ACT: AGENTIC CLASSIFICATION TREE

Anonymous authors
Paper under double-blind review

ABSTRACT

When used in high-stakes settings, AI systems are expected to produce decisions that are transparent, interpretable, and auditable—a requirement increasingly expected by regulations. Decision trees such as CART provide clear and verifiable rules, but they are restricted to structured tabular data and cannot operate directly on unstructured inputs such as text. In practice, large language models (LLMs) are widely used for such data, yet prompting strategies such as chain-of-thought or prompt optimization still rely on free-form reasoning, limiting their ability to ensure trustworthy behaviors. We present the *Agentic Classification Tree* (ACT), which extends decision-tree methodology to unstructured inputs by formulating each split as a natural-language question, refined through impurity-based evaluation and LLM feedback via TextGrad. Experiments on text benchmarks show that ACT matches or surpasses prompting-based baselines while producing transparent and interpretable decision paths.

Unstructured Data: Tuberculosis diagnosis based on symptoms descriptions

Patient 1: "My main symptom is a runny nose with increased sweating, pain, fever, skin rashes, nasal congestion, sore throat, muscle pain, loss of appetite, and chills. The heavy pain is in my left temple, and I have pink rashes on the back of my neck."

Patient 2: "I am coughing blood with pain (heavy in my left temple), skin lesions, nasal congestion, extreme fatigue affecting activities, diffuse muscle pain, loss of appetite, chills, a heavy pain on my left temple, and pink rash on the right side of my neck."

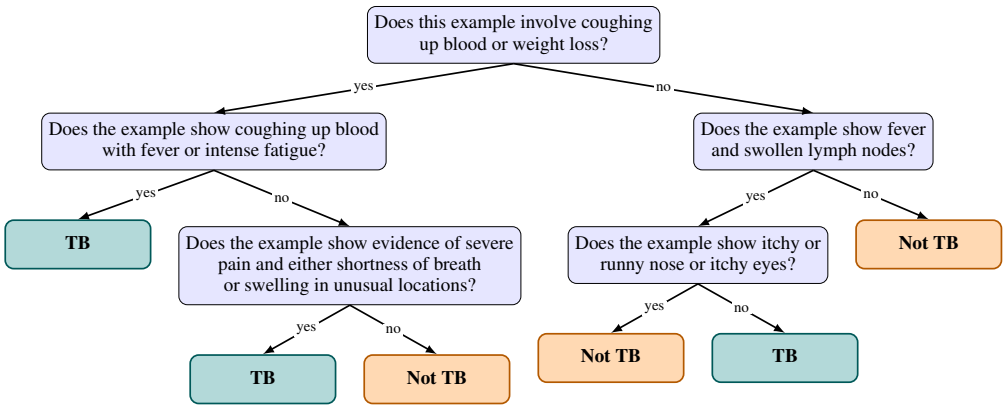

Figure 1: Example ACT decision tree for tuberculosis diagnosis using unstructured, free-text patient descriptions. A tree is automatically learned, with each node containing a binary natural language question, autonomously discovered via recursive prompt refinement to maximize label separation at each split. At inference, these questions are answered by a large language model (LLM) from the root node to the leaves of the tree. The final classification (TB or Not TB) corresponds to the majority label of training examples described by each leaf.

1 INTRODUCTION

As AI systems become increasingly integrated into high-stakes domains such as healthcare, education, legal decision-making, and finance, the need for transparency, interpretability, and auditability in AI decision-making has intensified. These requirements are grounded not only in practical considerations, but also in legal obligations: recent governance frameworks and regulations such as the EU AI Act[1], the OECD AI Principles[2], and the NIST AI Risk Management Framework[3] emphasize that AI-based decisions in high-stakes scenarios must be explainable and subject to human oversight. In this context, there is a growing need for models whose decision processes can be inspected, verified, and understood—not only by technical experts but also by stakeholders, auditors, and regulators.

Historically, complex decision-making tasks in high-stakes domains—such as medical diagnosis, fraud detection, or credit underwriting—have been addressed on interpretable AI systems such as expert systems (Jackson, 1986; Shortliffe, 1986; Leonard, 1993; Talebzadeh et al., 1995) and decision trees (Breiman et al., 1984). These models allow users to trace each decision through a sequence of explicit, human-understandable rules, directly supporting the requirements set out by regulatory frameworks. However, their applicability is largely limited to structured inputs such as tabular data, and are not directly applicable to unstructured data such as natural language or images.

Recent advancements in large language models (LLMs) have substantially enhanced the capacity for semantic understanding and reasoning within natural language processing (Brown et al., 2020; Achiam et al., 2023). Building on these models, agentic systems (Shinn et al., 2023; Yao et al., 2023; Schick et al., 2023) have been proposed to address complex tasks involving reasoning over diverse types of data, and are seen as a very promising direction for future AI systems (Chen et al., 2023; Wang et al., 2024).

Yet, despite their impressive capabilities, LLMs still suffer from several drawbacks—such as hallucinations, inconsistencies, unreliability, and difficulty to audit (Kryściński et al., 2019; Ji et al., 2023a; Tamkin et al., 2021). These issues constitute major obstacles to building reliable AI systems, limiting their adoption in sensitive or regulated domains. A growing body of work seeks to address these weaknesses by guiding LLM behavior toward more structured or self-consistent reasoning. Chain-of-Thought prompting Wei et al. (2022), for example, encourages multi-step reasoning by inserting explicit intermediate steps into prompts, improving consistency and sometimes reducing hallucinations. Other techniques—such as self-reflection Shinn et al. (2023); Madaan et al. (2023) and prompt optimization Zhou et al. (2022); Ji et al. (2023b); Yuksekgonul et al. (2024); Renze & Guven (2024)—introduce agentic feedback loops to iteratively refine model outputs. While these methods improve reliability in practice, they still rely on free-form text generation and implicit reasoning patterns that remain difficult to verify, audit, or formally constrain.

In this paper, we introduce a new type of classifier designed to combine the semantic reasoning capabilities of large language models (LLMs) with the transparency of decision trees. The idea is to follow a divide-and-conquer paradigm, allowing to reduce hallucination and reasoning errors by successive decomposition of the problem following a hierarchical decision logic.

We therefore propose to structure LLM and agentic reasoning to address complex classification tasks through verifiable decision patterns on *unstructured data* such as images or texts, in the form of an agentic decision tree. Adapting the traditional Decision tree algorithms such as CART Breiman et al. (1984) and C4.5 Salzberg (1994), we propose ACT, an agentic decision tree where each node is defined as a natural language binary question over the input data, iteratively partitioning the data into two subsets at each step. The best splits are found using TextGrad Yuksekgonul et al. (2024), a prompt-refinement technique providing textual feedback to an agent, to minimize the Gini criterion Breiman et al. (1984) at each split. Ultimately, ACT takes the form of a decision tree (cf. Figure 1), each leaf leading to a class prediction after successive splitting.

The benefits of ACT are therefore the following:

- **Transparent and structured decision process:** ACT enforces a systematic, rule-based structure that aligns the model's internal reasoning with its observable decision process. Its

[1] https://artificialintelligenceact.eu/

[2] https://oecd.ai/en/ai-principles

[3] https://www.nist.gov/itl/ai-risk-management-framework

tree-based architecture provides fully traceable and interpretable decision paths, facilitating auditability, human oversight, and intervention—akin to the properties of traditional expert systems.

- **Optimized decision process:** we show that decomposing some tasks into subquestions allows to improve performances of LLMs without retraining. By automatically finding the most relevant questions at each node, ACT is able to surpass existing methods.

