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.

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 (EU, 2024) or the NIST AI Risk Management Framework (NIST, 2023a) 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 have been addressed by 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 (Slack et al., 2019), directly supporting the requirements set out by regulatory frameworks (NIST, 2023b). 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 NLP (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., 2024; Wang et al., 2024). Yet, despite their impressive capabilities, LLMs still suffer from several drawbacks—such as hallucinations, unreliability and difficulty to audit (Kryscinski et al., 2020; 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, 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., 2023; Ji et al., 2023b; Renze & Guven, 2024; Yuksekgonul et al., 2025), introduce agentic feedback loops to iteratively refine model outputs. While these methods improve reliability, they still rely on free-form text generation and implicit reasoning that remain difficult to verify or audit.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

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 ... "

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 ... "

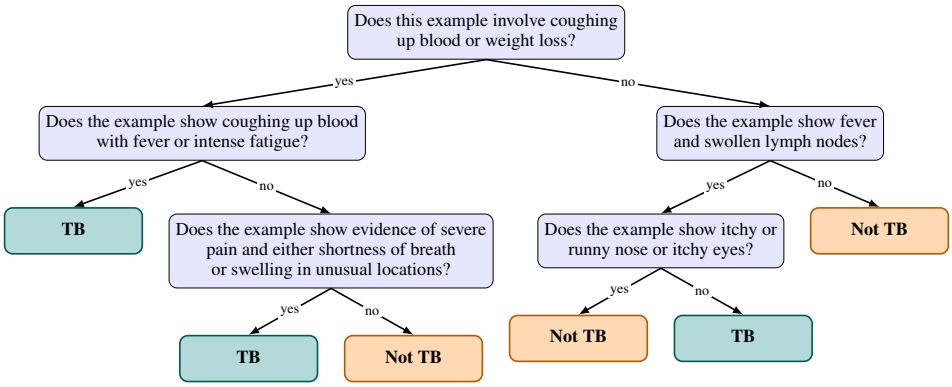

Figure 1: **Example ACT for tuberculosis (TB) diagnosis** using unstructured, free-text patient descriptions. A decision tree is automatically learned, with each node containing a binary natural-language question, autonomously discovered via recursive prompt refinement to maximize label separation at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

In this paper, we introduce a new type of classifier designed to combine the semantic reasoning capabilities of LLMs with the transparency of decision trees. The idea is to follow a divide-and-conquer paradigm by decomposing a difficult classification question (e.g., *Does the patient have disease X?*) into a series of easier questions (e.g., *Is symptom Y present?*), as shown in Figure 1. We therefore propose to structure LLM and agentic reasoning as an agentic decision tree, addressing complex classification tasks on *unstructured data* through verifiable decision patterns while reducing reasoning errors via hierarchical problem decomposition.

Adapting traditional decision tree algorithms such as CART (Breiman et al., 1984) and C4.5 (Salzberg, 1994), we propose the Agentic Classification Tree (ACT). In ACT, each node is defined as a binary natural-language 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., 2025), 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 tree-based architecture provides fully traceable and interpretable decision paths, facilitating auditability, human oversight and intervention.
- **Optimized decision process:** we show that decomposing some tasks into subquestions improves the performance of LLMs without retraining.

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 algorithm-based decision by disputing the conditions it relies on (Wachter et al., 2018; Venkatasubramanian & Alfano, 2020; Lyons et al., 2021).

Translating these desiderata into technical requirements is challenging given the opacity of modern architectures such as LLMs, with limited literature guidance on suitable explanations in these contexts. We therefore propose to define interpretability through the notion of **explicit decision paths** — structured sequences of semantically meaningful reasoning steps that human users can inspect, understand and verify. 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. Their transparent, rule-based structure allows users to trace predictions through 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 the word level (for text). 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, with approaches such as AutoPrompt (Shin et al., 2020), TextGrad (Yuksekgonul et al., 2025) and DSPy (Khattab et al., 2024) automatically optimizing prompt instructions to align 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 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 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 Agentic Classification Tree framework combines 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 them and our proposed ACT method.

3.1 TRADITIONAL DECISION TREES: CART AND C4.5

Classical decision-tree algorithms 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.

Table 1: Comparison between traditional decision tree methods and ACT.

	CART / C4.5	**ACT (ours)**
Data type	Tabular features	Unstructured inputs (text, images, etc.)
Node definition	Numerical: binary threshold, Categorical: binary subset or multiway split	Binary natural-language question queried by an LLM over the inputs.
Best split search	Greedy search over candidate feature splits.	Iterative question refinement via TEXTGRAD.
Split criterion	Gini (CART) or Information Gain Ratio (C4.5).	Gini impurity combined with semantic signals.
Inference	Each input follows the tree based on feature values.	Input is queried by the LLM at each node and routed according to its (yes/no) answers.

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. These methods are efficient and interpretable (Molnar et al., 2020) 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.

3.2 PROPOSED APPROACH: AGENTIC CLASSIFICATION TREE

ACT is designed to extend classical decision tree methods to unstructured data by replacing numeric or categorical feature splits with LLM-generated natural-language questions.

Node definition. Each node n in ACT is defined by a 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:

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

where $f_{\text{split}}(p, x_i)$ denotes the LLM's response to prompt p on input x_i. The resulting subsets $\mathcal{D}_L^n$ and $\mathcal{D}_R^n$ thus reflect the semantic partition induced by the model at node n. The splitting process is initialized with a neutral, generic question and in our experiments we use:

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

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. Unlike TextGrad, which begins with a task-specific question and iteratively refines it, ACT starts from a neutral prompt and discovers meaningful questions 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, the quality of a split induced by the current prompt $p^{(k)}$ is evaluated using standard impurity criteria $\delta(p^{(k)})$, e.g. weighted Gini impurity, 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 $g \in \{\mathcal{D}_L^n, \mathcal{D}_R^n\}$ 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 partition each subset into correct and misclassified groups: $X^n_{\text{correct}} = \{(x_i, y_i) \in g \mid y_i = \hat{y}_g\}$ and $X^n_{\text{error}} = \{(x_i, y_i) \in g \mid y_i \neq \hat{y}_g\}$.

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 the misclassification of some examples. 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. 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.

LLM Task: Semantic Purity Feedback for the `"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 on 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**.

Best split search. At each node, the objective is to refine the 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 an 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., 2025). Each refinement iteration proceeds in two stages:

1. *Guidance stage:* `TextGrad.feedback` takes the prompt $p^{(k)}$ and its evaluation $\mathcal{L}(p^{(k)})$ as input, and generates a natural-language editing instruction $\nabla_p \mathcal{L}(p^{(k)})$ specifying how to revise the prompt.
2. *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 runs 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 candidate prompts $p^{(0)}, \ldots, p^{(K)}$, we select the one with the lowest training impurity $p^* = \arg\min_k \delta(p^{(k)})$.

Algorithm 1: ACT: Agentic Classification Tree

Input: Dataset D, max depth d, number of iterations K
Output: Decision tree T
$T \leftarrow$ `GrowPromptTree`(D, d, K);
return T;

Function `GrowPromptTree(`D, d, K`)`**:**
  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 **do**
    $D_L \leftarrow \{(x_i, y_i) \in D \mid f_{\text{split}}(p^{(k)}, x_i) = \texttt{'yes'}\}$;
    $D_R \leftarrow D \setminus D_L$;
    $\delta(p^{(k)}) \leftarrow \frac{|D_L|}{|D|} \cdot \text{Gini}(D_L) + \frac{|D_R|}{|D|} \cdot \text{Gini}(D_R)$;
    if $k = K$ **then break**;
    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 question p^*;
  node.left $\leftarrow$ `GrowPromptTree`(D_L, d, K);
  node.right $\leftarrow$ `GrowPromptTree`(D_R, d, K);
  return node;

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

Tree construction. As summarized in Algorithm 1, 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, 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.

Tree interpretability. Decision trees, whose decisions rely on transparent structures and parameters, are widely considered to be interpretable (Molnar et al., 2020). Owing to its shallow decision-tree structure and natural-language node questions, ACT exhibits the same desirable properties. In practice, interpretability depends on both the tree depth (i.e., the number of questions along a decision path) and the clarity of the questions. Accordingly, we restrict the tree depth to 2–4 to keep decision paths concise and explicitly instruct TEXTGRAD during training to generate simple questions and limit the number of logical clauses (`and`/`or`) to at most $L = 2$, as detailed in Appendix B.2.

4 EXPERIMENTS

Datasets. We empirically evaluate ACT on five binary text-classification tasks: The **DIAGNO** dataset contains short clinical descriptions labeled as tuberculosis (TB) or not TB, and we use balanced train–test splits for controlled evaluation. **SPAM** consists of email messages annotated as spam or ham, also evaluated with balanced splits. **JAILBREAK** contains user prompts labeled according to whether they attempt to circumvent safety policies, and is used in its original moderately imbalanced form. **BANKCHURN** is a customer-churn dataset in which tabular attributes are serialized into short textual profiles following the protocol used by Arzaghi et al. (2025). This task emphasizes dataset-specific attribute relationships rather than broad world knowledge. **IMDB** is a standard sentiment dataset of movie reviews; we use a balanced subset for consistency.

