Decision Trees for Entity Identification: Approximation Algorithms and Hardness Results

15

VENKATESAN T. CHAKARAVARTHY, VINAYAKA PANDIT, and SAMBUDDHA ROY, IBM India Research Lab
PRANJAL AWASTHI, Carnegie Mellon University
MUKESH K. MOHANIA, IBM India Research Lab

We consider the problem of constructing decision trees for entity identification from a given relational table. The input is a table containing information about a set of entities over a fixed set of attributes and a probability distribution over the set of entities that specifies the likelihood of the occurrence of each entity. The goal is to construct a decision tree that identifies each entity unambiguously by testing the attribute values such that the average number of tests is minimized. This classical problem finds such diverse applications as efficient fault detection, species identification in biology, and efficient diagnosis in the field of medicine. Prior work mainly deals with the special case where the input table is binary and the probability distribution over the set of entities is uniform. We study the general problem involving arbitrary input tables and arbitrary probability distributions over the set of entities. We consider a natural greedy algorithm and prove an approximation guarantee of $O(r_K \cdot \log N)$, where N is the number of entities and K is the maximum number of distinct values of an attribute. The value r_K is a suitably defined Ramsey number, which is at most $\log K$. We show that it is NP-hard to approximate the problem within a factor of $\Omega(\log N)$, even for binary tables (i.e., $K = 2$). Thus, for the case of binary tables, our approximation algorithm is optimal up to constant factors (since $r_2 = 2$). In addition, our analysis indicates a possible way of resolving a Ramsey-theoretic conjecture by Erdös.

Categories and Subject Descriptors: F.2.2 [**Analysis of Algorithms and Problem Complexity**]: Nonnumerical Algorithms and Problems

General Terms: Algorithms, Theory

Additional Key Words and Phrases: Approximation algorithms, decision tree, Ramsey numbers

ACM Reference Format:
Chakaravarthy, V. T., Pandit, V., Roy, S., Awasthi, P., and Mohania, M. K. 2011. Decision trees for entity identification: Approximation algorithms and hardness results. ACM Trans. Algor. 7, 2, Article 15 (March 2011), 22 pages.
DOI = 10.1145/1921659.1921661 http://doi.acm.org/10.1145/1921659.1921661

A preliminary version of the article was presented at the ACM Symposium on Principles of Database Systems, [Chakaravarthy et al. 2007].
This work was done while P. Awasthi was at IBM India Research Lab, New Delhi.
Authors' addresses: V. T. Chakaravarthy, V. Pandit, and S. Roy, IBM India Research Lab, 4 Block C, Institutional Area, Vasanth Kunj, New Delhi – 110070, India; email: {vechakra, pvinayak, sambuddha}@in.ibm.com; P. Awasthi, Computer Science Department, Wean Hall 1313, Carnegie Mellon University, Pittsburgh, PA 15213; email: pawasthi@cs.cmu.edu; M. K. Mohania, IBM India Research Lab, 4 Block C, Institutional Area, Vasanth Kunj, New Delhi – 110070, India; email: mkmohania@in.ibm.com.
Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.
© 2011 ACM 1549-6325/2011/03-ART15 $10.00
DOI 10.1145/1921659.1921661 http://doi.acm.org/10.1145/1921659.1921661

1. INTRODUCTION

Decision trees for the purposes of identification and diagnosis have been studied for a long time now [Moret 1982]. Consider a typical medical diagnosis application. A hospital maintains a table containing information about diseases. Each row in the table is a disease and each column is a medical test and the corresponding entry specifies the outcome of the test for a person suffering from the given disease. Some of the medical tests are costly (e.g., MRI scans) and some require few days for the result to be known (e.g., blood cultures). When the hospital receives a new patient whose disease has not been identified, it would like to determine the shortest sequence of tests which can unambiguously determine the disease of the patient. Such a capability would enable it to achieve objectives like saving the expenditure of the patients, quickly determining the disease to start the treatment early, etc. Motivated by such applications, we consider the problem of constructing *decision trees for entity identification* from the given data.

Decision Trees for Entity Identification—Problem Statement. The input is a table $\mathcal{D}$ having N rows and m columns. Each row is called an *entity* and the columns are the *attributes* of these entities. Additionally, we are also given a probability distribution $\mathcal{P}$ over the set of entities. For each entity e, $\mathcal{P}$ specifies $p(e)$, the likelihood of the occurrence of e. A solution is a decision tree in which each internal node is labeled by an attribute and its branches are labeled by the values that the attribute can take. The entities are the leaves of the tree. The main requirement is that the tree should identify each entity correctly. The cost of the tree is the expected distance of an entity from the root, (i.e., $\Sigma_e p(e)d(e)$, where $d(e)$ is the distance of the entity e from the root). The goal is to construct a decision tree with the minimum cost. We call this the $\mathcal{WDT}$ problem (Here, $\mathcal{W}$ stands for "weight" and it stresses the fact that the entities are associated with probabilities/weights).

Example 1.1. Figure 1 shows an example table and two decision trees for it. In this example, the probability distribution over the entities is uniform, that is, $p(e_i) = 1/6$, for each entity e_i. In the first decision tree, the distance $d(e_1)$ is 2 and $d(e_4)$ is 3. The cost of the first decision tree is 14/6 and that of the second decision tree is 8/6. The second decision tree happens to be an optimal tree for this instance.

For a given table, the maximum number of distinct values that any attribute takes is called its *branching factor*. In the preceding example, the branching factor of the given table is 5, because every attribute takes at most 5 distinct values and the attribute B attains the maximum 5. Interesting special cases of the $\mathcal{WDT}$ problem can be obtained in two ways:

—the case in which every input instance is required to have a branching factor of at most K, for some constant K; we call this the K-$\mathcal{WDT}$ problem. Of particular interest is the 2-$\mathcal{WDT}$ problem, where the tables are binary.
—the case in which the probability distribution over the set of entities is known to be uniform; we call this the $\mathcal{UDT}$ problem (Here, $\mathcal{U}$ stresses the fact that the probabilities/weights are uniform).

The special case in which both of these restrictions apply is called the K-$\mathcal{UDT}$ problem.

Prior Results and Our Results. Much of the previous literature deals with the restricted 2-$\mathcal{UDT}$ problem. Hyafil and Rivest [1976] showed that the 2-$\mathcal{UDT}$ problem is NP-hard. Garey [1970, 1972] presented a dynamic programming-based algorithm for the 2-$\mathcal{UDT}$ problem that finds the optimal solution, but the algorithm runs in exponential time in the worst case.

Kosaraju et al. [1999] presented a greedy algorithm for the 2-$\mathcal{WDT}$ problem with an approximation ratio of $O(\log N)$; the approximation ratio remains the same for the

	A	B	C
e_1	1	1	1
e_2	1	2	0
e_3	1	3	1
e_4	2	4	0
e_5	2	4	1
e_6	2	5	0

Input Table

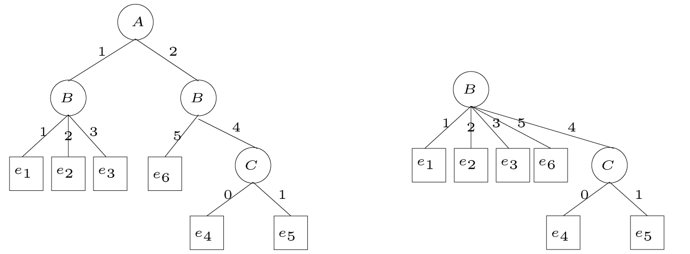

A Decision Tree of cost 14/6 Opt. Decision Tree of Cost 8/6

Fig. 1. Example decision trees.

special case of the 2-$\mathcal{UDT}$ problem. Independently, Dasgupta [2005] showed that the same greedy heuristic has an approximation ratio of $4 \log N$ for the 2-$\mathcal{UDT}$ problem. Recently, Heeringa and Adler [2005] gave an alternative analysis of the same greedy algorithm and obtained a slightly improved approximation ratio of $(1 + \ln N)$ for the 2-$\mathcal{UDT}$ problem (see also Heeringa [2006] and Adler and Heeringa [2008]). They also showed that it is NP-hard to approximate the 2-$\mathcal{UDT}$ problem within a ratio of $(1 + \epsilon)$, for some $\epsilon > 0$. We study the problem in its whole generality, namely the $\mathcal{WDT}$ problem, where the attributes can take multiple values and the input probability distribution can be arbitrary. This occurs commonly, for example, in medical diagnosis applications (e.g., blood-group can take multiple values; some diseases are more prevalent than others).