2 Problem Setup and Related Work

We consider a binary classification problem over an *unstructured* input space $\mathcal{X}$ (e.g., text or images), with labels $y \in \{0, 1\}$. Given a dataset $\mathcal{D} = \{(x_i, y_i)\}_{i=1}^N$ sampled i.i.d. from an unknown distribution, the goal is to learn a classifier $f : \mathcal{X} \to \{0, 1\}$ that is both accurate and interpretable. There is an abundant literature in XAI that focuses on defining desiderata for interpretable models Doshi-Velez & Kim (2017); Lipton (2018). Building on the regulatory requirements discussed in Section 1, we therefore aim to build a decision model satisfying the following properties:

- **Accuracy:** competitive predictive performance on unstructured inputs;
- **Transparency:** decisions should be produced via interpretable, step-by-step reasoning;
- **Contestability:** the ability to contest an algorithmic-based decision by disputing the conditions it relies on Wachter et al. (2017); Venkatasubramanian & Alfano (2020); Lyons et al. (2021).

Translating these desiderata into technical requirements is made especially difficult given the opaqueness and complexity of modern architectures such as LLMs and VLMs. There is a clear lack of literature support for what explanations satisfying these properties should look like in these contexts. As a result, in this work, we propose to define interpretability through the notion of **explicit decision paths**—structured sequences of semantically meaningful reasoning steps that can be inspected, understood, and verified by human users. Below, we ground this proposition in the existing XAI and LLM literatures.

Interpretable models and post-hoc explanations Interpretable models such as decision trees (Breiman et al., 1984; Salzberg, 1994) are traditionally viewed as a natural solution to meet the three aforementioned criteria. Offering transparent, rule-based decision paths that are easily understood and audited, their hierarchical structure allows users to trace predictions through explicit, human-readable splits, making them well-suited for regulated or high-stakes domains (Rudin, 2019). However, they are fundamentally limited to structured inputs, relying on predefined features and struggling to handle unstructured data such as raw text or images. Efforts to address interpretability in unstructured data contexts include post-hoc explanation techniques such as LIME (Ribeiro et al., 2016) and SHAP (Lundberg & Lee, 2017) proposing explanations in the form of feature attributions at the pixel (for images) or word (for text) level. Although very popular, these methods have been criticized for often failing to capture the model's true reasoning (Rudin, 2019; Laugel et al., 2019), consequently offering limited guarantees for auditability, consistency, or regulatory compliance.

More recent efforts in neural NLP have attempted to incorporate explanations directly into the model's output, for instance through self-explanation (Nye et al., 2021) or agentic feedback loops (Madaan et al., 2023). However, these approaches rely on free-form natural language generation, which—although interpretable to humans—remains unverifiable and unstructured. As a result, they fall short of providing a consistent and auditable decision procedure.

Workflow and prompt optimization Research on building trustworthy LLM and VLM-based systems has largely focused on workflow design and prompt refinement. Approaches such as AutoPrompt (Shin et al., 2020), TextGrad (Yuksekgonul et al., 2024), and DSPy (Khattab et al., 2023) automatically optimize the prompt instructions to better align the LLM behavior with downstream tasks. These methods show that LLMs can be guided toward more accurate or systematic outputs without retraining, but do not provide explanation patterns. Techniques such as Chain-of-Thought prompting (Wei et al., 2022), ReAct (Yao et al., 2023), and Reflexion (Shinn et al., 2023) structure reasoning through intermediate steps or self-feedback, providing both performance increases and

Table 1: Side-by-side comparison between traditional decision tree algorithms (CART, C4.5) and our proposed Agentic Classification Tree (ACT).

	CART / C4.5	**Agentic Classification Tree**
Data type	Tabular	Unstructured
Node definition	Numerical: binary threshold (both). Categorical: CART binary subset; C4.5 multiway split.	Binary natural-language question queried to an LLM over the inputs.
Best split search	Greedy exhaustive search over candidate splits.	Iterative refinement of the question through prompt optimization (TextGrad) with LLM feedback.
Split criterion	Gini (CART) or Information Gain Ratio (C4.5).	(Gini or IG) combined with semantic purity analysis.
Inference method	Each input follows the tree based on feature values.	Each input is queried by the LLM at each node and routed according to its answers.

transparent decision patterns. However, their reliance on free-form text generation makes them still vulnerable to hallucinations, questioning their efficacy in auditing and verification contexts.

We propose to address these challenges by enforcing a hierarchical tree-structured decision process over unstructured inputs. Defining each node as a natural-language question optimized through LLM feedback ensures both semantic adaptability and transparent reasoning paths. Our proposed framework, called Agentic Classification Tree (ACT), thus aims to combine the interpretability of decision trees with the semantic reasoning of LLMs, providing an accountable alternative to existing black-box prompt optimization methods.

3 METHODOLOGY

To build a structured agentic workflow in the form of a classification tree, we take inspiration from traditional decision tree algorithms such as CART (Breiman et al., 1984) and C4.5 (Salzberg, 1994). Table 1 provides a side-by-side comparison between traditional decision trees and our proposed ACT method.

3.1 TRADITIONAL DECISION TREES: CART AND C4.5

Classical decision-tree algorithms such as CART (Breiman et al., 1984) and C4.5 (Salzberg, 1994) recursively partition a dataset based on feature values. At each node, a split is defined by either a threshold on a numerical feature or a subset of categorical values. The quality of each split is measured by an impurity criterion—typically Gini impurity (CART) or Information Gain ratio (C4.5)—and the algorithm greedily selects the split that yields the greatest impurity reduction, continuing recursively until a stopping criterion is met (e.g., maximum depth or node purity).

These methods are efficient and highly interpretable for structured, tabular data, but cannot operate directly on unstructured inputs such as text or images. This limitation motivates our proposed ACT, which replaces feature-based splits with optimized natural language queries evaluated by an LLM.

3.2 PROPOSED APPROACH: AGENTIC CLASSIFICATION TREE (ACT)

The Agentic Classification Tree (ACT) is designed to extend classical decision tree methods to unstructured data. Like CART and C4.5, ACT recursively partitions the input space, but instead of relying on numeric or categorical features, it leverages large language models (LLMs) to define natural-language queries that guide data splits.

Node definition Each node x in ACT is defined by a natural-language prompt p to semantically partition data via an LLM. Given an instance x_i, the LLM is queried with p and responds with a binary answer (`yes` or `no`), determining the branch assignment. This process defines the following split:

$$\mathcal{D}_L^x = \{(x_i, y_i) \in \mathcal{D}^x \mid f_{\text{split}}(p, x_i) = \texttt{yes}\}, \quad \mathcal{D}_R^x = \mathcal{D}^x \setminus \mathcal{D}_L^x. \tag{1}$$

where $f_{\text{split}}(p, x_i)$ denotes the LLM's response to prompt p on input instance x_i. The resulting subsets $\mathcal{D}_L^x$ and $\mathcal{D}_R^x$ thus reflect the semantic partition induced by the model at node x.

To initialize the splitting process, we begin with a neutral, generic question such as:

```
"Based on the provided example, does it belong to the positive
class?  (yes/no)"
```

This provides an unbiased starting point. Since f_{split} is neither trained on the task nor informed of the target labels, the initial partition is not expected to be meaningful. This is by design: withholding task-specific information ensures that decision prompts are derived from the training data rather than from prior knowledge. In contrast to approaches such as TextGrad (Yuksekgonul et al., 2024), which typically begin with a task-specific question and iteratively refine it, ACT starts from a neutral prompt and discovers semantically meaningful questions from scratch through data-driven refinement.