Table 2: **Comparison of classification methods** across five datasets and four LLMs. Results show accuracy and F1-score on the test set (in %) for six baselines, as well as the proposed ACT with varying hyperparameters (d: tree depth, k: optimization steps per node).

Method	Gemma3 4b		GPT-4.1 Nano		GPT-4.1 Mini		Qwen3 4b		Avg	
	Acc	F1	Acc	F1	Acc	F1	Acc	F1	Acc	F1
DIAGNO										
CoT (0-shot)	61.5	69.5	63.3	54.7	61.2	51.4	63.2	41.7	62.3	54.3
DSPy (BFSR, 8 demos)	64.0	68.2	68.7	61.2	64.5	59.4	63.3	54.9	65.1	60.9
TEXTGRAD	63.3	56.7	64.5	53.0	65.8	55.5	64.0	60.0	64.4	56.3
TF-IDF + CART (d=3)	– –	– –	– –	– –	– –	– –	– –	– –	78.8	78.9
Con. tag.	65.5	63.9	59.3	58.1	75.3	78.6	58.8	64.7	64.7	66.3
Rule Fit	– –	– –	– –	– –	– –	– –	– –	– –	80.2	79.3
ACT (d=3, k=10)	65.3	66.5	66.8	67.0	77.3	75.2	64.8	61.3	68.6	67.5
ACT (d=4, k=20)	**68.3**	**70.8**	**70.3**	**70.8**	**82.3**	**81.8**	**70.3**	**73.5**	72.8	74.2
SPAM										
CoT (0-shot)	73.3	78.6	95.5	95.3	95.7	95.5	93.2	93.6	89.4	90.8
DSPy (BFSR, 8 demos)	95.0	95.2	98.3	98.3	98.2	98.2	97.2	97.2	97.2	97.2
TEXTGRAD	95.0	94.7	97.3	97.3	96.7	96.6	96.0	96.0	96.3	96.2
TF-IDF + CART (d=5)	– –	– –	– –	– –	– –	– –	– –	– –	92.7	93.1
Con. tag.	83.2	85.5	64.2	45.8	82.3	85.0	62.3	39.6	73.0	64.0
Rule Fit	– –	– –	– –	– –	– –	– –	– –	– –	97.0	96.9
ACT (d=2, k=5)	95.8	95.8	98.7	98.6	98.2	98.2	96.7	96.6	97.4	97.3
ACT (d=3, k=10)	**98.5**	**98.5**	**99.2**	**99.2**	**99.5**	**99.5**	**98.8**	**98.8**	99.0	99.0
JAILBREAK										
CoT (0-shot)	77.9	82.2	85.5	83.5	93.6	93.3	81.1	77.7	84.5	84.2
DSPy (BFSR, 8 demos)	89.6	89.3	91.6	**91.6**	94.8	94.8	85.5	83.9	90.4	89.9
TEXTGRAD	**91.6**	**91.3**	90.4	90.3	94.8	95.1	88.8	90.0	91.4	91.7
TF-IDF + CART (d=4)	– –	– –	– –	– –	– –	– –	– –	– –	92.8	92.4
Con. tag.	82.7	82.2	58.6	31.8	82.7	79.8	76.3	70.1	75.1	66.0
Rule Fit	– –	– –	– –	– –	– –	– –	– –	– –	96.8	96.8
ACT (d=3, k=10)	82.3	82.7	85.5	85.1	95.2	95.5	82.7	82.6	86.4	86.5
ACT (d=4, k=20)	90.8	90.0	**92.0**	91.5	**98.8**	**98.8**	**92.0**	**92.7**	93.4	93.3
BANKCHURN										
CoT (0-shot)	47.8	48.1	48.5	25.9	52.8	60.5	50.2	37.3	49.8	43.0
DSPy (BFSR, 8 demos)	52.8	55.6	52.3	42.8	57.2	58.5	54.0	52.2	54.1	52.3
TEXTGRAD	52.0	60.0	53.3	60.3	55.3	61.0	52.8	61.4	53.4	60.7
TF-IDF + CART (d=5)	– –	– –	– –	– –	– –	– –	– –	– –	63.2	66.8
Con. tag.	57.2	60.2	52.0	59.8	59.7	60.8	54.3	61.6	55.8	60.6
Rule Fit	– –	– –	– –	– –	– –	– –	– –	– –	63.0	59.8
ACT (d=3, k=10)	56.3	42.3	58.8	67.4	60.5	64.6	57.5	**62.7**	58.3	59.3
ACT (d=4, k=20)	**65.2**	**70.7**	**65.0**	**69.1**	**68.7**	**68.6**	**62.0**	60.0	65.2	67.1
IMDB										
CoT (0-shot)	92.3	91.9	94.2	94.1	95.7	95.6	94.2	94.0	94.1	93.9
DSPy (BFSR, 8 demos)	**93.7**	**93.6**	**95.5**	**95.5**	**96.8**	**96.8**	94.8	94.8	95.2	95.2
TEXTGRAD	93.2	92.9	95.2	95.1	95.8	95.8	94.7	94.5	94.7	94.6
TF-IDF + CART (d=3)	– –	– –	– –	– –	– –	– –	– –	– –	66.8	72.8
Con. tag.	74.8	76.9	56.2	31.7	50.0	66.7	51.3	63.8	58.1	59.8
Rule Fit	– –	– –	– –	– –	– –	– –	– –	– –	75.2	73.7
ACT (d=3, k=10)	93.2	93.2	92.8	92.8	95.7	95.7	93.8	93.9	93.9	93.9
ACT (d=3, k=20)	**93.7**	**93.6**	94.7	94.7	96.3	96.2	**95.5**	**95.5**	95.1	95.0

Models. We evaluate our method on four language models: Gemma3-4B (Gemma Team et al., 2025), GPT-4.1-nano (Achiam et al., 2023), GPT-4.1-mini (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 assess the effect of model scale, we report additional scaling results with full GPT-4.1 (Achiam et al., 2023).

Baselines. We compare ACT against six baselines: **Chain-of-Thought (CoT, zero-shot)** (Wei et al., 2022). **DSPy (BFSR, 8 demos)** (Khattab et al., 2024) following the setup of Yuksekgonul et al. (2025): Bootstrapped Few-Shot Random Search (BFSR) generates reasoning traces and answers for all training samples, filters for correct answers, constructs candidate prompts with up to 8 correct demonstrations (input, reasoning trace and output), evaluates the candidate prompts on the whole training set and selects the best-performing one. **TextGrad** (Yuksekgonul et al., 2025) refines a task-specific prompt through textual feedback, optimizing for accuracy given the task description and class labels. We use a batch size of 50 samples and run 12 epochs with 3 optimization steps each, yielding 36 refinement iterations where each candidate is evaluated on a validation set and the best is retained. **TF-IDF + CART** trains a CART classifier (depths 2–5) on TF-IDF text representations (Salton et al., 1975; Breiman et al., 1984), exhaustively searching for optimal

Table 3: **Interpretability statistics** averaged over the 4 LLMs on DIAGNO. For reference, the generic seed question *"Based on ... positive class? (yes/no)"* has 77 characters.

ACT (d, k)	AVG. Q. LEN	AVG. PATH
(3, 10)	92.00	2.35
(3, 20)	93.15	2.51
(4, 10)	78.05	3.03
(4, 20)	80.53	3.06

Table 4: **Symptoms comparison** between ACT and medical sources.

IN COMMON (5): COUGHING UP BLOOD, FEVER, WEIGHT LOSS, SWOLLEN LYMPH NODES, FATIGUE
ACT ONLY (2): SHORTNESS OF BREATH, SEVERE PAIN
MEDICAL SOURCES ONLY (4): CHEST PAIN, HEADACHE, INCREASED SWEATING, LOSS OF APPETITE

splits without relying on a language model. This offers structural transparency but limited semantic interpretability. The best-performing depth per dataset is reported in Table 2, with full results in Appendix E. **LLM Concept Bottleneck + CART (Con. tag.)** is a concept-bottleneck baseline where an LLM extracts salient domain concepts and encodes each text as a binary concept vector, on which a shallow CART is then trained. **RuleFit (TF-IDF rules)** (Friedman & Popescu, 2008) is a rule-ensemble method that learns sparse, human-readable if–then rules over TF-IDF features.

Configurations. The three prompt-based baseline methods (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? (yes/no)*"), without access to the task description or class names. Consequently, ACT must uncover the task by iteratively refining each node's question over k steps, making the comparison conservative with respect to the considered baselines. We vary two hyperparameters of the proposed ACT algorithm: the tree depth d and the number of prompt refinement steps per node k. To control input length and computational cost, f_{purity} uses at most m randomly selected well-classified and m randomly selected misclassified instances per group (with $m = 50$ in our experiments).

