DOI 10.5162/iCCC2026/P32

Visualizing Score Contributions in BBQ-Trees with Applications to Medical Screening Data

Alexander Stahl
FG Datenbank- und Informationssysteme
Brandenburgische Technische Universität
Cottbus, Germany
stahlale@b-tu.de

***Abstract*—Interpretable decisionmaking is essential in medical and health-related applications. BBQ-Trees (BBQTs) combine decision-tree structure with quantum-inspired logic, with advantages to accuracy and interpretability. However, their evaluation scheme can make per-sample explanations challenging. We propose a visualization method that maps subtree contributions to colors in order to reveal which decision paths drive an individual prediction. Using publicly available medical datasets as illustrative examples, we show that the method makes BBQT predictions more transparent without altering their predictive behavior.**

***Index Terms*—explainable AI, decision trees, visualization, clinical decision support, medical screening data**

I. Introduction

Machine learning (ML) is becoming increasingly popular in high-stakes domains such as medicine, where predictive models assist clinicians in diagnostic and therapeutic decision making. As summarized by Rahmani et al. [1], current developments show both the potential of ML in clinical practice but also the challenges associated with integrating complex models into safety-critical applications. While black-box models have gained widespread popularity and are increasingly deployed in many settings, recent works argue for the need for explainable models in high-stakes applications. As one of such applications, clinical decision support systems require models that are not only accurate but also interpretable. This is not only to foster trust among health practitioners and maintain accountability, but also to comply with regulatory standards.

Classical decision trees (DT) satisfy many of these requirements due to their transparent structure. However, they often struggle to represent curved or otherwise non-linear decision boundaries without becoming overly deep or fragmented, which can reduce interpretability [2], [3]. One alternative model, the *BBQ-Tree (BBQT)*, generalizes the DT and the so-called *Quantum Logic Decision Tree (QLDT)*. The model's QLDT components provide gradual, *quantum-inspired* truth values that allow the model to capture trend information and curved decision boundaries, whereas the DT components retain the capabilities of making classical axis-parallel attribute splits for accuracy boosts [4].

A inherited limitation of BBQTs is that all branches contribute to the final evaluation, due to the gradual truth semantics. Unlike classical DTs, where irrelevant branches are simply not traversed, BBQTs evaluate all decision paths. For bigger trees, this global evaluation can complicate the interpretability of the model and per-sample explanations, as it becomes more difficult to identify which parts of the tree were most influental for a particular prediction. In this work, we argue that local interpretability for BBQTs can be recovered by focusing on how the model evaluates a specific input sample. Instead of explaining all global interactions, we analyze the subtree scores that arise during the BBQT evaluation and visualize their relative contributions. Our contribution is a coloring method that maps subtree evaluation scores to a heat-based color scheme, which can highlight which parts of the BBQT contribute most to an individual prediction. This produces a locally interpretable representation of BBQT decisions, making the model's internal reasoning accessible to end users such as clinicians.

II. Related Work

Interpretable machine learning has produced a wide range of approaches for explaining model behavior, both for tree-based methods and for general black-box models. A substantial body of post-hoc explanation techniques aims to provide model-agnostic interpretability, while other work focuses on enhancing the transparency of tree structures themselves. Several tools and methods improve the visual interpretability of classical DTs. Common techniques include shading nodes by class distribution or prediction confidence, highlighting the path taken by an individual sample (e.g., `dtreeviz` [5]), and using leaf-level heatmaps to represent class purity, as implemented in `treeheatr` [6].

More general post-hoc methods such as LIME and SHAP approximate local feature importance by perturbation or additive attribution models. For neural architectures, techniques such as Grad-CAM and attention visualizations provide insight into salient regions or internal weight patterns. For an overview of existing methods, see Molnar [7]. While these approaches can highlight important features, they do not necessarily reveal the underlying *reasoning structure* of the model. In logic-based models, particularly tree-based ones, interpretability benefits from the ability to inspect the explicit decision structure, even when the tree becomes large or deep. This motivates staying within the family of tree-based models rather than relying solely on model-agnostic post-hoc explanations.