Best split criterion ACT optimizes each decision node by iteratively refining a natural-language question $p^{(k)}$ to achieve an effective semantic partition of the data. To guide this process, it combines quantitative evaluation (e.g., Gini impurity) with qualitative feedback from the LLM, yielding both statistical and semantic insights for prompt refinement. The procedure consists of two components:

1. **Quantitative Impurity Evaluation.** As in classical decision tree algorithms such as CART and C4.5, the quality of a split induced by the current prompt $p^{(k)}$ is evaluated using standard impurity criteria (e.g., weighted Gini impurity (Breiman et al., 1984) or information gain ratio (Salzberg, 1994)). We denote this score by $\delta(p^{(k)})$, which serves as an objective function quantifying how well the semantic split separates the class labels.
2. **Semantic Purity Analysis via LLM.** While the quantitative evaluation (e.g., Gini impurity) measures class-label heterogeneity within each partition, it does not reveal the *semantic* factors underlying the impurity. To make these sources explicit, we analyze each child node independently by contrasting its class-conditional subsets (correct vs. misclassified examples). For this purpose, we align the predicted answer (`yes`/`no`) with the ground-truth label ($y_i \in \{0, 1\}$) when defining correct and incorrect routing. Although this introduces an artificial correspondence between branch outcomes and class labels, it proved to be more effective in practice: the LLM more readily exploits feedback framed in terms of correct versus incorrect predictions than feedback based on contrasting the two branches directly (e.g., retaining the majority class and pruning the minority one). This alignment does not prevent the emergence of negative queries (e.g., "*Is feature A absent?*"), which the LLM can in principle generate and refine, though such formulations are typically harder to elicit consistently.

   Formally, each subset $g \in \{\mathcal{D}_L^x, \mathcal{D}_R^x\}$ is associated with the model's predicted answer $\hat{y}_g$ to the current prompt. For convenience, we denote by $\hat{y}_g \in \{0, 1\}$ the numeric encoding of this answer ($\hat{y}_g = 1$ for "yes", $\hat{y}_g = 0$ for "no"). We then further partition each subset into two groups:

$$X_{\text{correct}}^x = \{(x_i, y_i) \in g \mid y_i = \hat{y}_g\}, \quad X_{\text{error}}^x = \{(x_i, y_i) \in g \mid y_i \neq \hat{y}_g\}.$$

   Specifically, X_{correct} consists of input–label pairs where the LLM's predicted label $\hat{y}_g$ matches the true label y_i, while X_{error} consists of misclassified pairs where $y_i \neq \hat{y}_g$.

   Next, we prompt the LLM to analyze these two groups and identify key semantic **characteristics** that distinguish them—i.e., features whose presence or absence could explain why some examples are misclassified.

   Based on this contrast, the LLM—denoted f_{purity}—returns concise, actionable feedback s_g that is then used to refine the node's splitting question in the next optimization step to reduce impurity.

LLM Task: Semantic Purity Feedback for `"yes"` group (D_L)
Below are two groups of samples for which a model answered either `"yes"` or `"no"` in response to a natural language prompt to predict their class label.
Provide feedback about key **characteristics** that are present or absent in the group where the model's predicted label matched the true label, and in the group where the prediction was incorrect.
For the following examples, the model answered `"yes"`:

- **Well-classified examples** (true label = `"yes"`):
  [List of correctly classified inputs]
- **Misclassified examples** (true label = `"no"`):
  [List of misclassified inputs]

The feedback you provide must be clear and concise. Focus on the one or two most important **characteristics**.

By performing this semantic analysis independently on subsets D_L and D_R, we obtain targeted feedback s_{D_L} and s_{D_R} that are then used to guide subsequent question-refinement iterations.

Best split search At each decision node, the objective is to refine the natural-language prompt $p^{(k)}$ so that the resulting partition minimizes the weighted Gini impurity $\delta(p^{(k)})$. As described previously, this quantitative criterion is complemented by semantic feedback (s_{D_L}, s_{D_R}), obtained from LLM-based analyses of misclassified versus correctly classified examples.

To jointly evaluate statistical impurity and semantic error patterns, we define a semantic–impurity objective:

$$\mathcal{L}(p^{(k)}) = f_{\text{loss}}\big(p^{(k)}, s_{D_L}, s_{D_R}, \delta(p^{(k)})\big),$$

where f_{loss} is instantiated as a large language model (LLM). Given the current question, impurity statistics, and semantic feedback, the LLM returns a natural-language diagnosis of its limitations, highlighting semantic factors that contribute to impurity.

We realize the refinement process using the differentiable prompting framework *TextGrad* (Yuksekgonul et al., 2024). Each refinement iteration proceeds in two stages:

(i) *Guidance stage:* `TextGrad.feedback` takes as input the prompt $p^{(k)}$ together with its evaluation $\mathcal{L}(p^{(k)})$, and generates a natural-language editing instruction $\nabla_p \mathcal{L}(p^{(k)})$ that specifies how the prompt should be revised.

(ii) *Revision stage:* `TextGrad.step` applies this instruction to produce an updated prompt:

$$p^{(k+1)} = \texttt{TextGrad.step}\Big(p^{(k)}, \nabla_p \mathcal{L}(p^{(k)})\Big).$$

The optimization loop is repeated for up to K steps or until convergence. At each iteration, the current prompt $p^{(k)}$ is used to partition the data, evaluate impurity, and analyze semantic error patterns via the LLM, as described above. These signals jointly inform the loss $\mathcal{L}(p^{(k)})$, which is then used by `TextGrad` to generate a refined prompt $p^{(k+1)}$. Among all prompts generated during the optimization process, $p^{(0)}, p^{(1)}, \ldots, p^{(K)}$, we select the one with the lowest impurity on the training set, i.e., $p^* = \arg\min_k \ \delta(p^{(k)})$.

Tree construction ACT builds the decision tree recursively in a top-down manner. At each node, the best prompt is optimized via the procedure described above. If the resulting subset is pure or meets a predefined stopping criterion (e.g., minimum node size — the number of training examples in the node —, Gini impurity below a threshold, or maximum depth), a leaf node is created. As in classical decision trees, the leaf is assigned the majority class label among the examples it contains. The complete ACT procedure is summarized in Algorithm 1.

Algorithm 1: ACT: Agentic Classification Tree

Input: Dataset D
Output: Decision tree T
$T \leftarrow$ `GrowPromptTree`(D) ;
return T ;

Function `GrowPromptTree(`D`)`**:**
  Initialize prompt $p^{(0)}$; `// default initialization`
  if *D is pure or stopping criterion is met* **then**
    return Leaf node with majority class label;
  for $k = 0$ **to** $K - 1$ **do**
    $D_L \leftarrow \{(x_i, y_i) \in D \mid f_{\text{split}}(p^{(k)}, x_i) = \texttt{yes}\}$;
    $D_R \leftarrow D \setminus D_L$;
    $\delta(p^{(k)}) \leftarrow \frac{|D_L|}{|D|} \cdot \text{Gini}(D_L) + \frac{|D_R|}{|D|} \cdot \text{Gini}(D_R)$;
    foreach $(g, \hat{y}_g) \in \{(D_L, \texttt{'yes'}), (D_R, \texttt{'no'})\}$ **do**
      $X_{\text{correct}} \leftarrow \{(x_i, y_i) \in g \mid y_i = \hat{y}_g\}$;
      $X_{\text{error}} \leftarrow \{(x_i, y_i) \in g \mid y_i \neq \hat{y}_g\}$;
      $s_g \leftarrow f_{\text{purity}}(p^{(k)}, X_{\text{correct}}, X_{\text{error}})$; `// Intra-group semantic analysis`
    $\mathcal{L}(p^{(k)}) \leftarrow f_{\text{loss}}(p^{(k)}, s_{D_L}, s_{D_R}, \delta(p^{(k)}))$; `// LLM-evaluated semantic loss`
    $\nabla_p \mathcal{L}(p^{(k)}) \leftarrow$ `TextGrad.feedback(`$p^{(k)}, \mathcal{L}(p^{(k)})$`)` ;
    $p^{(k+1)} \leftarrow$ `TextGrad.step(`$p^{(k)}, \nabla_p \mathcal{L}(p^{(k)})$`)` ;
  $p^* \leftarrow \arg\min_{p^{(k)}} \delta(p^{(k)})$; `// Prompt with lowest impurity`
  Split D using p^* into D_L and D_R ;
  node $\leftarrow$ Create prompt node with final question p^* ;
  node.left $\leftarrow$ `GrowPromptTree(`D_L`)` ;
  node.right $\leftarrow$ `GrowPromptTree(`D_R`)` ;
  return node ;

Function `Gini(`D`)`**:**
  return $1 - \sum_k p_k^2$, where p_k is the class proportion in D;

4 EXPERIMENTS

We empirically evaluate the proposed Agentic Classification Tree (ACT) against state-of-the-art baseline methods on multiple binary text classification datasets.