4.1 EXPERIMENTAL RESULTS AND DISCUSSION

Test set accuracy. As shown in Table 2, ACT with appropriately selected hyperparameters consistently matches or surpasses the four LLM-based baselines (CoT, TEXTGRAD, Con. Tag. and DSPy) across datasets and LLMs. Averaged over all datasets and LLMs, the best-performing ACT configuration improves test accuracy by 4.7 percentage points over DSPy and by 5.1 points over TEXTGRAD, while providing interpretable decision paths. When evaluated by F1-score, ACT outperforms all four LLM-based baselines in 14 of the 20 dataset-model configurations.

Compared to TF-IDF + CART, ACT consistently achieves similar or higher accuracy on all tasks except for DIAGNO, where only the GPT-4.1-mini variant surpasses this baseline. However, TF-IDF+CART remains fundamentally limited by its lack of semantic interpretability. The resulting decision trees (cf. Figures 7 and 8) are typically difficult to interpret compared to the natural-language questions learned by ACT (cf. Figures 1, 5, and 6). This stems from TF-IDF+CART selecting individual words or n-grams as decision criteria without regard to context or semantic meaning, causing the model to struggle with negation, irony, rephrasing and adversarially crafted text—features particularly prevalent in SPAM, JAILBREAK and IMDB.

Compared to RuleFit, ACT with GPT-4.1-mini consistently achieves higher accuracy on all tasks, but RuleFit outperforms ACT on average for DIAGNO and JAILBREAK. Like TF-IDF + CART, RuleFit offers limited interpretability because its rules operate on TF-IDF embeddings (e.g., nasal ≤ 0.238), whereas ACT provides human-readable yes/no questions. Additional experiments (Appendix E) show that ACT can remain competitive with strong non-interpretable baselines (XGBoost, BERT, RoBERTa) while providing interpretable and human-readable decision paths.

Interpretability. To quantify ACT's interpretability, we report the average path length at inference and average node-question length in Table 3 (full results in Appendix F). Predictions require only about three nodes even for depth-4 trees and questions remain easily understandable, indicating that soft constraints on question complexity, particularly the limit L on logical clauses, are effective. These

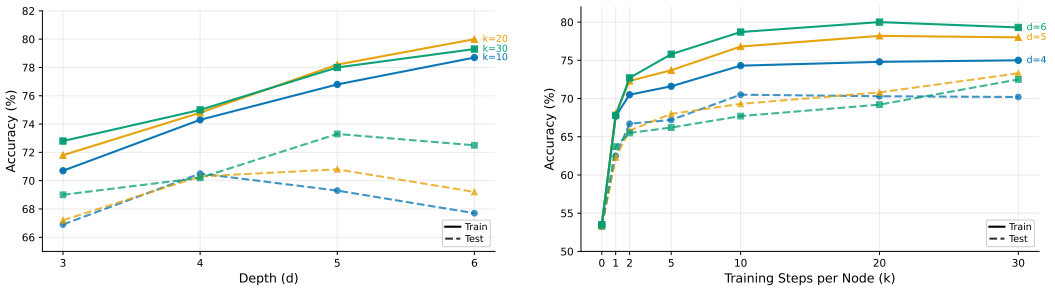

Figure 2: **Ablation study of ACT hyperparameters** on DIAGNO using GPT-4.1 Nano. Left: effect of tree depth d. Right: effect of optimization steps k.

quantitative results, together with manual inspection of generated trees (Figure 1; Appendix D.2), confirm that ACT's divide-and-conquer approach of decomposing complex classification tasks into simpler subquestions yields interpretable decision paths while improving accuracy. Furthermore, Table 4 compares the symptoms for Tuberculosis identified by ACT with established medical sources (WHO, 2024), revealing strong alignment and demonstrating ACT's efficacy in high-stakes scenarios (cf. Appendix D.1).

Ablation studies. As shown in Figure 2, tree depth d and optimization steps per node k both significantly influence ACT's performance. As expected for tree-based methods, increasing depth beyond a certain point (e.g., from 5 to 6 for DIAGNO with GPT-4.1-Nano) leads to overfitting. Increasing k consistently improves performance, though gains plateau after 10–20 steps. Additional results, including ablations on the number m of examples provided to f_{purity} during semantic analysis, are in Appendix G.

Scaling analysis. We evaluate how ACT performance scales with model capacity across the GPT-4.1 family (Nano, Mini and full GPT-4.1). As shown in Figure 3, ACT's test accuracy consistently improves with stronger models, whereas DSPy and TextGrad do not exhibit consistent scaling behavior, indicating that ACT more effectively leverages the enhanced reasoning capabilities of larger models.

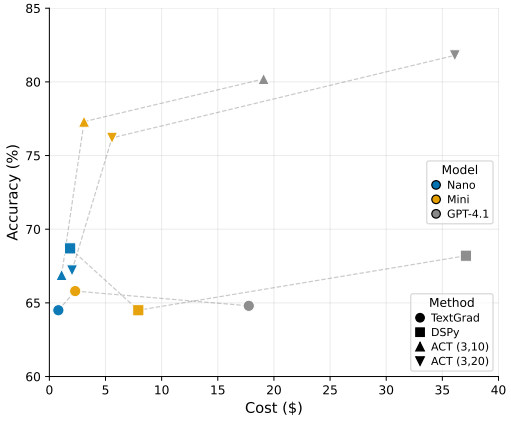

Figure 3: **Scaling analysis across the GPT-4.1 family.** Test accuracy (%) versus training cost (USD) for TEXTGRAD, DSPy, ACT ($d = 3, k = 10$) and ACT ($d = 3, k = 20$) on DIAGNO using Azure-hosted GPT-4.1 models.

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 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 tasks, improving computational efficiency through targeted prompt optimization strategies, and exploring applications beyond text classification to other structured and unstructured data modalities.

AI USE STATEMENT

In this work, we used generative AI tools to assist with the creation of the code base used to perform the experiments and with the preparation of this manuscript. Specifically, LLMs were used to help improve the structure and readability of the code and to assist with debugging. In the preparation of this manuscript, LLM-based tools were used to enhance the clarity and readability of the text. This involved assistance with grammar, phrasing, and stylistic improvements. We have reviewed all AI-assisted work. The code was carefully verified and tested, and the manuscript was re-read several times by the two lead authors. We have not used generative AI tools for any of the scientific content; the formulation of research questions, the experimental design, the results, and their interpretation were developed solely by the authors. No language model was used to generate original ideas, proofs, or analyses, and the rest of the required disclosure tasks are not applicable to this work. We take responsibility for the final content of this work, including text, claims or artifacts produced with the aid of generative AI.

ETHICS STATEMENT

Algorithmic decision making has a broad range of ethical implications, for example, allocative harms or representational harms (Shelby et al., 2023). Bias in algorithmic decision making especially often harms minorities (Lum & Isaac, 2016; Buolamwini & Gebru, 2018; Obermeyer et al., 2019). To reduce these harms and for other reasons, recent governance frameworks and regulations such as the EU AI Act (EU, 2024), the OECD AI Principles (OECD, 2019) and the NIST AI Risk Management Framework (NIST, 2023a) emphasize that AI-based decisions in high-stakes scenarios must be explainable and subject to human oversight. By providing explicit, interpretable decision paths that are accessible not only to technical experts but also to stakeholders, auditors and regulators, ACT offers an algorithmic decision-making framework that combines competitive accuracy with transparency and contestability.

Nonetheless, two important limitations must be considered. First, while ACT reduces reasoning errors, it remains LLM-based and therefore may still exhibit hallucinations, inconsistencies or unreliability. Second, as noted by O'Neil (2017), transparency and contestability, even if guaranteed by the algorithm, are only effective in practice when embedded within appropriate institutional or regulatory frameworks that provide meaningful avenues for recourse. We therefore urge practitioners deploying ACT in real-world settings to ensure adequate human oversight and mechanisms that allow decisions to be contested in a timely and effective manner.

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.

Mina Arzaghi, Alireza Dehghanpour Farashah, Florian Carichon, and Golnoosh Farnadi. Intrinsic meets extrinsic fairness: Assessing the downstream impact of bias mitigation in large language models. *arXiv preprint arXiv:2509.16462*, 2025.

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.

Joy Buolamwini and Timnit Gebru. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Proceedings of the 1st Conference on Fairness, Accountability and Transparency*, volume 81 of *Proceedings of Machine Learning Research*, pp. 77–91. PMLR, 23–24 Feb 2018.

Tianqi Chen and Carlos Guestrin. XGBOOST: A scalable tree boosting system. In *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, pp. 785–794, 2016.

Weize Chen, Yusheng Su, Jingwei Zuo, Cheng Yang, Chenfei Yuan, Chi-Min Chan, Heyang Yu, Yaxi Lu, Yi-Hsin Hung, Chen Qian, Yujia Qin, Xin Cong, Ruobing Xie, Zhiyuan Liu, Maosong Sun, and Jie Zhou. Agentverse: Facilitating multi-agent collaboration and exploring emergent behaviors. In *International Conference on Learning Representations*, volume 2024, pp. 20094–20136, 2024.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, volume 1 (long and short papers)*, pp. 4171–4186, 2019.

Deysi. Spam detection dataset, 2023. URL `https://huggingface.co/datasets/Deysi/spam-detection-dataset`. Accessed: August 10, 2026.