We present two approximation algorithms for the $\mathcal{UDT}$ problem. The first one is a simple algorithm that uses any given α-approximation algorithm for the 2-$\mathcal{UDT}$ problem as a black box and provides an $\alpha \lceil \log K \rceil$ approximation for the K-$\mathcal{UDT}$ problem. In particular, using the algorithm of Heeringa and Adler [2005] as the black box, we obtain an algorithm with an approximation ratio of $\lceil \log K \rceil (1 + \ln N)$. Our second algorithm for the $\mathcal{UDT}$ problem uses a greedy heuristic and has an approximation ratio of $r_K(1 + \ln N)$, where r_K is a suitably defined Ramsey number which is at most $(2 + 0.64 \log K)$. Our analysis builds on that of Heeringa and Adler [2005] and uses additional combinatorial arguments. The highlight of our analysis is that it establishes connections to Ramsey numbers and a conjecture by Erdös (see what follows for more details). Furthermore, notice that the second algorithm offers a constant factor improvement over the first algorithm.

Remark 1.2. We note that subsequent to our work, Chakaravarthy et al. [2009] considered a sligtly different greedy heuristic for the $\mathcal{UDT}$ problem and showed an approximation ratio of $4 \log N$.

Next we consider the general $\mathcal{WDT}$ problem. We first observe that by combining the black-box approach with the algorithm of Kosaraju et al. [1999], we get an $O(\log K \log N)$ approximation ratio for the $\mathcal{WDT}$ problem. We also show how extend our analysis for the $\mathcal{UDT}$ problem to handle weights and obtain an algorithm with an approximation ratio of $O(r_K \log N)$. This provides an alternative way of getting the result obtained via the black-box approach.

We next focus on the hardness of approximating various versions of the problem. We show that it is NP-hard to approximate the 2-$\mathcal{WDT}$ problem within a ratio of $\Omega(\log N)$. This implies that the $O(\log N)$-approximation algorithm of Kosaraju et al. [1999] for the 2-$\mathcal{WDT}$ problem is optimal up to constant factors. We also improve the hardness of approximation for the unweighted version of the problem. We show that it is NP-hard

Problem	Apprx. Ratio	Hardness of Apprx.
2-$\mathcal{UDT}$	$1 + \ln N$ [Heeringa and Adler 2005]	$(2 - \epsilon)$ for any $\epsilon > 0$
$\mathcal{UDT}$	$r_K(1 + \ln N)$ [This paper] Subsequently $4 \log N$ [Chakaravarthy et al. 2009]	$(4 - \epsilon)$ for any $\epsilon > 0$
2-$\mathcal{WDT}$	$O(\log N)$ [This paper, [Kosaraju et al. 1999]]	$\Omega(\log N)$
$\mathcal{WDT}$	$O(r_K \log N)$	$\Omega(\log N)$

Fig. 2. Summary of results.

to approximate the $\mathcal{UDT}$ and the 2-$\mathcal{UDT}$ problem, within a ratio of $(4 - \epsilon)$ and $(2 - \epsilon)$, respectively, for any $\epsilon > 0$. The results are summarized in Figure 2.

Ramsey Numbers and Connections to Erdös's Conjecture. Our analysis of the approximation algorithms has interesting connections with Ramsey theory and an unresolved conjecture by Erdös. Ramsey theory, treated at length in the book by Graham et al. [1990], deals with coloring the edges of complete graphs (or hypergraphs) with a specified number of colors satisfying certain constraints. For our purposes, we need the following specific type of Ramsey numbers.

For $n > 0$, let G_n denote the complete graph on n vertices. A k-coloring of G_n is a coloring of the *edges* of G_n using k colors. For $k > 0$, R_k is defined to be the smallest number n such that any k-coloring of G_n contains a monochromatic triangle.[1] The inverses of the Ramsey numbers are more convenient for our purposes. For $n > 0$, we define r_n to be the smallest number k such that we can color the edges of G_n using only k colors without inducing any monochromatic triangle.

The exact values of the Ramsey numbers for $k > 3$ are not known. However, it is known that for any k, $\frac{3^k+1}{2} \le R_k \le 1 + k!e$ (see West [2001], Nesetril and Rosenfeld [2001], and Schur [1916]). Erdös made the conjecture that for some constant α, for all k, $R_k \le \alpha^k$.

In terms of the inverse Ramsey numbers, the previous bounds translate as follows: (i) for any n, $r_n \le 2 + 0.64 \log n = O(\log n)$; (ii) $r_n = \Omega(\log n / \log \log n)$. The Erdös conjecture now reads

$$r_n = \Omega(\log n).$$

Our results provide interesting approaches to address the conjecture. Exhibit a constant $c > 0$ and show that for all $K \ge 2$, it is NP-hard to approximate the K-$\mathcal{WDT}$ problem within a factor of $c \log K \log N$. Notice that this would prove the conjecture under the assumption that NP $\neq$ P. However, we note that if the recent $O(\log N)$-approximation algorithm for $\mathcal{UDT}$ by Chakaravarthy et al. [2009] can be extended to the weighted case, the preceding approach will be ruled out. Another way of proving the conjecture would be to construct a family of bad instances for our algorithm (which is a simple greedy heuristic). We discuss the details later in the article.

Applications and Related Work. Decision trees for entity identification (as defined in this article) have been used for medical diagnosis (as described earlier), species identification in biology, fault detection, etc. [Moret 1982]. Taxonomists release field guides to help identify species based on their characteristics. These guides are often presented in the form of decision trees labeled by species characteristics. Typically, a field biologist identifies the species of a specimen at hand by referring to such guides (hopefully with as few look-ups as possible). Taxonomists refer to such decision trees as "identification keys" and an article on identification keys can be found in Wikipedia.[2]

[1]A monochromatic triangle is a triplet of vertices such that all the three edges between them have the same color. In Ramsey theory, R_k is denoted $R(3, 3, \ldots, 3)$, where "3" is repeated k times. For example, it is known that $R_1 = 3$, $R_2 = 6$, $R_3 = 17$ [Radziszowski 1994].
[2]http://en.wikipedia.org/wiki/Dichotomous_key.

Computer programs and algorithms for identification and diagnosis applications have been developed for nearly four decades (e.g., Pankhurst [1970], Reynolds et al. [2003], and Wijtzes et al. [1997]).

Murthy [1998] and Moret [1982] present excellent surveys on the use of decision trees in such diverse fields as machine learning, pattern recognition, taxonomy, switching theory, and boolean logic.

2. PRELIMINARIES

In this section, we define the $\mathcal{WDT}$ problem and its special cases. We also develop some notation used in the article.

Let $\mathcal{D}$ be a relational table having N tuples and m attributes. We call each tuple an *entity*. Let $\mathcal{E}$ and $\mathcal{A}$ denote the set of entities and attributes, respectively. For $x \in \mathcal{E}$ and $a \in \mathcal{A}$, $x.a$ denotes the value of the entity x on the attribute a. For $a \in \mathcal{A}$, $\mathcal{V}_a$ denotes the set of distinct values taken by a in $\mathcal{D}$. Let $K = \max_{a\in\mathcal{A}}\{|\mathcal{V}_a|\}$. Notice that $K \leq N$. We call K the *branching factor* of $\mathcal{D}$.