Datasets. We evaluate ACT on three binary classification tasks: medical diagnosis (DIAGNO), spam detection (SPAM), and jailbreak prompt classification (JAILBREAK). For DIAGNO and SPAM, we construct balanced splits to ensure controlled evaluation. JAILBREAK is used in its original form, which is moderately imbalanced (details in Appendix B.1).

Models. We evaluate our method on four language models: Gemma-4B Team et al. (2025), GPT-Nano-4.1 Achiam et al. (2023), GPT-Mini-4.1 Achiam et al. (2023), and Qwen3-4B Yang et al. (2025). These models were selected to represent diverse architectural approaches while prioritizing computationally efficient smaller models. To examine the effect of model size, we include two variants of the GPT-4.1 family with different parameter scales.

Baselines. We evaluate ACT against four classification baselines:

- **Chain-of-Thought (CoT, zero-shot)** (Wei et al., 2022): standard zero-shot prompting with step-by-step reasoning, without access to labeled examples or demonstrations.
- **DSPy (8-shot in-context learning)** (Khattab et al., 2023): a demonstration-based approach that conditions the model on eight curated input–label examples to elicit task-relevant behavior.
- **TextGrad** (Yuksekgonul et al., 2024): a prompt optimization method that refines a single task-specific prompt using textual feedback. The prompt includes the task description and class labels, and is optimized for accuracy. We follow the original setup, with no demonstrations provided.

Table 2: Comparison of classification methods across three text datasets (DIAGNO, SPAM, JAILBREAK) and four LLMs. Results show training and test accuracy (%) for baseline methods (CoT, DSPy, TextGrad), traditional CART with TF-IDF features, and the proposed ACT with varying hyperparameters (d: tree depth, k: optimization steps per node).

Dataset	Method	Gemma3 4b		GPT-4.1 Nano		GPT-4.1 Mini		Qwen3 4b		Avg Test
		Train	Test	Train	Test	Train	Test	Train	Test	(%)
DIAGNO	CoT (0-shot)	61.0	61.5	63.5	63.3	64.3	61.2	63.5	63.2	62.3
	DSPy (8 demos)	66.2	64.0	71.5	68.7	68.3	64.5	61.2	58.8	64.0
	TEXTGRAD	65.7	63.3	67.5	64.5	69.2	65.8	65.0	64.0	64.4
	TF-IDF + CART (d=3)	–	–	–	–	–	–	–	–	78.8
	ACT (d=3, k=10)	71.3	65.3	70.7	66.9	78.5	77.3	68.5	64.8	68.6
	ACT (d=3, k=20)	72.5	67.0	71.8	67.2	78.3	76.2	69.2	65.5	69.0
	ACT (d=4, k=10)	72.5	66.3	74.3	**70.5**	79.7	74.3	70.8	66.7	69.5
	ACT (d=4, k=20)	73.7	**68.3**	74.8	70.3	85.0	**82.3**	75.5	**70.3**	72.8
SPAM	CoT (0-shot)	75.0	73.3	96.3	95.5	95.5	95.7	92.5	93.2	89.4
	DSPy (8 demos)	95.2	95.0	98.3	98.3	98.8	98.2	96.7	96.5	97.0
	TEXTGRAD	93.3	92.6	97.8	97.3	98.0	96.7	96.2	96.0	95.7
	TF-IDF + CART (d=3)	–	–	–	–	–	–	–	–	89.3
	ACT (d=2, k=5)	97.0	95.8	98.7	98.6	98.5	98.2	97.1	96.6	97.3
	ACT (d=2, k=10)	97.7	96.5	99.7	**99.2**	100.0	99.2	98.6	98.5	98.4
	ACT (d=3, k=5)	98.3	98.2	98.3	97.0	100.0	**99.4**	98.8	98.5	98.3
	ACT (d=3, k=10)	99.3	**98.5**	99.2	98.8	100.0	99.1	99.0	**98.8**	98.8
JAILBREAK	CoT (0-shot)	81.6	77.9	84.1	85.5	88.1	90.4	78.0	80.1	83.5
	DSPy (8 demos)	85.8	79.8	88.0	88.0	94.8	94.8	85.3	83.5	86.5
	TEXTGRAD	92.9	**91.3**	92.4	88.2	95.6	94.7	82.3	81.5	88.9
	TF-IDF + CART (d=3)	–	–	–	–	–	–	–	–	91.3
	ACT (d=3, k=10)	85.3	82.3	85.5	85.5	96.9	95.4	84.8	82.7	86.5
	ACT (d=3, k=20)	85.2	84.3	92.3	**92.0**	97.7	96.2	84.2	84.7	89.3
	ACT (d=4, k=10)	91.4	90.0	92.7	87.0	98.1	97.2	85.8	85.1	89.8
	ACT (d=4, k=20)	92.1	90.6	93.2	88.2	99.2	**98.9**	87.7	**88.0**	91.4

- **TF-IDF + CART**: a CART trained on TF-IDF representations of the raw texts (Salton et al., 1975; Breiman et al., 1984) that does not rely on a language model, offering structural transparency but limited semantic interpretability, as it lacks natural-language reasoning.

The three LLM-based baseline methods above (CoT, DSPy, TextGrad) are initialized with task-specific prompts that explicitly state the classification objective and define the target classes (e.g., for DIAGNO, whether a case indicates Tuberculosis). In contrast, ACT begins from a generic, task-agnostic query ("Based on the provided example, does it belong to the positive class?"), without access to the task description, class names, or any domain-specific information. Consequently, ACT must uncover the task through successive node's question-refinement iterations, making the comparison conservative with respect to considered baselines. We vary two hyperparameters of the proposed ACT algorithm: the maximum tree depth d and the number of prompt refinement steps per node k, exploring different configurations in our experiments. To control input length and computational cost, f_{purity} is restricted to at most m randomly selected well-classified and m randomly selected misclassified instances per group (with $m = 50$ in our experiments). To ensure a fair comparison, all methods are evaluated using identical underlying language models.

4.1 Experimental Results and Discussion

As shown in Table 2, ACT with appropriately selected hyperparameters consistently matches or surpasses the CoT, TEXTGRAD, and DSPy baselines across datasets and LLMs. Concretely, averaged over all datasets and models, the best-performing ACT configuration ($d = 4$, $s = 20$) improves test accuracy by 5.2 percentage points over DSPy (8-shot) and by 4.7 points over TEXTGRAD, while providing transparent and interpretable question-based decision paths.