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

European Parliament and Council of the European Union. Regulation (eu) 2024/1689 of the european parliament and of the council of 13 june 2024 laying down harmonised rules on artificial intelligence (artificial intelligence act). Official Journal of the European Union, L series, 12 July 2024, 2024. `http://data.europa.eu/eli/reg/2024/1689/oj`.

Arsene Fansi Tchango, Rishab Goel, Zhi Wen, Julien Martel, and Joumana Ghosn. Ddxplus: A new dataset for automatic medical diagnosis. In *Advances in Neural Information Processing Systems*, volume 35, pp. 31306–31318. Curran Associates, Inc., 2022.

Jerome H. Friedman and Bogdan E. Popescu. Predictive learning via rule ensembles. *The Annals of Applied Statistics*, 2(3):916 – 954, 2008. doi: 10.1214/07-AOAS148.

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.

Peter Jackson. *Introduction to expert systems*. Addison-Wesley Pub. Co., Reading, MA, 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 LLM hallucination via self reflection. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pp. 1827–1843, Singapore, 2023b. Association for Computational Linguistics. doi: 10.18653/v1/2023.findings-emnlp.123.

Omar Khattab, Arnav Singhvi, Paridhi Maheshwari, Zhiyuan Zhang, Keshav Santhanam, Sri Vardhamanan A, Saiful Haq, Ashutosh Sharma, Thomas Joshi, Hanna Moazam, Heather Miller, Matei Zaharia, and Christopher Potts. DSPy: Compiling declarative language model calls into state-of-the-art pipelines. In *International Conference on Learning Representations*, pp. 54928–54958, 2024.

Wojciech Kryscinski, Bryan McCann, Caiming Xiong, and Richard Socher. Evaluating the factual consistency of abstractive text summarization. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 9332–9346, Online, 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.emnlp-main.750.

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.

Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. RoBERTa: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*, 2019.

Kristian Lum and William Isaac. To predict and serve? *Significance*, 13(5):14–19, 2016. doi: https://doi.org/10.1111/j.1740-9713.2016.00960.x.

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.

Andrew L. Maas, Raymond E. Daly, Peter T. Pham, Dan Huang, Andrew Y. Ng, and Christopher Potts. Learning word vectors for sentiment analysis. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pp. 142–150, Portland, Oregon, USA, 2011. Association for Computational Linguistics. Accessed: August 10, 2026 via `https://huggingface.co/datasets/stanfordnlp/imdb`.

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.

Christoph Molnar, Giuseppe Casalicchio, and Bernd Bischl. Interpretable machine learning – a brief history, state-of-the-art and challenges. In *ECML PKDD 2020 Workshops*, pp. 417–431, Cham, 2020. Springer International Publishing. ISBN 978-3-030-65965-3.

National Institute of Standards and Technology. Artificial intelligence risk management framework (AI RMF 1.0). Technical Report NIST AI 100-1, National Institute of Standards and Technology, Gaithersburg, MD, 2023a. URL `https://www.nist.gov/itl/ai-risk-management-framework`.

National Institute of Standards and Technology. NIST AI RMF playbook. `https://airc.nist.gov/airmf-resources/playbook/`, 2023b.

Ninaa510. Diagnosis-text dataset, 2024. URL `https://huggingface.co/datasets/ninaa510/diagnosis-text`. Accessed: August 10, 2026.

Maxwell Nye, Anders Johan Andreassen, Guy Gur-Ari, Henryk Michalewski, Jacob Austin, David Bieber, David Dohan, Aitor Lewkowycz, Maarten Bosma, David Luan, Charles Sutton, and Augustus Odena. Show your work: Scratchpads for intermediate computation with language models. *arXiv preprint arXiv:2112.00114*, 2021.

Ziad Obermeyer, Brian Powers, Christine Vogeli, and Sendhil Mullainathan. Dissecting racial bias in an algorithm used to manage the health of populations. *Science*, 366(6464):447–453, 2019. doi: 10.1126/science.aax2342.

OECD. Recommendation of the council on artificial intelligence. OECD Legal Instruments, OECD/LEGAL/0449, 2019. Amended 2024. `https://oecd.ai/en/ai-principles`.

Cathy O'Neil. *Weapons of Math Destruction.* Penguin Books, 2017.

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.* Springer, 1994.

Timo Schick, Jane Dwivedi-Yu, Roberto Dessi, Roberta Raileanu, Maria Lomeli, Eric Hambro, Luke Zettlemoyer, Nicola Cancedda, and Thomas Scialom. Toolformer: Language models can teach themselves to use tools. In *Advances in Neural Information Processing Systems*, volume 36, pp. 68539–68551. Curran Associates, Inc., 2023.

Renee Shelby, Shalaleh Rismani, Kathryn Henne, AJung Moon, Negar Rostamzadeh, Paul Nicholas, N'Mah Yilla-Akbari, Jess Gallegos, Andrew Smart, Emilio Garcia, and Gurleen Virk. Sociotechnical harms of algorithmic systems: Scoping a taxonomy for harm reduction. In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society*, AIES '23, pp. 723–741. Association for Computing Machinery, 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. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 4222–4235. Association for Computational Linguistics, 2020.

Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. Reflexion: language agents with verbal reinforcement learning. In A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine (eds.), *Advances in Neural Information Processing Systems*, volume 36, pp. 8634–8652. Curran Associates, Inc., 2023.

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

D Slack, SA Friedler, CD Roy, and C Scheidegger. Assessing the local interpretability of machine learning models. In *NeurIPS Workshop on Human-Centric Machine Learning (HCML)*, 2019.

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

Houman Talebzadeh, Sanda Mandutianu, and Christian F Winner. Countrywide loan-underwriting expert system. *AI magazine*, 16(1):51–64, 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.

Gaurav Topre. Bank customer churn dataset. `https://www.kaggle.com/datasets/gauravtopre/bank-customer-churn-dataset`, 2025. Accessed: August 10, 2026.

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.*, 31:841–887, 2018.

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, brian ichter, Fei Xia, Ed Chi, Quoc V Le, and Denny Zhou. Chain-of-thought prompting elicits reasoning in large language models. In *Advances in Neural Information Processing Systems*, volume 35, pp. 24824–24837, 2022.

WHO. World health organization - tuberculosis fact-sheet, 2024. URL https://www.who.int/news-room/fact-sheets/detail/tuberculosis. Accessed: August 10, 2026.

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, Pan Lu, Zhi Huang, Carlos Guestrin, and James Zou. Optimizing generative AI by backpropagating language model feedback. *Nature*, 639(8055):609–616, 2025. doi: 10.1038/s41586-025-08661-4.

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*, 2023.

A Computational Cost Analysis

Computational cost analysis. Figure 4 compares the training and inference costs of DSPy, TextGrad and two ACT configurations on the DIAGNO task. We report wall-clock time and total token usage (input plus output). During training, all methods have access to the full training set and at inference time all methods are evaluated on the full test set. All experiments were conducted using Gemma-3-4B-IT on a single H200 GPU with 8 threads per program and results are averaged over five independent runs.

The observed differences reflect the underlying program-construction strategies of each method. TextGrad exhibits relatively high wall-clock time despite lower token usage, as it relies on fewer but substantially longer prompts. DSPy, by contrast, evaluates a larger number of medium-sized prompts derived from few-shot reasoning traces. ACT issues many short prompts to estimate Gini impurity at each node, complemented by a smaller number of longer prompts for semantic feedback and node-level optimization. In its cheapest configuration, ACT is competitive with both baselines in overall computational cost. Increasing tree depth raises training cost but yields higher accuracy while preserving interpretability. At inference time, TextGrad is the most economical, while ACT remains within the same order of magnitude and provides higher accuracy with interpretable decision paths. DSPy shows similar wall-clock time to ACT but incurs higher token usage due to the substantial size of its few-shot prompts.

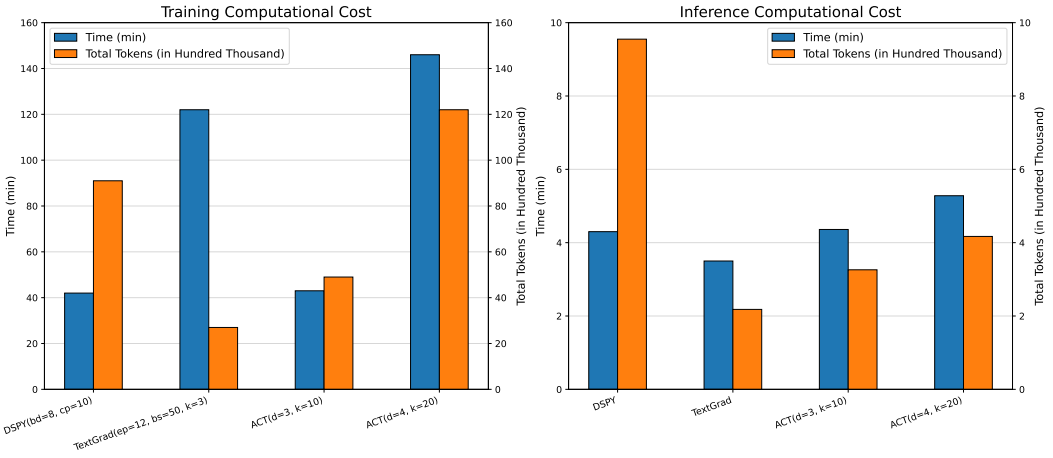

Figure 4: **Training and inference computational cost comparison on DIAGNO.** Each plot reports the average wall-clock time and token usage over 5 runs for TextGrad, DSPy (BFSR, 8-shot), and ACT under both the cheapest (d = 3, k = 10) and most expensive (d = 4, k = 20) configurations, using Gemma-3-4B on a single H200 GPU with 8 threads per program. Training cost reflects program construction dynamics: TextGrad processes fewer but very long prompts; DSPy evaluates 13 medium-sized candidate programs derived from few-shot traces; and ACT issues many short node-level prompts for Gini estimation supplemented with a smaller number of larger prompts for semantic feedback and node optimization. Inference cost shows that TextGrad is the most economical, while ACT remains within the same order of magnitude and provides higher accuracy with more interpretable outputs. DSPy exhibits similar wall-clock time to ACT but requires more tokens, largely due to the substantial size of the few-shot prompts.

