changer: Editing Generalized Additive Models with Interactive Visualization. *arXiv:2112.03245*, Dec. 2021.
- [47] R. Xin, T. Takagi, C. Zhong, C. Zhi, M. Seltzer, and C. Rudin. Efficient Construction of the Rashomon Set. 2022.
- [48] J. Yuan, B. Barr, K. Overton, and E. Bertini. Visual Exploration of Machine Learning Model Behavior with Hierarchical Surrogate Rule Sets. *arxiv:2201.07724*, Jan. 2022.
- [49] A. Zeileis, K. Hornik, and P. Murrell. Escaping RGBland: Selecting colors for statistical graphics. *Comput Stat Data Anal*, 53, 2009.
- [50] X. Zhao, Y. Wu, D. L. Lee, and W. Cui. iForest: Interpreting Random Forests via Visual Analytics. *IEEE TVCG*, 25, Jan. 2019.