Additionally, we compare our method against CART with TF-IDF features at depth 3, which exhaustively searches for optimal splits based on the TF-IDF matrix. This baseline achieves strong performance (78.8% on DIAGNO, 89.3% on SPAM, 91.3% on JAILBREAK). ACT consistently outperforms TF-IDF+CART on SPAM and JAILBREAK for both GPT-4.1 variants, while only Mini exceeds TF-IDF+CART on DIAGNO. However, the TF-IDF+CART approach remains fundamen-

tally limited by its lack of semantic interpretability. The resulting decision trees (see, e.g., Fig. 5 and 6 in Appendix) are typically hard to understand for users, especially when compared to the natural-language questions learned with ACT (cf. Figure 1 for DIAGNO, Figures 3 and 4 for SPAM and JAILBREAK). This stems in part from TF-IDF+CART selecting single words or n-grams as decision criteria, without any regard to context or semantic meaning. As a result, these models struggle with negation, rephrasing, and adversarially crafted text—features that are particularly prevalent in SPAM and JAILBREAK. These limitations explain ACT's consistent advantage in these domains. Furthermore, Appendix B.2 compares the most relevant symptoms for Tuberculosis identified by ACT on the DIAGNO dataset with those from established medical sources, highlighting their strong alignment and the efficacy of ACT in high-stakes scenarios.

4.2 Ablation Studies

The proposed ACT algorithm has two key hyperparameters: the depth of the tree d and the number of optimization steps per node during training k. Additional hyperparameters with less influence include the number m of examples X_0 and X_1 provided to the LLM when performing semantic analysis (f_{purity}) and the maximum number of logical operators L permitted in each generated question (cf. Appendix A.2 for more discussion). As illustrated in Figure 2, the depth d and optimization steps k significantly influence ACT performance. Across datasets and LLMs, test accuracy peaks at depths 4–5, with deeper trees exhibiting decreased test accuracy despite increased training accuracy at depth 6 —an indication of overfitting. Conversely, increasing optimization steps consistently improves performance, though gains plateau after 10–20 steps depending on the configuration, indicating that only a modest number of steps k is sufficient to identify the best splitting question.

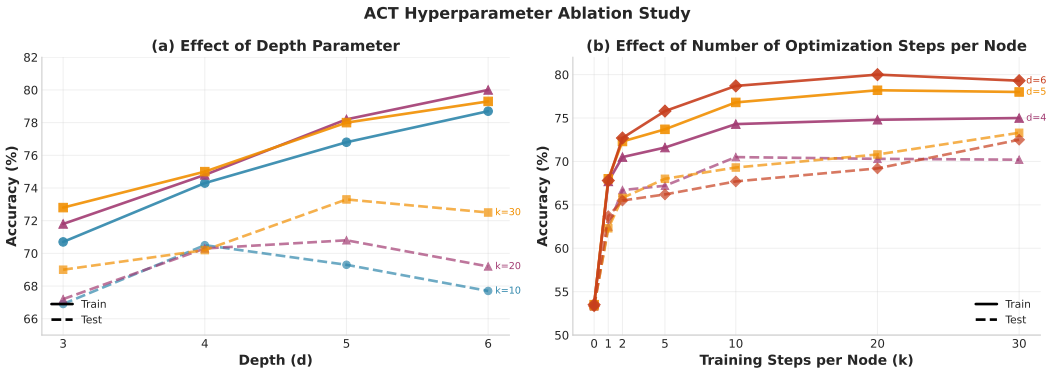

Figure 2: Ablation study of ACT hyperparameters on the DIAGNO dataset using GPT-4.1 Nano. (a) Effect of depth parameter d on model accuracy for different numbers of optimization steps per node during training k. (b) Effect of the number of optimization steps per node during training k on model accuracy for different depth values d. Solid lines represent training accuracy; dashed lines represent test accuracy.

5 Conclusion

We introduced the Agentic Classification Tree (ACT), a novel framework that combines the interpretability of traditional decision trees with the semantic reasoning capabilities of large language models (LLMs). Unlike conventional decision trees that rely on rigid feature-based splits, ACT dynamically optimizes natural-language prompts at each node, using LLM responses to perform semantically meaningful binary splits. Our experiments on text-based binary classification tasks demonstrate that ACT achieves competitive accuracy while providing interpretable, language-based decision logic. This approach highlights the potential of combining classical machine learning structures with modern language model reasoning, opening new avenues for interpretable and effective decision-making in complex, text-rich domains. Future work includes extending ACT to multi-class classification and regression task, improving computational efficiency through targeted prompt optimization strategies, and exploring applications beyond text classification to other structured and unstructured data modalities.

REFERENCES

Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.

Leo Breiman, J. H. Friedman, Richard A. Olshen, and C. J. Stone. *Classification and Regression Trees*. Wadsworth, 1984. ISBN 0-534-98053-8.

Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.

Weize Chen, Yusheng Su, Jingwei Zuo, Cheng Yang, Chenfei Yuan, Chen Qian, Chi-Min Chan, Yujia Qin, Yaxi Lu, Ruobing Xie, et al. Agentverse: Facilitating multi-agent collaboration and exploring emergent behaviors in agents. *arXiv preprint arXiv:2308.10848*, 2(4):6, 2023.

Deysi. Spam detection dataset. `https://huggingface.co/datasets/Deysi/spam-detection-dataset`, 2023. Accessed: 2025-09-17.

Finale Doshi-Velez and Been Kim. Towards a rigorous science of interpretable machine learning. *arXiv preprint arXiv:1702.08608*, 2017.

Arsene Fansi Tchango, Rishab Goel, Zhi Wen, Julien Martel, and Joumana Ghosn. Ddxplus: A new dataset for automatic medical diagnosis. *Advances in neural information processing systems*, 35: 31306–31318, 2022.

Peter Jackson. Introduction to expert systems. 1986.

Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang, Andrea Madotto, and Pascale Fung. Survey of hallucination in natural language generation. *ACM computing surveys*, 55(12):1–38, 2023a.

Ziwei Ji, Tiezheng Yu, Yan Xu, Nayeon Lee, Etsuko Ishii, and Pascale Fung. Towards mitigating hallucination in large language models via self-reflection. *arXiv preprint arXiv:2310.06271*, 2023b.

Omar Khattab, Arnav Singhvi, Paridhi Maheshwari, Zhiyuan Zhang, Keshav Santhanam, Sri Vardhamanan, Saiful Haq, Ashutosh Sharma, Thomas T Joshi, Hanna Moazam, et al. Dspy: Compiling declarative language model calls into self-improving pipelines. *arXiv preprint arXiv:2310.03714*, 2023.

Wojciech Kryściński, Bryan McCann, Caiming Xiong, and Richard Socher. Evaluating the factual consistency of abstractive text summarization. *arXiv preprint arXiv:1910.12840*, 2019.

Thibault Laugel, Marie-Jeanne Lesot, Christophe Marsala, Xavier Renard, and Marcin Detyniecki. The dangers of post-hoc interpretability: unjustified counterfactual explanations. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence*, pp. 2801–2807, 2019.

Kevin J Leonard. Detecting credit card fraud using expert systems. *Computers & industrial engineering*, 25(1-4):103–106, 1993.

Zachary C Lipton. The mythos of model interpretability: In machine learning, the concept of interpretability is both important and slippery. *Queue*, 16(3):31–57, 2018.

Scott M Lundberg and Su-In Lee. A unified approach to interpreting model predictions. *Advances in neural information processing systems*, 30, 2017.

Henrietta Lyons, Eduardo Velloso, and Tim Miller. Conceptualising contestability: Perspectives on contesting algorithmic decisions. *Proceedings of the ACM on Human-Computer Interaction*, 5 (CSCW1):1–25, 2021.

Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegreffe, Uri Alon, Nouha Dziri, Shrimai Prabhumoye, Yiming Yang, et al. Self-refine: Iterative refinement with self-feedback. *Advances in Neural Information Processing Systems*, 36:46534–46594, 2023.

Ninaa510. Diagnosis-text dataset. https://huggingface.co/datasets/ninaa510/diagnosis-text, 2024. Accessed: 2025-09-17.

Maxwell Nye, Anders Johan Andreassen, Guy Gur-Ari, Henryk Michalewski, Jacob Austin, David Bieber, David Dohan, Aitor Lewkowycz, Maarten Bosma, David Luan, et al. Show your work: Scratchpads for intermediate computation with language models. 2021.