B Additional Implementation Details of the ACT Algorithm

B.1 Illustrative Example: Prompt Refinement via Semantic Feedback

We describe the iterative prompt refinement procedure used in ACT, using an illustrative case from the `DIAGNO` 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 D_L) and those routed to the `no` branch (denoted D_R).
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., 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 fever?"*, better aligned with the semantic distinctions in the data.

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

Gini impurity: 0.495

Group predicted "YES" (s_{D_L}):

- **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" (s_{D_R}):

- **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" (s_{D_L}):

- **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" (s_{D_R}):

- **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

B.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. In all experiments reported in the paper, we set $L = 2$.

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 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 `<...>`.

B.3 Details on Datasets

We use five publicly available text classification datasets spanning diverse domains. DIAGNO, SPAM, and JAILBREAK are included because they represent tasks for which structured reasoning is expected to be beneficial, aligning with the types of problems ACT is designed to address. BANKCHURN, a tabular dataset serialized into text, is included because it relies less on pretraining-aligned task semantics and more on learning dataset-specific relationships between attributes and the churn label. IMDB, a widely used benchmark for sentiment classification, is added both to broaden the evaluation and to facilitate comparison with prior work. All datasets are relatively small and mostly balanced,

reflecting the low-data prompt-optimization regime in which ACT and comparable methods are typically evaluated (Yuksekgonul et al., 2025).

- **DIAGNO** (Ninaa510, 2024): a medical diagnosis dataset (tuberculosis vs. allergic sinusitis), a derived dataset based on DDXPlus (Fansi Tchango et al., 2022). We randomly sampled a balanced subset using a fixed random seed, comprising 600 training, 100 validation, and 600 test samples, each with a 50/50 class split between tuberculosis and allergic sinusitis.
- **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).
- **BANKCHURN** (Topre, 2025): a customer–attrition prediction dataset originally provided in tabular form. Following (Arzaghi et al., 2025), we serialize each record into a short textual profile describing the customer attributes. We constructed a balanced subset with 600 training, 100 validation, and 600 test samples, following the same procedure as for DIAGNO.
- **IMDB** (Maas et al., 2011): a sentiment analysis dataset consisting of movie reviews labeled as positive or negative. We constructed a balanced subset with 600 training, 100 validation, and 600 test samples, following the same procedure as for DIAGNO.

C 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:

$$\mathcal{D}_L^x = \{(x_i, y_i) \in \mathcal{D}^x \mid x_{i,j} \leq s\}, \quad \mathcal{D}_R^x = \{(x_i, y_i) \in \mathcal{D}^x \mid x_{i,j} > s\}. \tag{1}$$

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:

$$\delta_j^s = \frac{|\mathcal{D}_L^x(j,s)|}{|\mathcal{D}^x|} \cdot \mathrm{Gini}(\mathcal{D}_L^x(j,s)) + \frac{|\mathcal{D}_R^x(j,s)|}{|\mathcal{D}^x|} \cdot \mathrm{Gini}(\mathcal{D}_R^x(j,s)), \tag{2}$$

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

$$\text{Gini}(D) = 1 - \sum_{k \in \mathcal{Y}} p_k^2, \tag{3}$$

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).

D Additional Qualitative Results

D.1 Qualitative analysis of the questions generated by ACT

Table 5: Symptoms comparison between ACT and medical sources.

Symptoms 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		

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 tuberculosis diagnosis. These sources include websites of the World Health Organization[1] and the British National Health Service[2], and a literature review on tuberculosis diagnosis (Storla et al., 2008).

Table 5 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,

[1] https://www.who.int/news-room/fact-sheets/detail/tuberculosis

[2] https://www.nhs.uk/conditions/tuberculosis-tb/

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 stem from the binary setting considered, which opposes tuberculosis to allergic sinusitis: symptoms such as loss of appetite appear in both classes (see Figure 1) and are therefore not discriminative for this particular split, even though they are clinically relevant for tuberculosis. Overall, this tends to show that ACT is indeed able to identify relevant symptoms to make its predictions, showcasing its utility and reliability.

D.2 Agentic Classification Tree for SPAM and JAILBREAK Dataset

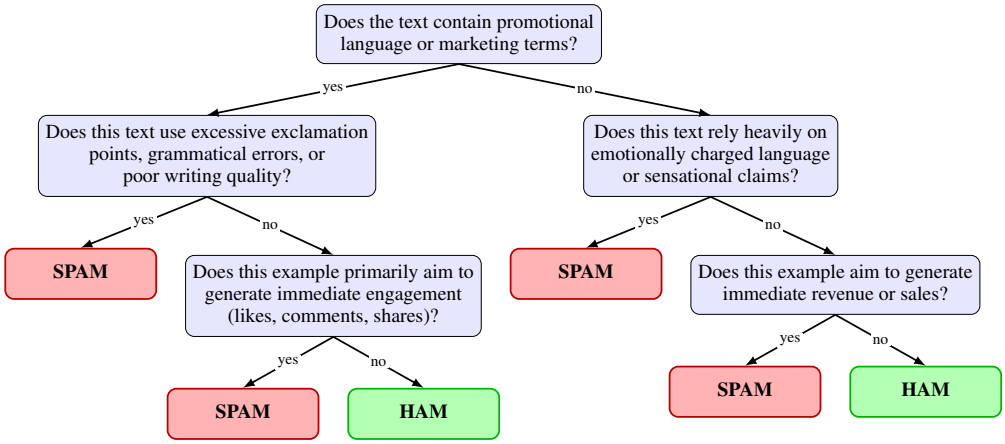

Figure 5: Decision tree of depth 3 generated by the Gemma3-4B model for spam email classification. The tree distinguishes between spam and legitimate email (ham) through hierarchical semantic questions about promotional content, writing quality, and intent. Each internal node represents a binary question optimized through the ACT framework to maximize class separation, with leaf nodes indicating the final classification based on the majority class of training examples.

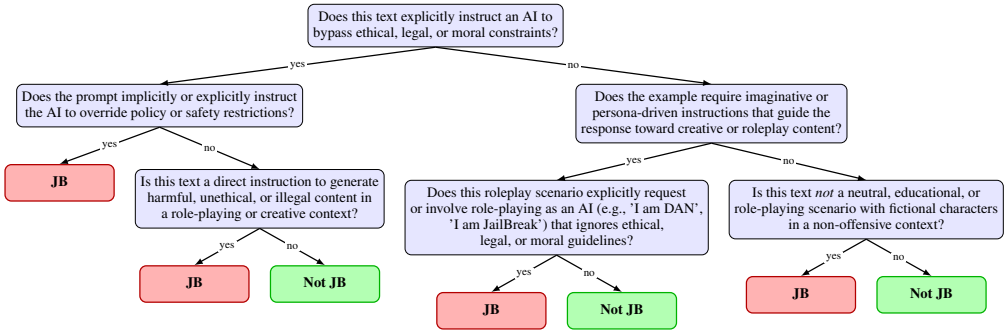

Figure 6: Decision tree of depth 3 generated by the Qwen3-4b model for jailbreak prompt classification. The tree recursively partitions inputs through binary natural-language questions optimized to distinguish between jailbreak attempts (JB) and legitimate prompts (Not JB). Each internal node contains 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.

D.3 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 7, and in Figure 8 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.

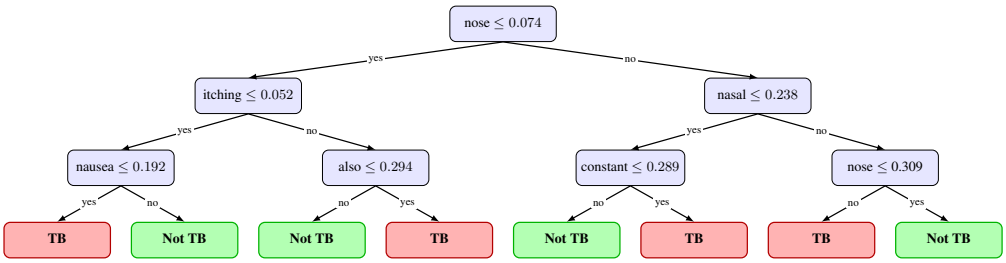

Figure 7: Decision tree of depth 3 generated by training a CART model after performing a TF-IDF preprocessing step on the DIAGNO dataset.