A *decision tree* T for the table $\mathcal{D}$ is a rooted tree satisfying the following properties. Each internal node u is labeled by an attribute a and has at most K children. Every branch (edge) out of u is labeled by a distinct value from the set $\mathcal{V}_a$. The entities are the leaves of the tree and thus the tree has exactly N leaves. The main requirement is that the tree should identify every entity correctly. In other words, for any entity x, the following traversal process should correctly lead to x. The process starts at the root node. Let u be the current node and a be the attribute label of u. Take the branch out of u labeled by $x.a$ and move to the corresponding child of u. The requirement is that this traversal process should reach the entity x.

Observe that the values of the attributes are used only for taking the correct branch in the traversal process. So, we can map each value of an attribute to a distinct number from 1 to K and assume that $\mathcal{V}_a$ is a subset of $\{1, 2, \ldots, K\}$. In the rest of the article, we assume that for any $x \in \mathcal{E}$ and $a \in \mathcal{A}$, $x.a \in \{1, 2, \ldots, K\}$.

For a tree T, we use "$u \in T$" to mean that u is an internal node in T. We denote by $\langle x, y\rangle$, an unordered pair of distinct entities.

Let T be a decision tree for $\mathcal{D}$. For an entity $x \in \mathcal{E}$, *path length* of x is defined to be the number of internal nodes in the path from the root to x; it is denoted $\ell_T(x)$. The sum of all path lengths is called *total path length* and is denoted $|T|$, that is, $|T| = \Sigma_{x\in\mathcal{E}}\ell_T(x)$. Let $w(\cdot)$ be a *weight function* that assigns a real number $w(x) > 0$, for each $x \in \mathcal{E}$. We define the *cost* of T with respect to $w(\cdot)$ as follows.

$$\text{cost}(T, w) = \sum_{x\in\mathcal{E}} w(x)\ell_T(x)$$

We will denote cost(T, w) as $w(T)$.

As mentioned in the Introduction, input to the $\mathcal{WDT}$ problem includes a probability distribution $\mathcal{P}$ over $\mathcal{E}$ specifying the likelihood of the occurrence of each entity and the goal is to construct a tree having the minimum expected path length. We view probabilities as weights and assume that the distribution is specified as a weight function $p(\cdot)$ that associates a weight $p(x) > 0$, for each entity x. Notice that when an entity is chosen at random according to the previous distribution, the expected path length is given by $p(T) = \text{cost}(T, p)$. We assume that the probabilities $p(x)$ are given as rational numbers. We can easily write these numbers in such way that for any entity x, $p(x) = w(x)/L$, where $w(x) \geq 1$ is an integer and L is an integer giving the common denominator. And so, without loss of generality, we assume the probability distribution will be given as an integer weight function $w(\cdot)$ over a set of entities, that is, for all $x \in \mathcal{E}$, $w(x) \geq 1$ is an integer. Notice that $p(T) = w(T)/L$ and hence, finding an optimal T under $p(\cdot)$ and $w(\cdot)$ are equivalent.

$\mathcal{WDT}$ *Problem.* The input is a relational table $\mathcal{D}$ and a probability distribution $\mathcal{P}$ represented as an integer weight function $w(\cdot)$. The goal is to construct a decision tree T having the minimum cost $w(T)$.

For a positive integer K, the K-$\mathcal{WDT}$ problem is a special case of the $\mathcal{WDT}$ problem where the input table is required to have a branching factor of at most K. Notice that in the K-$\mathcal{WDT}$ problem, the input is a table whose entries are drawn from the set $\{1, 2, \ldots, K\}$.

Of particular interest is the special case called $\mathcal{UDT}$ in which the probability distribution is uniform. In this problem, the weight function is given by $w(x) = 1$, for all $x \in \mathcal{E}$. Note that the cost of a tree T is $w(T) = |T|$. For an integer $K \geq 2$, the special case of the $\mathcal{UDT}$ where the input table is required to have a branching factor of at most K is called the K-$\mathcal{UDT}$ problem.

3. APPROXIMATION ALGORITHMS AND ANALYSIS

In this section, we present an algorithm for the $\mathcal{WDT}$ problem and prove an approximation ratio of $O(r_K \log N)$, where K refers to the branching factor of the input table. As mentioned in the Introduction, our analysis builds on that of Heeringa and Adler [2005] for the 2-$\mathcal{UDT}$ problem. In order to achieve our result, we have to extend their ideas to deal with two issues. Firstly, the attributes can be multivalued as opposed to binary; secondly, the entities can have arbitrary weights. For ease of exposition, we first show how to address the issue of attributes being multivalued. Then, we deal with the case of arbitrary weights. Specifically, Section 3.1 presents an algorithm and analysis for the $\mathcal{UDT}$ problem. These ideas are generalized in Section 3.2 to obtain an algorithm for the $\mathcal{WDT}$ problem.

3.1. The Unweighted Case: $\mathcal{UDT}$ Problem

This section deals with the $\mathcal{UDT}$ problem. Here, the probability distribution is uniform and so, the weights of all the entities are 1. The goal is to find a tree T with the minimum cost $|T|$.

We present two approximation algorithms for $\mathcal{UDT}$. The first one uses any given α-approximation algorithm for 2-$\mathcal{UDT}$ as a black box and provides an $\alpha\lceil\log K\rceil$ approximation for the K-$\mathcal{UDT}$ problem. In particular, using the algorithm of Heeringa and Adler [2005] as the black box, we obtain an algorithm with an approximation ratio of $\lceil\log K\rceil(1 + \ln N)$. Our second algorithm for the $\mathcal{UDT}$ problem uses a greedy heuristic and has an approximation ratio of $r_K(1 + \ln N)$. Recall that $r_K \leq 2 + 0.64 \log K$. Thus, the second algorithm offers a constant factor improvement over the first algorithm. The first approach has the advantage that any improvement in the approximation ratio for the 2-$\mathcal{UDT}$ problem automatically yields an improvement for the K-$\mathcal{UDT}$ problem. On the other hand, the second approach has the advantage that any improvement in the upper bound for r_K improves the approximation ratio.

3.1.1. The Black-Box Algorithm. Let $\mathcal{A}$ be an α-approximation algorithm for the 2-$\mathcal{UDT}$ problem. We show how to get a $(\alpha\lceil\log K\rceil)$-approximation algorithm for the K-$\mathcal{UDT}$ problem. The idea is to encode the given $\mathcal{UDT}$ instance as a 2-$\mathcal{UDT}$ instance and then invoke the algorithm $\mathcal{A}$ on the encoded instance.

Given an $N \times m$ table $\mathcal{D}$ having a branching factor of K, we construct an $N \times m\lceil\log K\rceil$ binary table $\mathcal{D}_2$ as follows. Each attribute in $\mathcal{D}$ is represented by $\lceil\log K\rceil$ attributes in $\mathcal{D}_2$. The former attribute is called the *original* attribute and the latter attributes are called as its *derived* attributes. The values appearing in an original attribute are represented in binary in the corresponding derived attributes. Invoke the algorithm $\mathcal{A}$ on the binary table $\mathcal{D}_2$ and let $\mathcal{T}_2$ be the decision tree returned by the algorithm. We

Procedure Greedy(E)
Input: $E \subseteq \mathcal{E}$, a set of entities in $\mathcal{D}$
Output: A decision tree T for the set E
Begin

1. If $|E| = 1$,
   Return a tree with $x \in E$ as a singleton node.
2. Let $\widehat{a}$ be the attribute that distinguishes the maximum number of pairs in E:
   $\widehat{a} = \text{argmax}_{a \in \mathcal{A}} |\{\langle x, y\rangle | x.a \neq y.a\}|$
3. Create the root node r with $\widehat{a}$ as its attribute label.
4. For $1 \leq i \leq K$,
   A. Let $E_i = \{x \in E | x.\widehat{a} = i\}$
   B. $T_i = \text{Greedy}(E_i)$
   C. Let r_i be the root of T_i. Add T_i to T by adding a branch from r to r_i with label i.
5. Return T with r as the root.

End

Fig. 3. The greedy algorithm.

obtain a decision tree $\mathcal{T}$ for $\mathcal{D}$ from $\mathcal{T}_2$ by replacing the attributes in its internal nodes with their original attributes in $\mathcal{D}$ and labeling appropriately. Notice that $|\mathcal{T}| \leq |\mathcal{T}_2|$.