Matthew Renze and Erhan Guven. Self-reflection in llm agents: Effects on problem-solving performance. *arXiv preprint arXiv:2405.06682*, 2024.

Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. " why should i trust you?" explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, pp. 1135–1144, 2016.

Cynthia Rudin. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature machine intelligence*, 1(5):206–215, 2019.

Gerard Salton, Anita Wong, and Chung-Shu Yang. A vector space model for automatic indexing. *Communications of the ACM*, 18(11):613–620, 1975.

Steven L Salzberg. C4. 5: Programs for machine learning by j. ross quinlan. morgan kaufmann publishers, inc., 1993, 1994.

Timo Schick, Jane Dwivedi-Yu, Roberto Dessì, Roberta Raileanu, Maria Lomeli, Luke Zettlemoyer, Nicola Cancedda, and Thomas Scialom. Toolformer: Language models can teach themselves to use tools, 2023. *arXiv preprint arXiv:2302.04761*, 2023.

Xinyue Shen, Zeyuan Chen, Michael Backes, Yun Shen, and Yang Zhang. " do anything now": Characterizing and evaluating in-the-wild jailbreak prompts on large language models. In *Proceedings of the 2024 on ACM SIGSAC Conference on Computer and Communications Security*, pp. 1671–1685, 2024.

Taylor Shin, Yasaman Razeghi, Robert L Logan IV, Eric Wallace, and Sameer Singh. Autoprompt: Eliciting knowledge from language models with automatically generated prompts. *arXiv preprint arXiv:2010.15980*, 2020.

Noah Shinn, Federico Cassano, Beck Labash, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. Reflexion: Language agents with verbal reinforcement learning, 2023. *URL https://arxiv.org/abs/2303.11366*, 1, 2023.

Edward H Shortliffe. Medical expert systems—knowledge tools for physicians. *Western Journal of Medicine*, 145(6):830, 1986.

Dag Gundersen Storla, Solomon Yimer, and Gunnar Aksel Bjune. A systematic review of delay in the diagnosis and treatment of tuberculosis. *BMC public health*, 8(1):15, 2008.

Houman Talebzadeh, Sanda Mandutianu, and Christian F Winner. Countrywide loan-underwriting expert system. *AI magazine*, 16(1):51–51, 1995.

Alex Tamkin, Miles Brundage, Jack Clark, and Deep Ganguli. Understanding the capabilities, limitations, and societal impact of large language models. *arXiv preprint arXiv:2102.02503*, 2021.

Gemma Team, Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, Ramona Merhej, Sarah Perrin, Tatiana Matejovicova, Alexandre Ramé, Morgane Rivière, et al. Gemma 3 technical report. *arXiv preprint arXiv:2503.19786*, 2025.

Suresh Venkatasubramanian and Mark Alfano. The philosophical basis of algorithmic recourse. In *Proceedings of the 2020 conference on fairness, accountability, and transparency*, pp. 284–293, 2020.

Sandra Wachter, Brent Mittelstadt, and Chris Russell. Counterfactual explanations without opening the black box: Automated decisions and the gdpr. *Harv. JL & Tech.*, 31:841, 2017.

Lei Wang, Chen Ma, Xueyang Feng, Zeyu Zhang, Hao Yang, Jingsen Zhang, Zhiyuan Chen, Jiakai Tang, Xu Chen, Yankai Lin, et al. A survey on large language model based autonomous agents. *Frontiers of Computer Science*, 18(6):186345, 2024.

Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837, 2022.

An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, et al. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*, 2025.

Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. React: Synergizing reasoning and acting in language models. In *International Conference on Learning Representations (ICLR)*, 2023.

Mert Yuksekgonul, Federico Bianchi, Joseph Boen, Sheng Liu, Zhi Huang, Carlos Guestrin, and James Zou. Textgrad: Automatic" differentiation" via text. *arXiv preprint arXiv:2406.07496*, 2024.

Yongchao Zhou, Andrei Ioan Muresanu, Ziwen Han, Keiran Paster, Silviu Pitis, Harris Chan, and Jimmy Ba. Large language models are human-level prompt engineers. In *The Eleventh International Conference on Learning Representations*, 2022.

A Appendix

A.1 Illustrative Example: Prompt Refinement via Semantic Feedback

We describe the iterative prompt refinement procedure used in ACT, using an illustrative case from the `DIAGNO3` medical diagnosis task. At each decision node, the objective is to construct a binary natural-language question that separates the data into semantically meaningful groups, thereby reducing weighted Gini impurity.

At each iteration k, the following steps are performed:

1. The current prompt $p^{(k)}$ is evaluated over the training set. For each input x_i, the LLM outputs a binary decision (`yes`/`no`).
2. The data is partitioned into two subsets: those routed to the `yes` branch (denoted **sdL**) and those routed to the `no` branch (denoted **sdG**).
3. For each subset, the correctly classified and misclassified examples are separated. These are then passed to the LLM, which is prompted to identify semantic features that distinguish the two groups.
4. Based on this contrastive analysis, the LLM generates feedback indicating which aspects of the current prompt are ambiguous or insufficient. This feedback is then used to revise the question, yielding a new candidate $p^{(k+1)}$.

This iterative refinement continues until a stopping criterion is met (e.g., convergence or maximum number of updates). In practice, we observe that prompts evolve from generic task-agnostic formulations to more specific, domain-relevant queries. For instance, an initial prompt such as *"Does this example belong to the positive class?"* may be refined into a targeted question like *"Does the example mention coughing up blood or high fever?"*, better aligned with the semantic distinctions in the data.

Iteration 0 — Initial Generic Prompt Prompt $p^{(0)}$: `Based on the provided context, does this example belong to the positive class? (yes/no)`

Gini impurity: 0.495

Group predicted "YES" (sdL):

- **Correct predictions:** Examples often exhibited severe and systemic symptoms such as coughing blood, weight loss, shortness of breath, rashes, and fatigue.
- **Misclassified "no" cases:** Contained milder or localized symptoms like eye itching, mild sore throat, or isolated skin rashes, which were insufficiently discriminated by the general prompt.

Group predicted "NO" (sdG):

- **Correct predictions:** Mostly examples with mild and localized symptoms—itchy nose, minor pain, swelling—without strong systemic indicators.
- **Misclassified "yes" cases:** Actually showed critical signs such as coughing up blood, high fever, significant pain, or weight loss, which the current question failed to detect.

Global semantic feedback: *"Focusing the question on specific, highly discriminative features—such as the presence of severe systemic symptoms like coughing blood, weight loss, or high fever—can enhance class separation. This targeted approach reduces ambiguity, improves alignment with key features, and minimizes misclassification."*

Revision $\rightarrow$ $p^{(1)}$**:** `Does this example show severe systemic symptoms such as coughing blood, weight loss, or high fever? (yes/no)`

Iteration 1 — Emphasizing a Single Discriminative Feature **Prompt** $p^{(1)}$**:** `Does this example show severe systemic symptoms such as coughing blood, weight loss, or high fever? (yes/no)`

Gini impurity: 0.470

Group predicted "YES" (sdL):

- **Correct predictions:** Frequently involved coughing up blood alongside systemic symptoms like nasal congestion, eye itching, fatigue, fever, and skin lesions—indicating strong respiratory or infectious patterns.
- **Misclassified "no" cases:** Contained only mild respiratory issues (e.g., itchy or runny nose) and lacked the critical signal of coughing up blood. The model overgeneralized and flagged weakly indicative cases as positive.

Group predicted "NO" (sdG):

- **Correct predictions:** Examples were diverse but lacked specific severe indicators like high fever or respiratory distress. These cases included generalized symptoms (e.g., skin rashes, fatigue) that didn't clearly indicate the positive class.
- **Misclassified "yes" cases:** Contained classic positive indicators such as high fever, muscle pain, fatigue, and systemic lesions. The model failed to recognize the more complex symptom constellation as positive.