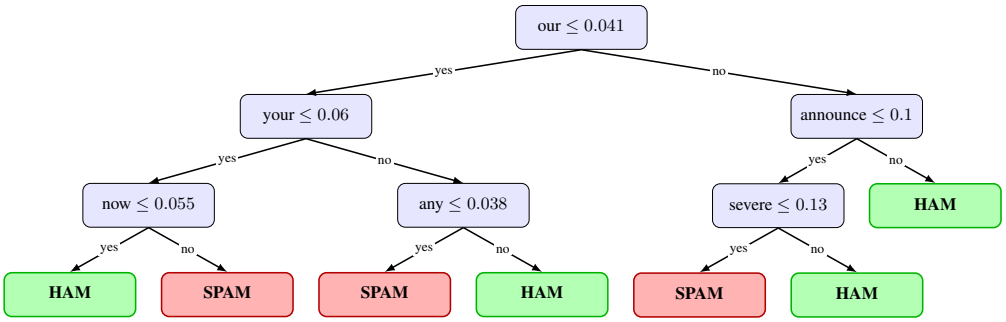

Figure 8: Decision tree of depth 3 generated by training a CART model after performing a TF-IDF preprocessing step on the SPAM dataset.

E Additional Results and Baseline Implementation Details

E.1 Additional Baselines and ACT Configuration Results

To complement the results shown in Table 2 we expand this table in Table 6 and Table 7, which include additional results for the ACT $(d = 3, k = 20)$ and $(d = 4, k = 10)$ configurations, as well as strong but not interpretable baselines:

- **TF-IDF + XGBoost**: a gradient-boosted tree ensemble trained on TF-IDF representations of the texts (Chen & Guestrin, 2016). We tune both the maximum tree depth and the number of trees (see Table 8), and report the best-performing configuration per dataset.
- **Fine-tuned BERT and RoBERTa classifiers**: supervised Transformer baselines initialized from pretrained BERT and RoBERTa encoders (Devlin et al., 2019; Liu et al., 2019) and fine-tuned end-to-end on each dataset for text classification. These models provide strong predictive performance references, but do not yield explicit, human-readable decision rules.

E.2 Full Tree-Based Baselines Results

To complement the prompting-based baselines used in the main paper, we report in Table 8 an extended comparison with purely symbolic lexical models. These include TF–IDF + CART with depths ranging from 2 to 5, and TF–IDF + XGBoost with multiple depth and tree-count configurations. These models provide a transparent, non-LLM control group and clarify the performance gap between ACT and traditional supervised baselines across all five datasets.

Table 6: Comparison of classification methods across the first three datasets and four LLMs. Results show accuracy and F1-score on the test set (in %) for nine baselines, as well as the proposed ACT with varying hyperparameters (d: tree depth, k: optimization steps per node).

Method	Gemma3 4b		GPT-4.1 Nano		GPT-4.1 Mini		Qwen3 4b		Avg	
	Acc	F1	Acc	F1	Acc	F1	Acc	F1	Acc	F1
DIAGNO										
CoT (0-shot)	61.5	69.5	63.3	54.7	61.2	51.4	63.2	41.7	62.3	54.3
DSPy (BFSR, 8 demos)	64.0	68.2	68.7	61.2	64.5	59.4	63.3	54.9	65.1	60.9
TEXTGRAD	63.3	56.7	64.5	53.0	65.8	55.5	64.0	60.0	64.4	56.3
TF-IDF + CART (d=3)	−−	−−	−−	−−	−−	−−	−−	−−	78.8	78.9
TF-IDF + XGBoost (d=4, nt=200)	−−	−−	−−	−−	−−	−−	−−	−−	83.0	83.5
BERT	−−	−−	−−	−−	−−	−−	−−	−−	90.3	91.1
RoBERTa	−−	−−	−−	−−	−−	−−	−−	−−	89.8	90.4
Con. tag.	65.5	63.9	59.3	58.1	75.3	78.6	58.8	64.7	64.7	66.3
Rule Fit	−−	−−	−−	−−	−−	−−	−−	−−	80.2	79.3
ACT (d=3, k=10)	65.3	66.5	66.8	67.0	77.3	75.2	64.8	61.3	68.6	67.5
ACT (d=3, k=20)	67.0	67.9	67.2	69.0	76.2	76.9	65.5	61.8	69.0	68.9
ACT (d=4, k=10)	66.3	**71.1**	**70.5**	70.1	74.3	70.1	66.7	66.8	69.5	69.5
ACT (d=4, k=20)	**68.3**	70.8	70.3	**70.8**	**82.3**	**81.8**	**70.3**	**73.5**	72.8	74.2
SPAM										
CoT (0-shot)	73.3	78.6	95.5	95.3	95.7	95.5	93.2	93.6	89.4	90.8
DSPy (BFSR, 8 demos)	95.0	95.2	98.3	98.3	98.2	98.2	97.2	97.2	97.2	97.2
TEXTGRAD	95.0	94.7	97.3	97.3	96.7	96.6	96.0	96.0	96.3	96.2
TF-IDF + CART (d=5)	−−	−−	−−	−−	−−	−−	−−	−−	92.7	93.1
TF-IDF + XGBoost (d=3, nt=200)	−−	−−	−−	−−	−−	−−	−−	−−	96.3	96.3
BERT	−−	−−	−−	−−	−−	−−	−−	−−	99.5	99.5
RoBERTa	−−	−−	−−	−−	−−	−−	−−	−−	99.3	99.3
Con. tag.	83.2	85.5	64.2	45.8	82.3	85.0	62.3	39.6	73.0	64.0
Rule Fit	−−	−−	−−	−−	−−	−−	−−	−−	97.0	96.9
ACT (d=2, k=5)	95.8	95.8	98.7	98.6	98.2	98.2	96.7	96.6	97.4	97.3
ACT (d=2, k=10)	96.5	96.6	98.8	98.8	99.2	99.2	98.5	98.5	98.3	98.3
ACT (d=3, k=5)	98.2	98.1	97.0	96.9	99.0	99.0	98.5	98.5	98.2	98.1
ACT (d=3, k=10)	**98.5**	**98.5**	**99.2**	**99.2**	**99.5**	**99.5**	**98.8**	**98.8**	99.0	99.0
JAILBREAK										
CoT (0-shot)	77.9	82.2	85.5	83.5	93.6	93.3	81.1	77.7	84.5	84.2
DSPy (BFSR, 8 demos)	89.6	89.3	91.6	**91.6**	94.8	94.8	85.5	83.9	90.4	89.9
TEXTGRAD	**91.6**	**91.3**	90.4	90.3	94.8	95.1	88.8	90.0	91.4	91.7
TF-IDF + CART (d=4)	−−	−−	−−	−−	−−	−−	−−	−−	92.8	92.4
TF-IDF + XGBoost (d=2, nt=200)	−−	−−	−−	−−	−−	−−	−−	−−	96.8	96.8
BERT	−−	−−	−−	−−	−−	−−	−−	−−	98.4	98.4
RoBERTa	−−	−−	−−	−−	−−	−−	−−	−−	99.6	99.6
Con. tag.	82.7	82.2	58.6	31.8	82.7	79.8	76.3	70.1	75.1	66.0
Rule Fit	−−	−−	−−	−−	−−	−−	−−	−−	96.8	96.8
ACT (d=3, k=10)	82.3	82.7	85.5	85.1	95.2	95.5	82.7	82.6	86.4	86.5
ACT (d=3, k=20)	84.3	86.5	91.2	90.5	96.0	96.6	87.6	86.6	89.8	90.1
ACT (d=4, k=10)	90.0	90.1	91.6	91.0	97.2	97.2	85.1	83.9	91.0	90.6
ACT (d=4, k=20)	90.8	90.0	**92.0**	91.5	**98.8**	**98.8**	**92.0**	**92.7**	93.4	93.3

Table 7: Comparison of classification methods across the two remaining datasets and four LLMs. Results show accuracy and F1-score on the test set (in %) for nine baselines, as well as the proposed ACT with varying hyperparameters (d: tree depth, k: optimization steps per node).