Given a tree T for $\mathcal{D}$, we can construct a tree T_2 for $\mathcal{D}_2$ such that $|T_2| \leq \lceil \log K \rceil |T|$. In constructing a decision tree T_2 for the encoded instance $\mathcal{D}_2$, the main task is to take the correct branches of the internal nodes of T using the binary derived attributes. We achieve this by replacing each internal node with a complete binary tree of depth $\lceil \log K \rceil$ using the derived attributes of the original attribute of the internal node. Clearly, $|T_2| \leq \lceil \log K \rceil |T|$. This shows that $|\mathcal{T}_2^*| \leq \lceil \log K \rceil |\mathcal{T}^*|$ where $\mathcal{T}^*$ and $\mathcal{T}_2^*$ are the optimal decision trees for $\mathcal{D}$ and $\mathcal{D}_2$, respectively. Since $|\mathcal{T}_2| \leq \alpha |\mathcal{T}_2^*|$, the solution $\mathcal{T}$ returned by the black-box algorithm satisfies $|\mathcal{T}| \leq \alpha \lceil \log K \rceil |\mathcal{T}^*|$.

THEOREM 3.1. *Given a α-approximation algorithm for the 2-$\mathcal{UDT}$ problem, the black-box algorithm has an approximation ratio of $\alpha \lceil \log K \rceil$ for the $\mathcal{UDT}$ problem where K is the branching factor of the input table.*

In particular, we obtain an approximation ratio of $\lceil \log K \rceil (1 + \ln N)$ by using the Heeringa-Adler algorithm as a black box.

3.1.2. The Greedy Algorithm. In this section, we present a greedy algorithm for the $\mathcal{UDT}$ problem. The algorithm is similar in spirit to that of Heeringa and Adler [2005] for the 2-$\mathcal{UDT}$ problem. We build on their analysis and develop further combinatorial arguments to obtain our approximation ratio.

Given as input an $N \times m$ table $\mathcal{D}$ having branching factor at most K, the greedy algorithm produces a decision tree $\mathcal{T}$ as described in the following. Let $\mathcal{E}$ and $\mathcal{A}$ denote the set of entities and attributes of $\mathcal{D}$, respectively. The intuition is that any decision tree should distinguish every pair of distinct entities. So, a natural idea is to make the attribute that distinguishes the maximum number of pairs as the root of $\mathcal{T}$, where an attribute a is said to distinguish a pair $\langle x, y\rangle$, if $x.a \neq y.a$. Choosing such an attribute $\widehat{a}$ can be easily done in time $O(mN^2)$. Picking the attribute $\widehat{a}$ as the label for the root node partitions the set $\mathcal{E}$ into disjoint sets $E_1, E_2, \ldots, E_K$, where $E_i = \{x | x.\widehat{a} = i\}$. We recursively apply the same greedy procedure on each of these sets to obtain K decision trees and make these the subtrees of the root node. The greedy procedure is formally specified in Figure 3. We get the output tree $\mathcal{T}$ by calling $\mathcal{T} = \text{Greedy}(\mathcal{E})$.

THEOREM 3.2. *The greedy algorithm has an approximation ratio of* $(r_K(1+\ln N))$ *for the* $\mathcal{UDT}$ *problem, where* K *is the branching factor of the input table.*

We now analyze the greedy algorithm and prove Theorem 3.2. The analysis is divided into two parts. In the first part, we introduce certain combinatorial objects called tabular partitions and analyze the performance of the greedy algorithm using these objects. In the second part, we relate these objects to Ramsey colorings and complete the proof of Theorem 3.2.

3.1.3. Analysis Involving Tabular Partitions. Let $\mathcal{T}$ and $\mathcal{T}^*$ be the greedy and the optimal decision trees, respectively. In this section, we prove a relationship between $|\mathcal{T}|$ and $|\mathcal{T}^*|$ involving tabular partitions, defined in the following.

Definition 3.3 (*Tabular Partitions*). For any positive integer $n \geq 1$, a tabular partition P of n is a sequence $P_1, P_2, \ldots, P_n$ such that P_i is a partition of the set $\{1, 2, \ldots, n\} - \{i\}$. We require that for any distinct $1 \leq i, j \leq n$, if A is the set in P_i containing j and B is the set in P_j containing i, then $A \cap B = \emptyset$. Let the length of a partition P_i denote the number of sets in it. We define the compactness of P as $\mathtt{comp}(P) = \max_i(\text{length of } P_i)$, for $1 \leq i \leq n$. We define C_n to be the smallest number such that there exists a tabular partition of n having compactness C_n.

THEOREM 3.4. $|\mathcal{T}| \leq C_K(1+\ln N)|\mathcal{T}^*|$.

We next focus on proving the previous result. In Section 3.1.4, we shall show that $C_K \leq r_K$ and obtain Theorem 3.2 by combining the two results. We start with some notations and observations. Let T be any decision tree for $\mathcal{D}$ and u be an internal node of T. We define $E^T(u) \subseteq \mathcal{E}$ to be the set of entities in the subtree of T under u.

PROPOSITION 3.5. *For any decision tree* T *of* $\mathcal{D}$*, we have* $|T| = \Sigma_{u\in T}|E^T(u)|$.

PROOF. Each entity x contributes a cost equal to its distance from the root. Let us distribute this cost uniformly among the internal nodes on the path from x to the root. Observe that the total cost accumulated at an internal node u is equal to $|E^T(u)|$. Thus $|T| = \Sigma_{u\in T}|E^T(u)|$. □

Consider a decision tree T and a pair $\langle x, y\rangle$ of entities. We say that a node $u \in T$ *separates* the pair $\langle x, y\rangle$, if the traversal for both x and y passes through u, but x and y take different branches from u. Formally, u is said to separate [3] $\langle x, y\rangle$, if $x, y \in E^T(u)$ and $x.a \neq y.a$, where a is the attribute label of u. For any pair $\langle x, y\rangle$ of entities, there exists a unique separator in T that separates x and y. We define $\mathtt{SEP}(u)$ to be the set of all pairs separated by u. The separators with respect to the greedy tree $\mathcal{T}$ will be important in our analysis. For each pair $\langle x, y\rangle$, we denote by $s_{x,y}$ the separator of $\langle x, y\rangle$ in $\mathcal{T}$ and let $S_{x,y}$ denote $E^{\mathcal{T}}(s_{x,y})$.

From Proposition 3.5, we see that each node $u \in \mathcal{T}$ contributes a cost of $|E^{\mathcal{T}}(u)|$ towards the total cost $|\mathcal{T}|$ and separates the pairs in $\mathtt{SEP}(u)$. We distribute the cost $|E^{\mathcal{T}}(u)|$ equally among the pairs in $\mathtt{SEP}(u)$. For each pair $\langle x, y\rangle \in \mathtt{SEP}(u)$, we define the cost $c_{x,y} = |E^{\mathcal{T}}(u)|/|\mathtt{SEP}(u)|$. Since each pair has a unique separator, the costs $c_{x,y}$ are well defined.

It is easy to see that $|E^{\mathcal{T}}(u)| = \sum_{\langle x,y\rangle\in \mathtt{SEP}(u)} c_{x,y}$ and by Proposition 3.5, we have $|\mathcal{T}| = \sum_{\langle x,y\rangle} c_{x,y}$, where the summation is taken over all (unordered) pairs of distinct entities. Notice that each pair $\langle x, y\rangle$ also has a unique separator in $\mathcal{T}^*$. So, we rewrite the

[3] We note that the separator of $\langle x, y\rangle$ is nothing but the least common ancestor of x and y.

preceding summation by partitioning the set of all pairs according to their separators in $\mathcal{T}^*$ and obtain the following equation.

$$|\mathcal{T}| = \sum_{z \in \mathcal{T}^*} \sum_{\langle x,y \rangle \in \mathrm{SEP}(z)} c_{x,y} \tag{1}$$

For each $z \in \mathcal{T}^*$, we define $\alpha(z)$ to be the term corresponding to z in the summation given in Eq. (1). Clearly, $\alpha(z) = \sum_{\langle x,y \rangle \in \mathrm{SEP}(z)} c_{x,y}$. The following lemma gives an upper bound on $\alpha(z)$.