DOI 10.5162/iCCC2026/P32

Our method extends the idea of path- or node-level visualization by mapping subtree evaluation scores to color intensities. This provides a sample-specific explanation of BBQT predictions, which bridges the gap between the interpretability of logic-base models and post-hoc explanation techniques.

III. Background: BBQT and QLDT

This section reviews essential components of BBQTs necessary for understanding the research problem we address and the method we introduce. The training process is omitted, as it is not central to this study. All material in this section is drawn from our earlier work [4] and is restated here for completeness.

BBQTs extend QLDTs by incorporating classical Boolean split types into the quantum-inspired framework. The core evaluation principle of QLDTs remains largely intact. Each internal node represents a logical condition formulated in a continuous quantum-inspired logic, where attribute values are interpreted as degrees of truth indicating how strongly a feature expresses a certain trend (e.g., being "high").

BBQTs generalize both QLDTs and classical DTs by allowing three types of splits:

- Q-splits: conditions with gradual truth values
- B-splits: classical DT splits with learned thresholds
- $B_{1/2}$-splits: B-Splits with a fixed threshold at 0.5

This combination preserves several interpretability advantages of QLDTs. Q-splits capture meaningful trends (e.g., "low" hemoglobin or "elevated" glucose in a patient) rather than depending solely on arbitrary numeric cutoffs. As the underlying logic framework limits the usage of attribute per path, BBQTs often remain relatively shallow. In addition, the learning process tends to place the most informative and interpretable conditions near the top of the tree, which supports the interpretation of results.

For the evaluation of a given sample o, each node in a BBQT computes a truth value $[c_{ij}]^o \in [0, 1]$ that quantifies how well the sample satisfies the corresponding condition. A BBQT can be viewed as a collection of paths from the root to each leaf. Within a path, the node values are combined by conjunction, implemented as the product of the node truth values. Class scores are obtained by summing the scores of all paths $\in \mathcal{P}_1$, i. e. paths leading to leaves of the class 1. Then, the score S_1 for class 1 is

$$S_1(o) = [e]^o = \sum_{i \in \mathcal{P}_1} \prod_{j \in \text{path}_i} [c_{ij}]^o.$$

The final prediction is made by comparing S_1 to a learned threshold $\tau \in [0, 1]$.

Unlike classical DTs, which consider only the single path selected by hard splits, BBQTs evaluate *all* paths simultaneously. As mentioned above, this global aggregation can make it more challenging to identify which parts of the tree contributed most to an individual prediction.

IV. Subtree-Based Coloring for Interpretability

This section introduces our subtree-based coloring method, beginning with established decision-tree practices and extending them to BBQ-Trees.

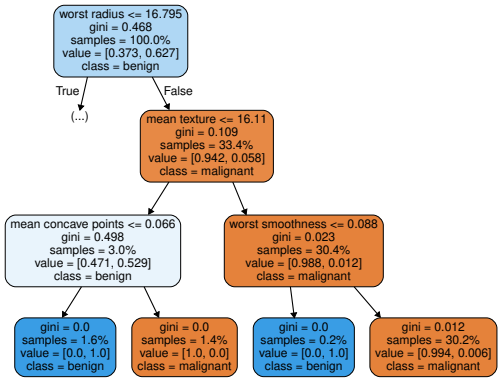

Fig. 1. Right branch of a DT trained on the *Breast Cancer Wisconsin* dataset. Nodes are colored by predicted class and class distribution.

Coloring DTs is a common practice in many machine learning libraries. In fact, standard plotting utilities typically produce color-coded trees by default. For example, scikit-learn's [8] visualization (Figure 1) assigns colors based on the predicted class and the class distribution within each node. In the illustrated example, a DT was trained to predict breast cancer outcomes using the *Breast Cancer Wisconsin* dataset [9]. For compactness, only the right branch of the tree is shown. Orange denotes *malignant* and blue *benign*, while saturation reflects the class distribution in each node. Such visualizations are useful for understanding the global behavior of the tree, but they do not provide a detailed explanation for an individual prediction. Classical DTs, however, operate strictly by traversing a single root-to-leaf path, and this path can be highlighted to indicate which decisions were made for a specific sample. We adapt this idea to BBQTs, where the evaluation is more complex and cannot be represented by a single traversed path.