Method	Gemma3 4b		GPT-4.1 Nano		GPT-4.1 Mini		Qwen3 4b		Avg	
	Acc	F1	Acc	F1	Acc	F1	Acc	F1	Acc	F1
BANKCHURN										
CoT (0-shot)	47.8	48.1	48.5	25.9	52.8	60.5	50.2	37.3	49.8	43.0
DSPy (BFSR, 8 demos)	52.8	55.6	52.3	42.8	57.2	58.5	54.0	52.2	54.1	52.3
TEXTGRAD	52.0	60.0	53.3	60.3	55.3	61.0	52.8	61.4	53.4	60.7
TF-IDF + CART (d=5)	– –	– –	– –	– –	– –	– –	– –	– –	63.2	66.8
TF-IDF + XGBoost (d=2, nt=200)	– –	– –	– –	– –	– –	– –	– –	– –	56.7	48.8
BERT	– –	– –	– –	– –	– –	– –	– –	– –	71.7	71.6
RoBERTa	– –	– –	– –	– –	– –	– –	– –	– –	72.3	72.3
Con. tag.	57.2	60.2	52.0	59.8	59.7	60.8	54.3	61.6	55.8	60.6
Rule Fit	– –	– –	– –	– –	– –	– –	– –	– –	63.0	59.8
ACT (d=3, k=10)	56.3	42.3	58.8	67.4	60.5	64.6	57.5	62.7	58.3	59.3
ACT (d=3, k=20)	58.0	47.8	**67.0**	67.5	**69.2**	65.3	**62.7**	**63.7**	64.2	61.1
ACT (d=4, k=10)	61.7	65.3	61.7	60.9	69.0	65.0	60.8	60.2	63.3	62.9
ACT (d=4, k=20)	**65.2**	**70.7**	65.0	**69.1**	68.7	**68.6**	62.0	60.0	65.2	67.1
IMDB										
CoT (0-shot)	92.3	91.9	94.2	94.1	95.7	95.6	94.2	94.0	94.1	93.9
DSPy (BFSR, 8 demos)	**93.7**	**93.6**	**95.5**	**95.5**	**96.8**	**96.8**	94.8	94.8	95.2	95.2
TEXTGRAD	93.2	92.9	95.2	95.1	95.8	95.8	94.7	94.5	94.7	94.6
TF-IDF + CART (d=3)	– –	– –	– –	– –	– –	– –	– –	– –	66.8	72.8
TF-IDF + XGBoost (d=2, nt=200)	– –	– –	– –	– –	– –	– –	– –	– –	76.2	76.7
BERT	– –	– –	– –	– –	– –	– –	– –	– –	90.2	90.1
RoBERTa	– –	– –	– –	– –	– –	– –	– –	– –	92.7	92.7
Con. tag.	74.8	76.9	56.2	31.7	50.0	66.7	51.3	63.8	58.1	59.8
Rule Fit	– –	– –	– –	– –	– –	– –	– –	– –	75.2	73.7
ACT (d=3, k=10)	93.2	93.2	92.8	92.8	95.7	95.7	93.8	93.9	93.9	93.9
ACT (d=3, k=20)	**93.7**	**93.6**	94.7	94.7	96.3	96.2	**95.5**	**95.5**	95.1	95.0

Table 8: Performance of TF-IDF + CART and TF-IDF + XGBoost across all datasets.

Dataset / Split	CART d=2	CART d=3	CART d=4	CART d=5	XGB (2,50)	XGB (2,200)	XGB (3,50)	XGB (3,200)	XGB (4,50)	XGB (4,200)
DIAGNO – Train	82.0	84.3	86.5	**88.8**	86.2	94.7	89.2	98.7	93.5	**99.8**
DIAGNO – Test	77.5	**78.8**	73.7	76.3	81.0	81.5	81.8	81.7	82.2	**83.0**
SPAM – Train	87.8	92.3	93.5	**94.7**	97.8	99.7	98.7	99.8	99.2	**100.0**
SPAM – Test	85.8	89.3	90.7	**92.7**	94.3	96.2	93.8	**96.3**	94.5	96.0
JAILBREAK – Train	88.3	91.1	93.1	**94.3**	95.6	97.9	96.5	98.5	97.1	**99.0**
JAILBREAK – Test	86.4	92.4	**92.8**	92.0	95.2	**96.8**	95.2	96.4	95.2	**96.8**
BANKCHURN – Train	59.7	63.0	63.7	**66.7**	71.5	84.5	79.3	90.0	84.2	**93.2**
BANKCHURN – Test	62.7	62.8	62.8	**63.2**	53.2	**56.7**	54.3	53.7	54.8	55.3
IMDB – Train	68.2	71.0	71.8	**76.2**	83.5	97.3	89.0	99.3	94.7	**99.8**
IMDB – Test	66.8	**66.8**	66.5	**66.8**	72.5	**76.2**	73.3	74.8	74.8	75.5

E.3 FULL-DATASET EVALUATION RESULTS

In this subsection, we report results obtained by running all prompting-based baselines, TF-IDF + CART and ACT on the full training datasets for the BANKCHURN, SPAM and DIAGNO tasks.

Full Datasets. Table 9 shows performance for ACT and baseline methods on the full datasets. For BANKCHURN, the dataset contains 7'731 training samples (20.57% labeled "yes") and 1'969 test samples (19.55% "yes"). A validation set of 300 samples was constructed by random sampling from the training set while preserving the class distribution. When using this validation split, the effective training size is 7'431 samples. For SPAM, the training set has 7'766 (50.46% "yes") and the test set has 2'725 (50.46% "yes") samples. The validation set of 409 samples was constructed by random sampling from the training set while preserving the class distribution. For DIAGNO, the full test set containing the Tuberculosis (TB) and Allergic sinusitis samples contains 997 (33.5% TB, resp. "yes") samples and the train set contains the maximum possible amount of samples (1994 with 33.5% TB, resp. "yes") such that it matches the test set distribution. The original train set distribution would be 50% of each class, but we are interested in evaluating the in-distribution test accuracy. In Table 9 the TF-IDF + CART configuration with the best test set accuracy among depths $3-10$ is reported and additionally Table 10 reports TF-IDF + CART results across all tested tree depths.

Results. On the heavily imbalanced BANKCHURN dataset, several baselines (e.g., TextGrad and DSPy with GPT-4.1-Nano) collapse to majority-class prediction, yielding F1-scores of 0, while ACT with Gemma-3-4b achieves the highest accuracy among all methods alongside an acceptable F1-score. On SPAM, where the label distribution mirrors the random subset used in the main paper, prompting-

Table 9: **Results for ACT and baselines** (CoT, TextGrad, DSPY and TF-IDF+CART) methods **on the full (and therefore imbalanced) datasets** for BANKCHURN, SPAM and DIAGNO. Metrics are test set accuracy and F1-score on the test set. Depending on the class imbalance, different ACT hyperparameter configurations (d for maximum depth and k for number of optimization steps per node) were evaluated.

Dataset	Model	Method	Train Acc	Test Acc	Test F1
DIAGNO	Gemma	CoT	49.2	50.1	55.0
		TextGrad	68.6	67.0	54.2
		DSPy	67.6	66.5	57.9
		ACT ($d = 4, k = 10$)	73.6	72.6	41.5
		ACT ($d = 5, k = 10$)	74.8	**72.8**	52.7
		ACT ($d = 8, k = 10$)	81.4	71.9	**59.4**
	Nano	CoT	69.0	68.6	48.1
		TextGrad	69.4	68.6	48.1
		DSPy	73.8	70.7	45.9
		ACT ($d = 4, k = 10$)	74.9	72.7	54.1
		ACT ($d = 5, k = 10$)	75.8	**73.7**	54.2
		ACT ($d = 8, k = 10$)	80.2	73.4	**56.1**
	–	TF-IDF + CART (best, d=7)	90.1	85.5	76.7
SPAM	Gemma	CoT	74.9	75.0	80.1
		TextGrad	97.5	97.2	97.2
		DSPy	97.1	97.4	97.5
		ACT ($d = 2, k = 10$)	96.5	96.3	96.4
		ACT ($d = 3, k = 10$)	97.9	97.6	97.6
		ACT ($d = 4, k = 10$)	98.7	97.9	98.0
		ACT ($d = 5, k = 10$)	99.1	**98.8**	**98.8**
	Nano	CoT	96.1	95.6	95.5
		TextGrad	98.1	98.2	98.2
		DSPy	98.1	98.2	98.2
		ACT ($d = 2, k = 10$)	99.6	99.5	99.5
		ACT ($d = 3, k = 10$)	99.6	99.5	99.5
		ACT ($d = 4, k = 10$)	99.7	**99.6**	**99.6**
		ACT ($d = 5, k = 10$)	99.6	99.5	99.5
	–	TF-IDF + CART (best, d=10)	95.3	94.9	95.2
BANKCHURN	Gemma	CoT	48.5	47.3	27.3
		TextGrad	78.0	78.9	2.4
		DSPy	55.5	56.0	33.2
		ACT ($d = 4, k = 20$)	75.9	75.0	22.9
		ACT ($d = 8, k = 10$)	82.6	**82.7**	**40.6**
	Nano	CoT	66.7	67.8	19.1
		TextGrad	79.4	80.5	0.0
		DSPy	79.4	80.3	0.0
		ACT ($d = 4, k = 20$)	80.2	80.7	9.1
		ACT ($d = 8, k = 10$)	80.7	**80.7**	**12.4**
	–	TF-IDF + CART (best, d=3)	79.7	80.7	4.5

based methods achieve results comparable to Table 2, confirming the effectiveness of prompt-based methods in limited data settings. TF-IDF+CART improves modestly with the full dataset but remains below the prompting-based baselines. On DIAGNO, TF-IDF+CART outperforms prompt-based baselines built on small LLMs, consistent with the random subset results in Table 2. However, additional results (not shown in Table 9 because only ACT and no baselines were evaluated) on ACT with GPT-4.1-Mini $(d = 8, k = 10)$ achieve 81.4% accuracy and an F1-score of 68.4, demonstrating that stronger backbone models close this gap. Crucially, unlike black-box prompt-based baselines or TF-IDF+CART, ACT produces interpretable and auditable decision paths. More broadly, larger training sets allow ACT to benefit from deeper trees without the overfitting observed on small subsets and performance differences between similarly-sized LLMs (e.g. Gemma-3-4b vs. GPT-4.1-Nano) become more pronounced on full datasets, particularly in F1-score on imbalanced tasks. Table 10 shows that increasing tree depth consistently improves training accuracy, yet TF-IDF+CART overfits on BANKCHURN and DIAGNO. On the heavily imbalanced BANKCHURN task, the best test accuracy is achieved at depth 3 but at the cost of a near-zero F1-score, while depth 10 recovers F1 at the expense of accuracy, highlighting the difficulty of class imbalance.