LEMMA 3.6. *For any $z \in \mathcal{T}^*$, $\alpha(z) \le C_K(1 + \ln|Z|)|Z|$, where $Z = E^{\mathcal{T}^*}(z)$.*

Assuming the correctness of Lemma 3.6, we first prove Theorem 3.4. The lemma is proved later in the section.

PROOF OF THEOREM 3.4. Replacing the inner summation in Eq. (1) by $\alpha(z)$ we have

$$|\mathcal{T}| \le C_K(1 + \ln N) \sum_{z \in \mathcal{T}^*} |E^{\mathcal{T}^*}(z)| = C_K(1 + \ln N)|\mathcal{T}^*|.$$

The first step is obtained by invoking Lemma 3.6 and the fact that $|Z| \le N$. Proposition 3.5 gives us the second step. $\square$

We now proceed to prove Lemma 3.6. Fix any $z \in \mathcal{T}^*$. Let us denote $Z = E^{\mathcal{T}^*}(z)$. Let a_z be the attribute label of z. The node z partitions the set Z into K sets $Z_1, Z_2, \ldots, Z_K$, where $Z_i = \{x \in Z | x.a_z = i\}$. We extend the preceding notations to sets of values. For any $A \subseteq \{1, 2, \ldots, K\}$, define $Z_A = \cup_{i \in A} Z_i$. We prove the following upper bound on $c_{x,y}$.

LEMMA 3.7. *Let $\langle x, y \rangle \in \mathrm{SEP}(z)$. Consider disjoint sets $A, B \subseteq \{1, 2, \ldots, K\}$ satisfying $y \in Z_A$ and $x \in Z_B$. Then,*

$$c_{x,y} \le \frac{1}{|S_{x,y} \cap Z_A|} + \frac{1}{|S_{x,y} \cap Z_B|}.$$

PROOF. We are given a pair $\langle x, y \rangle \in \mathrm{SEP}(z)$. Let $s = s_{x,y}$ be the separator of $\langle x, y \rangle$ in $\mathcal{T}$ and the attribute label of s be a_s. The cost $c_{x,y}$ is given by $|S_{x,y}|/|\mathrm{SEP}(s)|$, where $S_{x,y} = E^{\mathcal{T}}(s)$. The greedy algorithm chose the attribute a_s for the node s. Hypothetically, consider choosing the attribute a_z, instead. Let us denote the pairs separated by such a choice as X, that is, define $X = \{\langle x, y \rangle | x, y \in S_{x,y} \text{ and } x.a_z \ne y.a_z\}$. Notice that the greedy algorithm chose the attribute a_s, instead of a_z, because a_s distinguishes more pairs compared to a_z, meaning, $|\mathrm{SEP}(s)| \ge |X|$. It follows that $c_{x,y} \le |S_{x,y}|/|X|$. Partition $S_{x,y}$ into $S_1, S_2, \ldots, S_K$, where $S_i = \{x \in S_{x,y} | x.a_z = i\}$. Then,

$$|X| = \sum_{1 \le i < j \le K} |S_i| \cdot |S_j|.$$

Now we claim that

$$c_{x,y} \le \frac{1}{\sum_{i \in A} |S_i|} + \frac{1}{\sum_{j \in B} |S_j|}. \tag{2}$$

The claim can be proved as follows. Let $\overline{A} = A \cup (\{1, 2, \ldots, K\} - A - B)$ so that $\overline{A} \cup B = \{1, 2, \ldots, K\}$ and $\overline{A} \cap B = \emptyset$. Recall that $|S_{x,y}| = |S_1| + |S_2| + \cdots + |S_K|$. It follows that

$$\begin{aligned} \frac{1}{\sum_{i\in A}|S_i|} + \frac{1}{\sum_{j\in B}|S_j|} &\geq \frac{1}{\sum_{i\in \overline{A}}|S_i|} + \frac{1}{\sum_{j\in B}|S_j|} \\ &= \frac{S_{x,y}}{(\sum_{i\in \overline{A}}|S_i|)\cdot(\sum_{j\in B}|S_j|)} \\ &\geq \frac{S_{x,y}}{|X|} \\ &\geq c_{x,y}. \end{aligned}$$

This proves the claim in Eq. (2).

Observe that for any $1 \leq i \leq K$, $S_{x,y} \cap Z_i \subseteq S_i$ and hence $|S_{x,y} \cap Z_i| \leq |S_i|$. Therefore,

$$c_{x,y} \leq \frac{1}{\sum_{i\in A}|S_{x,y} \cap Z_i|} + \frac{1}{\sum_{j\in B}|S_{x,y} \cap Z_j|}.$$

Finally, since the sets $Z_{i'}$ and $Z_{j'}$ are disjoint for any distinct $1 \leq i' \leq j' \leq K$, it follows that the first term equals $1/|S_{x,y} \cap Z_A|$ and the second term equals $1/|S_{x,y} \cap Z_B|$. The lemma is proved. □

For each $\langle x, y\rangle$, we shall a choose a suitable pair of disjoint sets A and B and obtain an upper bound on $c_{x,y}$ by invoking Lemma 3.7. We make use of tabular partitions for choosing these sets; the motivation for doing so will become clear in the proof of Lemma 3.10. Let P^* be an optimal tabular partition of K having compactness C_K, given by the sequence $P_1, P_2, \ldots, P_K$. Consider any pair $\langle x, y\rangle \in \texttt{SEP}(z)$. Let $i = x.a_z$ and $j = y.a_z$ so that $x \in Z_i$ and $y \in Z_j$. Let $\widehat{A}$ be the set in the partition P_i that contains j and $\widehat{B}$ be the set in the partition P_j that contains i. Notice that, by the definition of tabular partitions, the sets $\widehat{A}$ and $\widehat{B}$ are disjoint. We invoke Lemma 3.7 with $\widehat{A}$ and $\widehat{B}$ as the required disjoint sets. (Observe that for any i and j, all the pairs in $Z_i \times Z_j$ will make use of the same disjoint sets while invoking the lemma. Thus the sets chosen depend only on the values $x.a_z$ and $y.a_z$). Therefore,

$$c_{x,y} \leq \frac{1}{|S_{x,y} \cap Z_{\widehat{A}}|} + \frac{1}{|S_{x,y} \cap Z_{\widehat{B}}|}.$$

We split the preceding cost into two parts and attribute the first term to x and the second term to y. Define

$$c^x_{x,y} = \frac{1}{|S_{x,y} \cap Z_{\widehat{A}}|} \text{ and } c^y_{x,y} = \frac{1}{|S_{x,y} \cap Z_{\widehat{B}}|}.$$

It follows that $c_{x,y} \leq c^x_{x,y} + c^y_{x,y}$. For any $x \in Z$, we imagine that x pays a cost $c^x_{x,y}$ to get separated from an entity $y \in Z$. We denote the accumulated cost as $\texttt{Acc}_z(x)$ and define it as

$$\texttt{Acc}_z(x) = \sum_{y:\langle x,y\rangle \in \texttt{SEP}(z)} c^x_{x,y}.$$

Now the lemma given next follows easily.

LEMMA 3.8. *For any z, $\alpha(z) \leq \Sigma_{x\in Z}\texttt{Acc}_z(x)$.*

Our next task is to obtain an upper bound on $\texttt{Acc}_z(x)$, so that we get a bound on $\alpha(z)$. The following lemma is useful for this purpose.