Global semantic feedback: *"The current question is too broad and includes features (like fever or weight loss) that are insufficiently discriminative. Focus the question on the most specific positive-class indicator — coughing up blood — and remove less reliable features. This should improve class separation by reducing false positives."*

Revision $\rightarrow$ $p^{(2)}$**:** `Does this example show evidence of coughing up blood? (yes/no)`

Iteration 2 — Broadening to Capture Severe Cases **Prompt** $p^{(2)}$**:** `Does this example show evidence of coughing up blood? (yes/no)`

Gini impurity: 0.468

A.2 Prompt Constraints during Prompt Optimization

During prompt optimization in ACT, additional constraints are applied to keep questions clear and interpretable while encouraging exploration. The goal of these constraints is twofold: (i) ensure that

generated questions remain simple, precise, and auditable, and (ii) incentivize diversity so that ACT explores many different semantic formulations during refinement. The maximum number of logical operators L allowed per question is a controllable hyperparameter in ACT.

LLM Task: Generate Optimized Question (with Constraints)

The following constraints must always be satisfied:

- The question has to be clear and easy.
- The question must focus on at most two **characteristics**.
- The question must include at least one **characteristic** different from the previous one.
- The content of the new question must be significantly different from the current one.
- Use at most L logical operators (`and`/`or`).
- The question has to be answerable with `yes` or `no` only.
- The question must finish with "(yes/no)?".
- Do not use vague words like `could`, `might`, or `possibly`.
- Do not use blanks or placeholder tags like `\___"` or `<...>`.

A.3 Review of Traditional Decision Tree Methods: CART and C4.5

Classification and Regression Trees (CART), proposed by Breiman et al. (1984), are widely used decision tree algorithms for supervised classification tasks. In this work, we focus exclusively on the classification setting.

In this subsection we define $\mathcal{D}$ as a tabular dataset $\mathcal{D} = \{(x_i, y_i)\}_{i=1}^N$, where each instance $x_i \in \mathbb{R}^d$ is associated with a target output y_i, CART constructs a binary decision tree by recursively partitioning the dataset based on threshold splits on numeric input features or binary partitions of categorical inputs, selected in a greedy manner. We detail these steps below.

Node Definition. At each internal node x of the tree, we denote by $\mathcal{D}^x \subseteq \mathcal{D}$ the subset of the training data that reaches node x. In particular, for the root node x_0, we have $\mathcal{D}^{x_0} = \mathcal{D}$, the full training dataset.

To split the data at node x, CART evaluates candidate partitions along a feature dimension j and a threshold $s \in \mathbb{R}$. The dataset $\mathcal{D}^x$ is then divided into two subsets:

$$\mathcal{D}_L^x = \{(x_i, y_i) \in \mathcal{D}^x \mid x_{i,j} \leq s\}, \quad \mathcal{D}_R^x = \{(x_i, y_i) \in \mathcal{D}^x \mid x_{i,j} > s\}. \tag{2}$$

These subsets correspond to the training data that will be used at the left and right child nodes of x, respectively. That is, if y_L and y_R denote the left and right children of node x, we define $\mathcal{D}^{y_L} = \mathcal{D}_L^x$ and $\mathcal{D}^{y_R} = \mathcal{D}_R^x$.

For categorical features, CART evaluates binary partitions over subsets of discrete values, typically by assigning certain categories to the left or right child nodes based on the impurity criterion.

For categorical attributes, CART partitions the data based on whether instances belong or do not belong to specific category subsets.

Best split criterion At each node x, CART selects the optimal feature and split threshold by minimizing a node-specific impurity score. This score is typically computed using the Gini impurity, which quantifies the heterogeneity of class labels within a subset.

Given a candidate split of the local dataset $\mathcal{D}^x$ into $\mathcal{D}_L^x$ and $\mathcal{D}_R^x$, the impurity of the split is defined as the weighted sum of the impurities of the two resulting subsets:

$$\delta_j^s = \frac{|\mathcal{D}_L^x(j,s)|}{|\mathcal{D}^x|} \cdot \mathrm{Gini}(\mathcal{D}_L^x(j,s)) + \frac{|\mathcal{D}_R^x(j,s)|}{|\mathcal{D}^x|} \cdot \mathrm{Gini}(\mathcal{D}_R^x(j,s)), \tag{3}$$

The Gini impurity of any subset $D \subseteq \mathcal{D}$ is given by:

$$\text{Gini}(D) = 1 - \sum_{k \in \mathcal{Y}} p_k^2, \tag{4}$$

where $\mathcal{Y}$ denotes the set of possible class labels, and p_k is the empirical proportion of class k in the dataset D.

A Gini impurity of 0 indicates perfect purity—i.e., all instances in the subset belong to a single class—while a value closer to 0.5 (in binary classification) indicates a highly mixed, impure subset.

Best split search The CART algorithm constructs the decision tree recursively using a greedy strategy: at each internal node x, the best split is selected by optimizing a local impurity criterion over the current subset $\mathcal{D}^x$. This split is chosen by evaluating its effect on the resulting child nodes—i.e., by computing the weighted impurity of the partition it induces. In practice, this local optimization is typically performed via exhaustive search over all candidate features and thresholds, though heuristic approximations may be used in high-dimensional settings.

The recursive construction proceeds until a stopping criterion is met, such as achieving node purity, reaching a minimum number of samples, or exceeding a predefined maximum depth $d_{\max}$.

Tree construction. Decision trees are constructed recursively. At each node, the split minimizing impurity is applied to partition the data, and the procedure recurses on the resulting subsets. Recursion terminates when purity is reached or a predefined stopping condition is met (e.g., maximum depth, minimum node size), and the node is labeled with the majority class.

C4.5. C4.5 (Salzberg, 1994) follows the same recursive tree-building process as CART but differs mainly in two respects: (i) it chooses splits by maximizing the *information gain ratio* rather than minimizing Gini impurity, and (ii) it allows multiway splits on categorical attributes, whereas CART restricts all splits to be binary.

While both CART and C4.5 have become canonical algorithms for structured data, they are not directly applicable to unstructured domains such as text or images. This limitation motivates our proposed Agentic Classification Tree (ACT).

B Additional results and datasets

B.1 Details on datasets

We consider three publicly available binary text classification datasets spanning different domains:

- **DIAGNO** (Ninaa510, 2024): a medical diagnosis dataset (tuberculosis vs. allergic sinusitis) , a derived dataset based on DDXPlus Fansi Tchango et al. (2022). We constructed a balanced subset with 600 training, 100 validation, and 600 test samples, ensuring a 50/50 split (300 tuberculosis and 300 allergic sinusitis cases).
- **SPAM**: an email spam detection dataset derived from Deysi (2023). We constructed a balanced subset with 600 training, 100 validation, and 600 test samples, following the same procedure as for DIAGNO.
- **JAILBREAK** (Shen et al., 2024): a jailbreak prompt classification dataset with 923 training, 102 validation, and 249 test samples (32 examples have been dropped due to high context tokens).

B.2 Qualitative analysis of the questions generated by ACT

In addition to its tree structure, another crucial element ensuring the interpretability of ACT lies in how relevant the generated questions are. For this purpose, we propose to compare these questions with the ones that domain experts would have asked, focusing on the DIAGNO dataset. We survey several relevant medical sources to collect a list of symptoms that are commonly used for tubercu-

Table 3: Symptoms comparison between ACT and medical sources.

Questions in common	ACT only	Medical sources only
Coughing up blood	Shortness of breath	Chest pain
Fever	Severe pain	Headache
Weight loss		Increased sweating
Swollen lymph nodes		Loss of appetite
Fatigue		

losis diagnosis. These sources include websites of the World Health Organization[4] and the British National Health Service[5], and a literature review on tuberculosis diagnosis Storla et al. (2008).