To visualize BBQTs for a given sample, we compute the score contribution of each subtree to the final class score. These contributions are then mapped to colors using a perceptually uniform colormap (e.g., `viridis`) to ensure that stronger contributions receive visually stronger emphasis.

Intermediate node scores follow the BBQT evaluation scheme. For a node associated with an attribute a, the model evaluates two subtrees: one conjoined with a and one with $\neg a$. The node score is their disjunction,

$$[a]^o * [c_{\text{subtree}_1}]^o + (1 - [a]^o) * [c_{\text{subtree}_2}]^o,$$

where $[\cdot]^o$ denotes the evaluation of a condition for sample o and c_{subtree} represents the condition encoded by a subtree. The same expression contains the edge scores: the first term of the sum corresponds to the edge representing a, and the second to the edge representing $\neg a$. B-split nodes are treated analogously, but using their respective evaluation rule.

Mapping these scores to color intensities produces a visual representation in which users can directly see which parts of the BBQT "push" the prediction higher or lower for an individual patient. Figure 2 presents an exemplary colored BBQT unrelated to any particular dataset. It shows the typical

DOI 10.5162/iCCC2026/P32

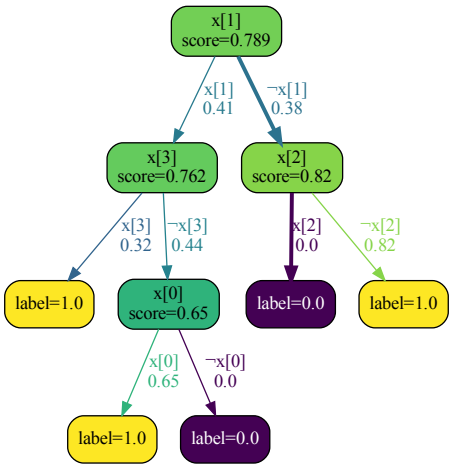

Fig. 2. Example BBQT with heatmap-based coloring scheme

structure and the evaluation for a sample predicted as class 1. Warm colors represent high subtree scores, while cold colors indicate low scores. Here, the root reaches a score of 0.789, reflecting the strong influence of the paths to 1-leaves.

V. Applicability of BBQ-Trees in Medical Settings

In this section, we examine the applicability of BBQTs in medical decision scenarios. We outline the data requirements necessary to use the model effectively and illustrate its behavior on commonly used medical datasets. This motivates improving BBQT interpretability for practitioners who rely on model outputs without being ML experts.

Like other classifiers for tabular data, BBQTs can be applied to structured patient data in settings where binary decisions are required (e.g., disease vs. no disease). Then, each patient corresponds to one data point in a multidimensional feature space. For reliable use of BBQTs, several practical criteria must be fulfilled:

- Balanced datasets for training
- Attribute values scaled to the unit interval $[0, 1]$
- Binary labels
- Non-correlated attributes (Not strict in practice.)

In general, BBQTs perform well when the target class can be represented by distinct regions near the corners of the feature space hypercube.

Table I presents results on four standard medical datasets: *Breast Cancer Wisconsin*, *Pima Indians Diabetes*, *Heart Disease Cleveland*, and *Chronic Kidney Disease* [9]–[12]. While selected metadata is included in the table, additional details are available in the respective descriptions in the cited sources. Across these datasets, accuracy ranges from 79% to 92.5%. For completeness, ROC AUC and other metrics are also reported.

All experiments were conducted using a grid search to identify suitable hyperparameters and evaluated using 4-fold cross validation. The train/test split ratio was set to 0.20. Standard scikit-learn tools were used alongside our custom BBQT implementation.