LEMMA 3.9. *Let $x \in \mathcal{E}$ be any entity and $Q \subseteq \mathcal{E}$ be any set of entities such that $x \notin Q$. Then,*

$$\sum_{y \in Q} \frac{1}{|S_{x,y} \cap Q|} \leq (1 + \ln |Q|).$$

PROOF. Let $t = |Q|$. We shall prove the following claim.

$$\sum_{y \in Q} \frac{1}{|S_{x,y} \cap Q|} \leq \sum_{i=1}^{t} \frac{1}{i}$$

The claim implies the lemma, since it is well known that $\sum_{i=1}^{t}(1/i) \leq (1 + \ln t)$, for all t. We prove the claim by applying induction on $|Q|$. For the base case of $|Q| = 1$, let $Q = \{y\}$, where $y \neq x$. Clearly, $y \in S_{x,y}$ and so, $|S_{x,y} \cap Q| = 1$, and the claim follows. Assuming that the claim is true for all sets of size at most $t-1$, we prove it for any set Q of size t. Let y^* be any entity in Q such that for all $y \in Q$, $s_{x,y}$ is a descendent of s_{x,y^*} (a node is considered to be a descendent of itself). If more than one such element exists, pick one arbitrarily. Intuitively, y^* is one among the first batch of entities in Q to get separated from x. The main observation is that $Q \subseteq S_{x,y^*}$ and so, $S_{x,y^*} \cap Q = Q$. Thus $1/|S_{x,y^*} \cap Q| = 1/|Q| = 1/t$. We apply the induction hypothesis on the set of remaining entities $Q' = Q - y^*$ and infer that

$$\sum_{y \in Q'} \frac{1}{|S_{x,y} \cap Q'|} \leq \sum_{i=1}^{t-1} \frac{1}{i}.$$

Clearly, $Q' \subseteq Q$ and hence $|S_{x,y} \cap Q'| \leq |S_{x,y} \cap Q|$, so, in the previous summation, if we replace the term $|S_{x,y} \cap Q'|$ by $|S_{x,y} \cap Q|$, then the resulting inequality is also true. We conclude that

$$\begin{aligned}
\sum_{y \in Q} \frac{1}{|S_{x,y} \cap Q|} &= \frac{1}{|S_{x,y^*} \cap Q|} + \sum_{y \in Q'} \frac{1}{|S_{x,y} \cap Q|} \\
&\leq \frac{1}{t} + \sum_{i=1}^{t-1} \frac{1}{i} \\
&= \sum_{i=1}^{t} \frac{1}{i}.
\end{aligned}$$

□

LEMMA 3.10. *For any $x \in Z$, $\texttt{Acc}_z(x) \leq C_K(1 + \ln |Z|)$.*

PROOF. Let $r = x.a_z$ and so $x \in Z_r$. Let $\widetilde{Z} = Z - Z_r$ be the rest of the entities in Z. Notice that $\texttt{Acc}_z(x) = \Sigma_{y \in \widetilde{Z}} c^x_{x,y}$. We perform the preceding summation by partitioning $\widetilde{Z}$ according to P_r, the r^{th} member of the optimal tabular partition $P^* = P_1, P_2, \ldots, P_K$. Let $P_r = s_1, s_2, \ldots, s_\ell$, where $\ell \leq C_K$. For $1 \leq i \leq \ell$, define $Q_i = \{y \in \widetilde{Z} | y.a_z \in s_i\}$. Thus, $\widetilde{Z} = Q_1 \cup Q_2 \cup \ldots \cup Q_\ell$ and hence,

$$\texttt{Acc}_z(x) = \sum_{1 \leq i \leq \ell} \sum_{y \in Q_i} c^x_{x,y}. \tag{3}$$

We derive an upper bound for each term in the outer sum using Lemma 3.9. Fix any $1 \leq i \leq \ell$. Notice that for any $y \in Q_i$, we have $c^x_{x,y} = 1/|S_{x,y} \cap Q_i|$, by definition. Moreover,

$x \notin Q_i$. Thus, by applying Lemma 3.9 on Q_i, we get

$$\sum_{y \in Q_i} c^x_{x,y} \leq (1 + \ln |Q_i|) \leq (1 + \ln |Z|). \tag{4}$$

We get the lemma by combining Eqs. (3) and (4), and the fact that $\ell \leq C_K$. □

PROOF OF LEMMA 3.6. The result is proved by combining Lemma 3.8 and Lemma 3.10.

$$\begin{aligned} \alpha(z) &\leq \sum_{x \in Z} \mathtt{Acc}_z(x) \\ &\leq \sum_{x \in Z} C_K(1 + \ln |Z|) \\ &= C_K(1 + \ln |Z|)|Z|. \end{aligned}$$

□

3.1.4. Tabular Partitions and Ramsey Colorings. In this section, we introduce the notion of directed Ramsey colorings and show that they are equivalent to tabular partitions. Throughout the discussion, for $n > 0$, let G_n and $\widetilde{G}_n$ denote the complete undirected and the complete directed graph on n vertices, respectively.

Definition 3.11. Let $n > 0$ be an integer. A directed Ramsey coloring of $\widetilde{G}_n$ is a coloring $\widetilde{\tau}$ of the edges such that for any triplet of distinct vertices x, y and z, if $\widetilde{\tau}(x, y) = \widetilde{\tau}(x, z)$ then $\widetilde{\tau}(y, x) \neq \widetilde{\tau}(y, z)$ (and by symmetry, $\widetilde{\tau}(z, x) \neq \widetilde{\tau}(z, y)$).

We define $\widetilde{R}_k$ to be the smallest number n such that $\widetilde{G}_n$ cannot be directed Ramsey colored using k colors.[4] The inverse of these numbers will be useful. Define $\widetilde{r}_n$ to be the minimum number of colors required to do a directed Ramsey coloring of $\widetilde{G}_n$.

We claim that for any n, there exists a tabular partition P of compactness k if and only if there exists a directed Ramsey coloring $\widetilde{\tau}$ of $\widetilde{G}_n$ that uses only k colors. A proof sketch follows. Let $P = P_1, P_2, \ldots P_n$. Fix $1 \leq x \leq n$. Arrange the sets in the partition P_x in an arbitrary manner, say $P_x = s_{x,1}, s_{x,2}, \ldots, s_{x,\ell}$, where $\ell \leq k$. The n-1 edges outgoing from the vertex x are colored according to the partition P_x. Meaning, for $1 \leq c \leq \ell$, for $y \in s_{x,c}$, we set $\widetilde{\tau}(x, y) = c$. For any y and z, if $\widetilde{\tau}(x, y) = \widetilde{\tau}(x, z)$, then it means that y and z belong to the same set in the partition P_x. By the property of tabular partitions, it should be the case that x and z belong to different sets in the partition P_y, implying that $\widetilde{\tau}(y, x) \neq \widetilde{\tau}(y, z)$. We conclude that $\widetilde{\tau}$ is a directed Ramsey coloring and that $\widetilde{\tau}$ uses only k colors. The converse is proved using a similar argument. The claim implies the following proposition.

THEOREM 3.12. *For any n, $C_n = \widetilde{r}_n$.*

Let us call an edge-coloring of G_n a Ramsey coloring if it does not induce any monochromatic triangles. For any n, a Ramsey coloring τ of G_n readily yields a directed Ramsey coloring $\widetilde{\tau}$ of $\widetilde{G}_n$. For each pair of vertices x and y, we set $\widetilde{\tau}(x, y) = \widetilde{\tau}(y, x) = \tau(x, y)$. It can easily be verified that $\widetilde{\tau}$ is indeed a directed Ramsey coloring of $\widetilde{G}_n$. The number of colors used in $\widetilde{\tau}$ is the same as that of τ. Therefore, we have the following proposition.

PROPOSITION 3.13. *For any n, $\widetilde{r}_n \leq r_n$.*

PROOF OF THEOREM 3.2. The result follows from Theorems 3.4 and 3.12, and Proposition 3.13. □

[4]Such a number exists, as shown in Theorem 5.1.

3.2. The Weighted Case: $\mathcal{WDT}$ Problem

In this section, we show how to deal with the weighted case, namely the $\mathcal{WDT}$ problem. Let D be the input $N \times m$ table over a set of entities $\mathcal{E}$ and a set of attributes $\mathcal{A}$, having a branching factor of K. Let $w(\cdot)$ be the input weight function that assigns an integer weight $w(x) \geq 1$ to each entity $x \in \mathcal{E}$. The problem is to construct an optimal decision tree $\mathcal{T}^*$ having the minimum cost with respect to $w(\cdot)$. We present an algorithm which generalizes the greedy algorithm for the $\mathcal{UDT}$ problem.