Table 10: **TF-IDF + CART results across all evaluated tree depths** on the full BANKCHURN, SPAM and DIAGNO datasets.

Dataset (Train/Test Set Size)	**Depth**	**Train Acc (%)**	**Test Acc (%)**	**Test F1 (%)**
BANKCHURN (7731/1969)	3	79.7	**80.7**	4.5
	4	80.0	79.8	3.4
	5	80.4	80.2	4.4
	6	80.8	80.0	5.7
	7	81.3	79.4	8.6
	8	82.1	78.4	10.2
	9	82.7	79.6	10.7
	10	**83.3**	77.6	**11.6**
SPAM (7766/2725)	3	90.9	90.2	91.2
	4	92.4	92.0	92.6
	5	93.5	93.4	93.9
	6	94.3	94.2	94.6
	7	94.9	94.8	95.1
	8	95.3	94.9	95.1
	9	95.3	94.9	95.2
	10	**95.3**	**94.9**	**95.2**
DIAGNO (1994/997)	3	82.4	80.0	71.0
	4	83.2	80.9	72.6
	5	84.9	81.6	68.3
	6	88.0	84.6	75.6
	7	90.1	**85.5**	**76.7**
	8	91.7	83.9	74.3
	9	92.7	83.6	73.3
	10	**94.1**	83.5	73.9

F INTERPRETABILITY STATISTICS FOR ACT

In Table 11 we report two interpretability indicators for ACT across all datasets and ACT configurations evaluated: the average question length (in characters) and the average path length (number of nodes traversed during inference). These values are averaged over the four LLMs used in our experiments. Across datasets, average question lengths range from 78 to 155 characters, reflecting concise node-level prompts, while average path lengths show that predictions typically require only about three nodes, even for depth-4 trees. We also observe a mild trade-off in which datasets with longer questions (e.g., JAILBREAK and SPAM) yield shorter paths, indicating that ACT can consolidate multiple relevant features into a single informative query. For reference, the generic seed question *"Based on the provided example, does it belong to the positive class? (yes/no)"* has a length of 77 characters.

Table 11: Interpretability statistics across datasets.

ACT (d, k)	Avg. Q. Len	Avg. Path
(3, 10)	92.00	2.35
(3, 20)	93.15	2.51
(4, 10)	78.05	3.03
(4, 20)	80.53	3.06

(a) DIAGNO

ACT (d, k)	Avg. Q. Len	Avg. Path
(2, 5)	118.00	1.33
(2, 10)	142.88	1.26
(3, 5)	94.18	1.36
(3, 10)	122.63	1.14

(b) SPAM

ACT (d, k)	Avg. Q. Len	Avg. Path
(3, 10)	136.00	1.98
(3, 20)	155.00	2.13
(4, 10)	142.00	2.41
(4, 20)	150.33	2.57

(c) JAILBREAK

ACT (d, k)	Avg. Q. Len	Avg. Path
(3, 10)	90.50	2.92
(3, 20)	89.90	2.93
(4, 10)	82.20	3.10
(4, 20)	89.77	2.96

(d) BANKCHURN

ACT (d, k)	Avg. Q. Len	Avg. Path
(3, 10)	81.30	1.97
(3, 20)	95.80	2.28

(e) IMDB

G ADDITIONAL ABLATION STUDIES AND SCALING ANALYSIS

G.1 ABLATION ON THE NUMBER OF FEEDBACK SAMPLES m

We conduct an ablation study on the number of labeled samples m used to generate semantic feedback during node-level prompt optimization. We evaluate $m \in \{5, 10, 20, 50\}$ on the DIAGNO dataset using GPT-4.1-nano for two ACT configurations: $(d = 3, k = 10)$ and $(d = 4, k = 20)$. The results in Tables 12 and 13 show how varying m affects training performance, test metrics, and token usage. As comparison, the results for DSPy are also shown.

Table 12: Effect of varying m for ACT $(d = 3, k = 10)$ on DIAGNO with GPT-4.1-nano.

m	Train Acc ↑	Test Acc ↑	Test F1 ↑	Tokens ($\times 100k$) ↓
5	69.8	67.0	67.6	**52.4**
10	71.2	67.5	71.3	53.5
20	71.3	**69.5**	**71.7**	54.2
50	71.8	66.8	67.0	55.2
DSPy	71.5	68.7	61.2	152.6

Table 13: Effect of varying m for ACT $(d = 4, k = 20)$ on DIAGNO with GPT-4.1-nano.

m	Train Acc ↑	Test Acc ↑	Test F1 ↑	Tokens ($\times 100k$) ↓
5	71.7	65.8	65.9	**124.7**
10	74.2	68.7	70.6	135.0
20	74.5	70.0	**71.3**	137.6
50	74.8	**70.3**	70.8	143.8
DSPy	71.5	68.7	61.2	152.6

G.2 Model Scaling and Cost-Normalized Analysis

To assess how model strength affects different prompting strategies, we conducted an additional scaling study using GPT-4.1. We evaluate CoT, DSPy, TextGrad, and two ACT configurations on the DIAGNO and BANKCHURN datasets, as these are the most challenging tasks. As shown in Table 14, depending on the task, stronger models do not consistently improve CoT, DSPy, or TextGrad, whereas ACT benefits from increased model capacity. These results indicate that ACT more effectively leverages the enhanced reasoning abilities of more powerful models.

Table 14: Performance scaling of CoT, DSPy, TextGrad, and ACT when moving from GPT-4.1-nano to GPT-4.1-mini to GPT-4.1. Values shown are test-set accuracies in %. ACT benefits more consistently from stronger models.

Method	GPT-4.1-nano	GPT-4.1-mini	GPT-4.1
DIAGNO			
CoT	63.3	61.2	62.2
DSPy	**68.7**	64.5	68.2
TextGrad	64.5	65.8	64.8
ACT (3, 10)	66.8	**77.3**	80.2
ACT (3, 20)	67.2	76.2	**81.8**
BANKCHURN			
CoT	48.5	52.8	57.2
DSPy	52.3	57.2	59.5
TextGrad	53.3	55.3	58.2
ACT (3, 10)	58.8	60.5	69.2
ACT (3, 20)	**67.0**	**69.2**	**72.5**

Figure 3 in the main paper visualizes the accuracy–cost trade-off across methods and model sizes, highlighting how performance scales with increasing computational expense. This section further reports token usage, monetary cost, test accuracy, and cost-normalized accuracy for each method, with Table 15 summarizing these metrics across the GPT-4.1 family on DIAGNO to facilitate comparison of efficiency and scalability.

Table 15 shows that TextGrad achieves the highest cost-normalized accuracy in the Nano and Mini settings, while ACT (3,10) remains competitive—ranking just behind TextGrad in smaller models and performing strongly on GPT-4.1. The (3,20) variant is less cost-efficient due to higher token usage, though it still slightly outperforms DSPy on Mini and GPT-4.1. As shown in Figure 3 in the main paper, ACT's test accuracy consistently scales with cost across all model sizes, whereas TextGrad and DSPy do not exhibit consistent cost-dependent accuracy gains.

Table 15: Cost normalization and model scaling across the GPT-4.1 family on the DIAGNO dataset.

Methods	Model	Cost ↓	Test Acc. ↑	Total Tokens (x100k) ↓	Cost Norm Acc. ↑
ACT (3, 10)	Nano	1.09	66.8	55.18	61.54
ACT (3, 20)	Nano	2.02	67.2	101.37	33.31
DSPY	Nano	1.85	**68.7**	152.58	37.23
TextGrad	Nano	**0.80**	64.5	**48.06**	**80.93**
ACT (3, 10)	Mini	3.09	**77.3**	48.08	24.98
ACT (3, 20)	Mini	5.58	76.2	87.58	13.65
DSPY	Mini	7.92	64.5	161.28	8.15
TextGrad	Mini	**2.30**	65.8	**36.37**	**28.60**
ACT (3, 10)	GPT-4.1	19.08	80.2	53.71	**4.20**
ACT (3, 20)	GPT-4.1	36.11	**81.8**	102.46	2.26
DSPY	GPT-4.1	37.10	68.2	153.89	1.84
TextGrad	GPT-4.1	**17.76**	64.8	**49.08**	3.65