VI. Demonstration of Subtree-Based Coloring

Figure 3 shows a BBQT trained on the *Autism Screening Data for Toddlers* dataset [13]. This example illustrates clearly how the proposed coloring method supports the interpretation of BBQT predictions. The model is relatively shallow, with depth $d = 4$ plus leaves, yet it already contains $\sum_{i=0}^{d-1} 2^i = 15$ internal nodes. Because a BBQT evaluates all paths globally, even such a small tree might become difficult for non-experts to interpret when presented without visual guidance.

Using the heatmap-based coloring scheme introduced above, nodes and edges with high contribution scores appear in warm colors (yellow), whereas low scores appear in cold colors (purple). In this dataset, most features correspond to binary (yes/no) questions, which results in contributions concentrated at the extremes. This leads to a clear visual separation between influential and non-influential parts of the tree.

The displayed coloring corresponds to the evaluation of a particular sample that the model classified as *ASD positive*. This can be seen immediately at the root node: its color is bright yellow, indicating a high final score (with 1 representing ASD and 0 representing no ASD). Three leaves in the tree predict the positive class, and therefore represent active minterms in the internal CQQL aggregation. By examining which of these leaves contributes most strongly to the root score, we can identify the driving factor behind the decision.

Only the rightmost leaf exhibits a consistently warm-colored path from leaf to root, indicating that this path provided the dominant contribution to the final prediction. The logical condition along this path is:

$$x[5] \wedge x[9] \wedge x[4] \wedge x[10] \geq 0.031.$$

According to the dataset documentation in [13], these features represent:

- $x[5]$: Does the child follow where an adult is looking?
- $x[9]$: Does the child stare at nothing with no apparent purpose?
- $x[4]$: Does the child point?
- $x[10] \geq 0.031$: Normalized age of the child should be greater or equal to 0.031 (~ 2.48 months)

Thus, these features contributed most significantly to the model's positive ASD prediction for this sample. As the dataset's class labels were produced through an automated screening tool called ASDTests [14], they do not possess the level of diagnostic reliability as validated assessments. However, as the purpose of this experiment is to examine the technical applicability of BBQTs rather than to make substantive clinical claims, we hold that this limitation does not compromise our methodological insights.

DOI 10.5162/iCCC2026/P32

TABLE I
EVALUATION RESULTS OF BBQT ON STANDARD MEDICAL DATASETS

Dataset	n_samples	n_features	best_cv_score	accuracy	f1	roc_auc
Breast Cancer Wisconsin	569	30	0.9639	0.9035	0.9197	0.9666
Pima Indians Diabetes	768	8	0.7945	0.7077	0.6666	0.7901
Heart Disease Cleveland	297	13	0.8163	0.7833	0.7346	0.8214
Chronic Kidney Disease	400	24	0.9466	0.925	0.9361	0.94