Weighted Greedy Algorithm. Refer to the greedy algorithm given in Figure 3. The main step in that algorithm is choosing an attribute that distinguishes the maximum number of pairs. We modify this step so that the weights are taken into account. Namely, we choose the attribute $\widehat{a}$

$$\widehat{a} = \operatorname{argmax}_{a \in \mathcal{A}} \sum_{\langle x,y\rangle \in S(a)} w(x)w(y),$$

where $S(a) = \{\langle x, y\rangle | x, y \in \mathcal{E} \text{ and } x.a \neq y.a\}$, is the set of pairs distinguished by the attribute a. We call the preceding procedure the weighted greedy algorithm.

Let $W = \Sigma_{x\in\mathcal{E}} w(x)$ denote the total weight of the entities. Let $\mathcal{T}$ and $\mathcal{T}^*$ denote the weighted greedy and the optimal trees, under the weight function $w(\cdot)$.

THEOREM 3.14. *$w(\mathcal{T}) \leq C_K(1 + \ln W)w(\mathcal{T}^*)$, where W is the sum of weights of all the entities.*

We prove the this theorem by adapting the proof of Theorem 3.2. Due to space constraints, we provide an outline of the proof.

Intuitively, we imagine that each entity x is replicated $w(x)$ times and modify the proof of Theorem 3.2 accordingly. We reuse notation from the previous proof. Let $\text{SEP}(u)$ be the set of all pairs separated by u. For each pair $\langle x, y\rangle$, we denote by $s_{x,y}$ the separator of $\langle x, y\rangle$ in $\mathcal{T}$ and let $S_{x,y}$ denote $E^{\mathcal{T}}(s_{x,y})$. Additional notation is introduced next.

For a set of entities $X \subseteq \mathcal{E}$, let $w(X)$ denote the total weight of the entities in X, that is, $w(X) = \Sigma_{x\in X} w(x)$. We also define weights on any set of pairs of entities: for a set of pairs $X \subseteq \mathcal{E} \times \mathcal{E}$, define $w(X) = \Sigma_{\langle x,y\rangle \in X} w(x)w(y)$.

Proposition 3.5 generalizes to the weighted case as follows.

PROPOSITION 3.15. *For a decision tree T of $\mathcal{D}$, $w(T) = \Sigma_{u\in T}\, w(E^T(u))$.*

For each pair of entities $\langle x, y\rangle$, define a cost $c_{x,y}$ as follows.

$$c_{x,y} = w(x)w(y)\left[\frac{w(S_{x,y})}{w(\text{SEP}(s_{x,y}))}\right]$$

By Proposition 3.15, we get the following equation, which is similar to Eq. (1).

$$w(\mathcal{T}) = \sum_{z\in\mathcal{T}^*} \sum_{\langle x,y\rangle \in \text{SEP}(z)} c_{x,y} \tag{5}$$

For each $z \in \mathcal{T}^*$, the inner summation in Eq. (5) is defined as the cost $\alpha(z) = \sum_{\langle x,y\rangle \in \text{SEP}(z)} c_{x,y}$. Our goal is to derive an upper bound on $\alpha(z)$.

Fix any $z \in \mathcal{T}^*$. Let us denote $Z = E^{\mathcal{T}^*}(z)$. Let a_z be the attribute label of z. The node z partitions the set Z into K sets $Z_1, Z_2, \ldots, Z_K$, where $Z_i = \{x \in Z | x.a_z = i\}$. We extend the preceding notations to sets of values. For any $A \subseteq \{1, 2, \ldots, K\}$, define $Z_A = \cup_{i\in A} Z_i$. The following lemma generalizes Lemma 3.7 to the weighted case.

LEMMA 3.16. *Let $\langle x, y\rangle \in \texttt{SEP}(z)$. Consider disjoint sets $A, B \subseteq \{1, 2, \ldots, K\}$ satisfying $y \in Z_A$ and $x \in Z_B$. Then,*

$$c_{x,y} \leq w(x)w(y)\left[\frac{1}{w(S_{x,y} \cap Z_A)} + \frac{1}{w(S_{x,y} \cap Z_B)}\right].$$

Consider any $\langle x, y\rangle \in \texttt{SEP}(z)$. Let P^* be an optimal tabular partition of K having compactness C_K, given by the sequence $P_1, P_2, \ldots, P_K$. Let $i = x.a_z$ and $j = y.a_z$ so that $x \in Z_i$ and $y \in Z_j$. Let $\widehat{A}$ be the set in the partition P_i that contains j and $\widehat{B}$ be the set in the partition P_j that contains i. Define

$$c^x_{x,y} = \frac{w(x)w(y)}{w(S_{x,y} \cap Z_{\widehat{A}})} \text{ and } c^y_{x,y} = \frac{w(x)w(y)}{w(S_{x,y} \cap Z_{\widehat{B}})}.$$

By Lemma 3.16, we have that $c_{x,y} \leq c^x_{x,y} + c^y_{x,y}$. For each entity $x \in E^{\mathcal{T}^*}(z)$, define $\texttt{Acc}_z(x)$ as

$$\texttt{Acc}_z(x) = \sum_{y:\langle x,y\rangle \in \texttt{SEP}(z)} c^x_{x,y}.$$

We wish to derive an upper bound on $\texttt{Acc}_z(x)$. The following lemma, which generalizes Lemma 3.9, is useful for this purpose.

LEMMA 3.17. *Let $x \in \mathcal{E}$ be any entity and $Q \subseteq \mathcal{E}$ be any set of entities such that $x \notin Q$. Then,*

$$\sum_{y \in Q} \frac{w(y)}{w(S_{x,y} \cap Q)} \leq (1 + \ln w(Q)).$$

The following is obtained by generalizing Lemma 3.10.

LEMMA 3.18. *For any $x \in Z$, $\texttt{Acc}_z(x) \leq w(x)C_K(1 + \ln w(Z))$.*

PROOF OF THEOREM 3.14. Consider any $z \in \mathcal{T}^*$ and let $Z = E^{\mathcal{T}^*}(z)$. Then, $\alpha(z) \leq \Sigma_{x\in Z}, \texttt{Acc}_z(x)$. Applying Lemma 3.18 and Proposition 3.15, we get that

$$\alpha(z) \leq C_K(1 + \ln w(Z))w(Z). \tag{6}$$

Replacing the inner summation in Eq. (5) by $\alpha(z)$ we have

$$\begin{aligned} w(\mathcal{T}) &\leq C_K(1 + \ln W) \sum_{z \in \mathcal{T}^*} w(E^{\mathcal{T}^*}(z)) \\ &= C_K(1 + \ln W)w(\mathcal{T}^*). \end{aligned} \tag{7}$$

The first step is obtained by invoking Eq. (6) and the fact that $w(Z) \leq w(\mathcal{E}) = W$. Proposition 3.15 gives us the second step. □

Theorem 3.14 shows that the approximation ratio of the weighted greedy algorithm is logarithmic in N, when the total weight W is polynomially bounded in N. Unfortunately, when the weights are arbitrarily large, the ratio could be worse. We overcome this issue by using the following rounding technique.

Rounded Greedy Algorithm. Let $\mathcal{D}$ be an input table having a branching factor of K and let w_{in} be the input integer weight function. Let $w^{\max}_{\text{in}} = \max_x w_{\text{in}}(x)$ denote the maximum weight. Define a new weight function $w(\cdot)$ as follows: for any entity $x \in \mathcal{E}$, define

$$w(x) = \left\lceil \frac{w_{\text{in}}(x)N^2}{w^{\max}_{\text{in}}} \right\rceil.$$

Run the weighted greedy algorithm with $w(\cdot)$ as the input weight function and obtain a tree $\mathcal{T}$. Return the tree $\mathcal{T}$.