Table 3 shows a side-by-side comparison of the symptoms identified by ACT as positively associated with tuberculosis, and those most commonly identified in the aforementioned medical sources, differentiating the symptoms identified by both ACT and the medical sources (left column), from those that were exclusive (center and right columns). We see that most of the symptoms appear in common. Some of the differences may be explained by ambiguity between some symptoms (e.g. "severe pain" vs. Chest pain and headaches). Others may come from binary classification setting considered, opposing tuberculosis cases to allergy, potentially resulting in some symptoms (e.g. loss of appetite) not being represented in the dataset.

Overall, this tends to show that ACT is indeed able to identify relevant symptoms to make its predictions, showcasing its utility and reliability.

B.3 Agentic Classification Tree for SPAM Dataset

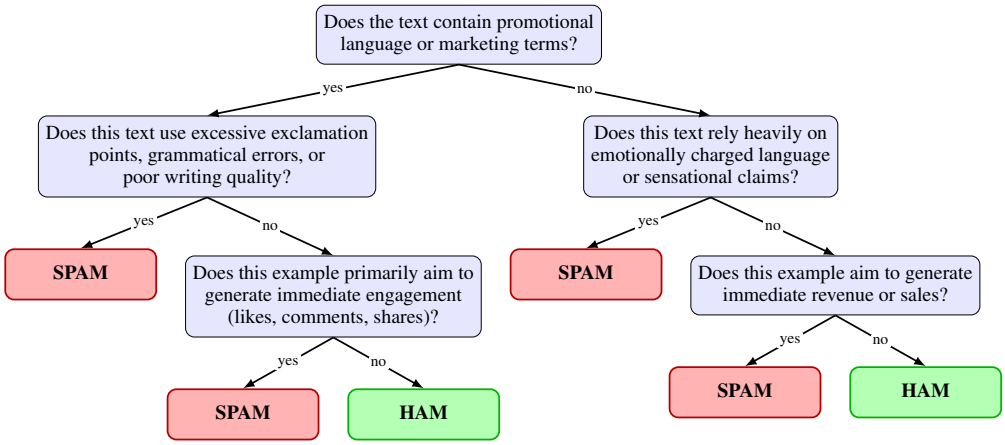

Figure 3: Decision tree of depth 3 generated by the Gemma3-4B model for spam email classification. The tree distinguishes between spam and legitimate email (ham) through hierarchical semantic questions about promotional content, writing quality, and intent. Each internal node represents a binary question optimized through the ACT framework to maximize class separation, with leaf nodes indicating the final classification based on the majority class of training examples.

[4]https://www.who.int/news-room/fact-sheets/detail/tuberculosis
[5]https://www.nhs.uk/conditions/tuberculosis-tb/

B.4 AGENTIC CLASSIFICATION TREE FOR JAILBREAK DATASET

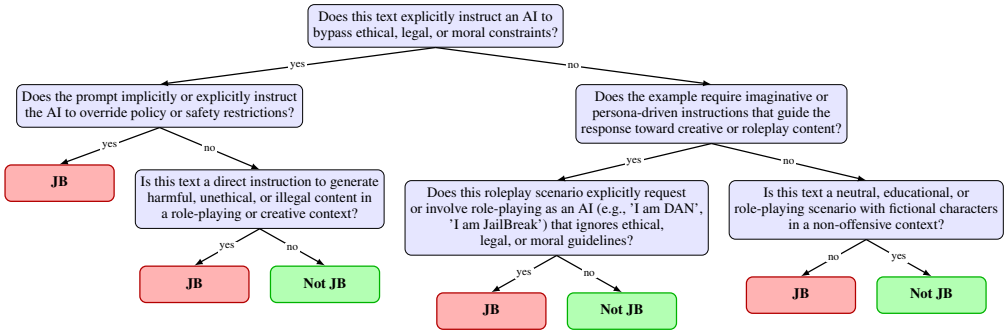

Figure 4: Decision tree of depth 3 generated by the Qwen3-4b model for jailbreak prompt classification. The tree recursively partitions inputs through binary natural-language questions optimized to distinguish between jailbreak attempts (JB) and legitimate prompts (Not JB). Each internal node contains a semantically meaningful question discovered through iterative prompt refinement, with terminal nodes indicating the final classification based on the majority class of training examples reaching that leaf.

B.5 QUALITATIVE COMPARISON WITH TF-IDF + CART

In this section we aim to illustrate the benefits of leveraging LLM agents to directly deal with unstructured data by comparing ACT with the traditional TF-IDF and CART combination. The tree resulting from this procedure for the DIAGNO dataset is shown in Figure 5, and in Figure 6 for the SPAM dataset. In both cases, we observe the limitation of this approach, as nodes built only on one word provide little help in understanding the model. Worse, these words may often end up being generic adverbs and prepositions (e.g. ”our”, ”any”), resulting ultimately in explanations that are both hard to understand and leverage.

On the other hand, the decision trees learned by ACT for these datasets can be easily understood, and verified, by a user.

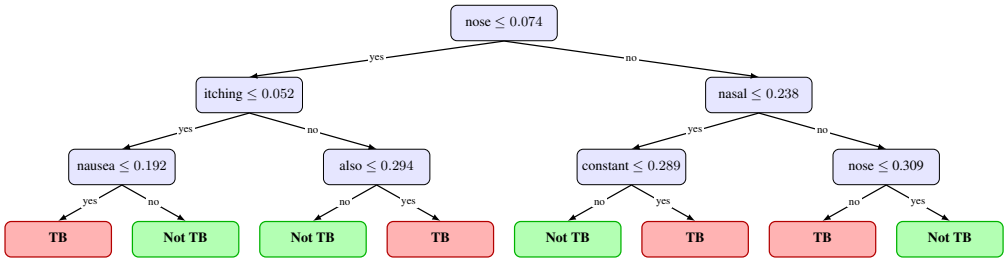

Figure 5: Decision tree of depth 3 generated by training a CART model after performing a TF-IDF preprocessing step on the DIAGNO dataset.

APPENDIX: USE OF LANGUAGE MODELS

During the preparation of this manuscript, we made limited use of large language models to enhance the clarity and readability of the text. This involved assistance with grammar, phrasing, and stylistic improvements, particularly in the abstract and selected explanatory sections. All scientific content, including the formulation of research questions, experimental design, results, and interpretations, was developed solely by the authors. No language model was used to generate original ideas, proofs, or analyses.

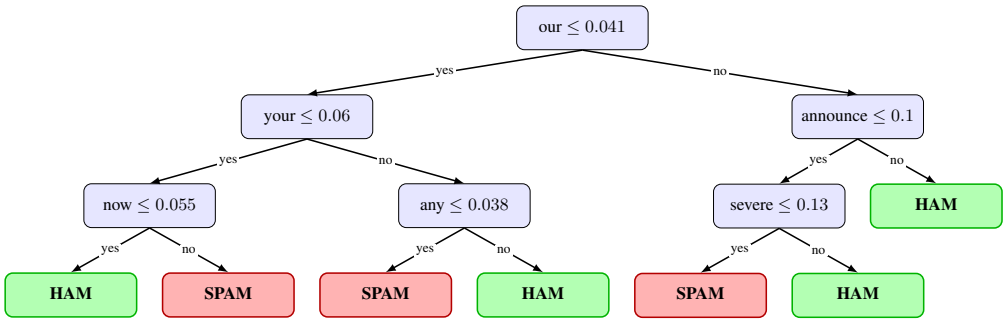

Figure 6: Decision tree of depth 3 generated by training a CART model after performing a TF-IDF preprocessing step on the SPAM dataset.