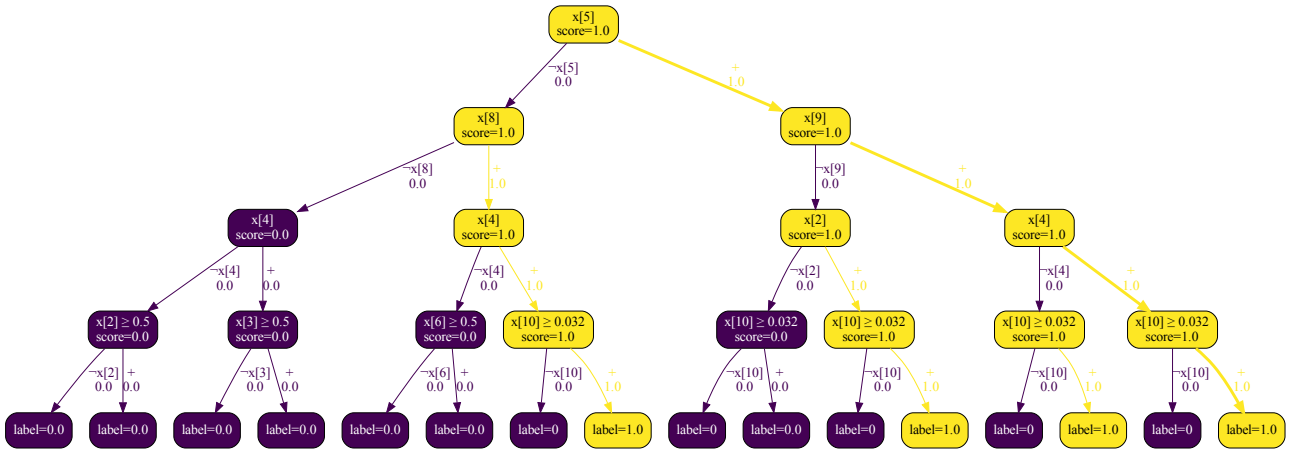

Fig. 3. BBQT for autism diagnosis with 89.63% accuracy, using highlighting for a single prediction

VII. CONCLUSION

This study shows a methodological contribution aimed at improving the interpretability of BBQ-Trees, a class of classification models. While BBQTs can represent complex, non-linear decision boundaries, their evaluation scheme can make it difficult to identify important factors of the decision. We address this by introducing a color-based visualization that maps subtree evaluation scores to color values, for example in a heatmap color scheme. This enables insights into the influences of each decision path in the trained model.

Our experiments on popular medical datasets demonstrate that BBQTs can be applied in settings where binary decisions must be made from structured patient data. Notably, the results are not intended as clinical findings. Instead, they are meant to illustrate the model's applicability and show how the proposed visualization can support interpretability in domains, such as medicine, where understanding decisions is important. The method is general and can be applied to other areas in which transparent model behavior is required.

REFERENCES

[1] A. M. Rahmani, E. Yousefpoor, M. S. Yousefpoor, Z. Mehmood, A. Haider, M. Hosseinzadeh, and R. Ali Naqvi, "Machine learning (ML) in medicine: Review, applications, and challenges," *Mathematics*, vol. 9, no. 22, 2021.

[2] L. Breiman, J. H. Friedman, R. A. Olshen, and C. J. Stone, *Classification and Regression Trees*. New York: CRC Press, 1984.

[3] S. Russell and P. Norvig, *Artificial Intelligence: A Modern Approach, Global Edition*. London: Pearson Education, 2021.

[4] A. Stahl and I. Schmitt, "BBQ-Tree: A unified classifier and regressor combining boolean and quantum logic decisions," *Information Systems*, vol. 136, 2026. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0306437925001188

[5] T. Parrt, D. Johnson, and J. Fairbanks, "dtreeviz: Decision tree visualization library," https://github.com/parrt/dtreeviz, 2019, version accessed: 2025-12-01.

[6] N. Yang and K. Baum, "treeheatr: Heatmaps for decision trees," https://cran.r-project.org/package=treeheatr, 2020, version accessed: 2025-12-01.

[7] C. Molnar, *Interpretable Machine Learning*, 2022, accessed: 2025-11-26.

[8] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, "Scikit-learn: Machine learning in Python," *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.

[9] M. O. S. N. Wolberg, William and W. Street, "Breast Cancer Wisconsin (Diagnostic)," UCI Machine Learning Repository, 1993, DOI: https://doi.org/10.24432/C5DW2B.

[10] U. M. L. Repository, "Pima indians diabetes database," https://www.kaggle.com/datasets/uciml/pima-indians-diabetes-database, 2016, accessed: 2025-11-29.

[11] S. W. P. M. Janosi, Andras and R. Detrano, "Heart Disease," UCI Machine Learning Repository, 1989, DOI: https://doi.org/10.24432/C52P4X.

[12] S. P. Rubini, L. and P. Eswaran, "Chronic Kidney Disease," UCI Machine Learning Repository, 2015, DOI: https://doi.org/10.24432/C5G020.

[13] F. Abdeljaoued, "Autism screening for toddlers," https://www.kaggle.com/datasets/fabdelja/autism-screening-for-toddlers, 2020, accessed: 2025-11-28.

[14] "Asdtests," https://asdtests.com, accessed: 2025-11-29.