Let $\mathcal{T}^*$ and $\mathcal{T}^*_{\text{in}}$ be the optimal decision trees under the weight functions $w(\cdot)$ and $w_{\text{in}}(\cdot)$, respectively. From Theorem 3.14, we have a good bound for $w(\mathcal{T})$ with respect to $w(\mathcal{T}^*)$. But, of course, we need to compare $w_{\text{in}}(\mathcal{T})$ and $w_{\text{in}}(\mathcal{T}^*_{\text{in}})$. We do this next.

THEOREM 3.19. $w_{\text{in}}(\mathcal{T}) \le 2C_K(1 + 3\ln N)w_{\text{in}}(\mathcal{T}^*_{\text{in}})$.

PROOF. Let $x \in \mathcal{E}$ be any entity and consider the path from the root to x in the tree $\mathcal{T}^*_{\text{in}}$. Notice that each internal node along this path separates at least one entity from x. (Otherwise, $\mathcal{T}^*_{\text{in}}$ contains a "dummy" node that does not separate any pairs and hence can be deleted to obtain a tree of lesser cost). So, the length of the path is at most N and hence, the following claim is true.

Claim 1: $|\mathcal{T}^*_{\text{in}}| \le N^2$.

We next compare $w_{\text{in}}(\mathcal{T}^*_{\text{in}})$ and $w(\mathcal{T}^*_{\text{in}})$. We have

$$\begin{aligned} w(\mathcal{T}^*_{\text{in}}) &= \sum_{x\in\mathcal{E}} w(x)\ell_{\mathcal{T}^*_{\text{in}}}(x) \\ &\le \sum_{x\in\mathcal{E}} \left(\frac{w_{\text{in}}(x)N^2}{w^{\max}_{\text{in}}} + 1\right)\ell_{\mathcal{T}^*_{\text{in}}}(x) \\ &= \frac{w_{\text{in}}(\mathcal{T}^*_{\text{in}})N^2}{w^{\max}_{\text{in}}} + |\mathcal{T}^*_{\text{in}}| \\ &\le \frac{w_{\text{in}}(\mathcal{T}^*_{\text{in}})N^2}{w^{\max}_{\text{in}}} + N^2 \\ &\le \frac{2w_{\text{in}}(\mathcal{T}^*_{\text{in}})N^2}{w^{\max}_{\text{in}}}. \end{aligned} \tag{8}$$

The second step is from the definition of $w(\cdot)$ and the fourth step is obtained from Claim 1. The last inequality is obtained by observing the fact that $w_{\text{in}}(\mathcal{T}^*_{\text{in}}) \ge w^{\max}_{\text{in}}$.

Notice that for any entity $x \in \mathcal{E}$, $1 \le w(x) \le N^2$ and so the total weight W under the function $w(\cdot)$ satisfies $W \le N^3$. So, Theorem 3.14 implies the following claim.

Claim 2: $w(\mathcal{T}) \le C_K(1 + 3\ln N)w(\mathcal{T}^*)$.

We can now compare $w_{\text{in}}(\mathcal{T})$ and $w_{\text{in}}(\mathcal{T}^*_{\text{in}})$. Note that $\mathcal{T}^*$ is the optimal tree under the function $w(\cdot)$ and hence, $w(\mathcal{T}^*) \le w(\mathcal{T}^*_{\text{in}})$. We obtain the lemma by combining the observation with Eq. (8) and Claim 2. □

By combining Theorem 3.19, Theorem 3.12, and Proposition 3.13, we get the following result.

THEOREM 3.20. *The approximation ratio of the rounded greedy algorithm is at most* $2r_K(1 + 3\ln N) = O(r_K \log N)$.

4. HARDNESS OF APPROXIMATION

In this section, we study the hardness of approximating the $\mathcal{WDT}$ and the $\mathcal{UDT}$ problems. We show that it is NP-hard to approximate the 2-$\mathcal{WDT}$ problem within a ratio of $\Omega(\log N)$. Therefore, our approximation algorithm for the 2-$\mathcal{WDT}$ problem is optimal up to constant factors. We also improve the previous hardness results for the $\mathcal{UDT}$ problem.

4.1. Hardness of Approximating the 2-$\mathcal{WDT}$ Problem

THEOREM 4.1. *It is NP-hard to approximate the* 2-$\mathcal{WDT}$ *problem within a factor of* $\Omega(\log N)$*, where* N *is the number of entities in the input.*

PROOF. We prove the result via a reduction from the set cover problem. It is known that approximating set cover within a factor of $\Omega(\log n)$ is NP-hard [Raz and Safra 1997].

Let $(U, \mathcal{S})$ be the input set cover instance, where $U = \{x_1, x_2, \ldots, x_n\}$ is a universe of items and $\mathcal{S}$ is a collection of sets $\{S_1, S_2, \ldots, S_m\}$ such that $S_i \subseteq U$, for each i. Without loss of generality, we can assume that for any pair of distinct items x_i and x_j, there exists a set $S_k \in \mathcal{S}$ containing exactly one of these two items. (If not, one of these items can be removed from the system.) Construct an instance of the 2-$\mathcal{WDT}$ problem having $N = n+1$ entities and m attributes. The set of entities is $\mathcal{E} = \{x_1, x_2, \ldots, x_n\} \cup \{\widehat{x}\}$, where each entity x_i corresponds to the item x_i and $\widehat{x}$ is a special entity. The set of attributes is $\mathcal{A} = \{S_1, S_2, \cdots, S_m\}$, so that each attribute S_i corresponds to the set S_i. The $N \times m$ table $\mathcal{D}$ is given as follows. For each entity x_i and attribute S_j, set $x_i.S_j = 1$, if $x_i \in S_j$ and otherwise, set $x_i.S_j = 0$. For the special entity $\widehat{x}$, set $\widehat{x}.S_j = 0$, for all attributes S_j. For each entity x_i, set the weight $w(x_i) = 1$. As for the special entity $\widehat{x}$, set its weight as $w(\widehat{x}) = N^3$. This completes the construction.

Let T be a decision tree for $\mathcal{D}$. Let C be the set of attributes found along the path from the root to the entity $\widehat{x}$. Recall that the length of the preceding path is denoted as $\ell_T(\widehat{x})$. Observe that C is a cover for $(U, \mathcal{S})$. We have $(|C| = \ell_T(\widehat{x})) \leq w(T)/N^3$. On the other hand, given a cover C, we can construct a decision tree T satisfying the following two properties: (i) the set of attributes along the path from the root to $\widehat{x}$ is exactly the set C so that $|\ell_T(\widehat{x})| = |C|$; (ii) for every other entity x_i, $\ell_T(x_i) \leq N$. (The second property is based on the fact that for any table containing N entities, it suffices to test at most N attributes in order to distinguish any entity from the rest). Thus $w(T) \leq |C|N^3 + N^2$. In particular, $w(T^*) \leq |C^*|N^3 + N^2$, where T^* and C^* are the optimal decision tree and optimal cover, respectively.

Based on the previous observations, we can prove the following claim. If there exists an $\alpha(N)$-approximation algorithm for the 2-$\mathcal{WDT}$ problem then for any $\epsilon > 0$, we can design an $(1 + \epsilon)\alpha(n)$-approximation algorithm for the set cover problem. Therefore, the hardness of set cover problem implies the claimed hardness result for the 2-$\mathcal{WDT}$ problem. □

4.2. Hardness for the $\mathcal{UDT}$ and the 2-$\mathcal{UDT}$ Problems

In this section, we present improved results of hardness of approximation for the $\mathcal{UDT}$ and the 2-$\mathcal{UDT}$ problems. For the 2-$\mathcal{UDT}$ problem, Heeringa and Adler [2005] showed a hardness of approximation of $(1 + \epsilon)$, for some $\epsilon > 0$. We show that for any $\epsilon > 0$, it is NP-hard to approximate the $\mathcal{UDT}$ and the 2-$\mathcal{UDT}$ problems within a factor of $(4 - \epsilon)$ and $(2 - \epsilon)$, respectively. Our reductions are from the Minimum Sum Set Cover ($\mathcal{MSSC}$) problem.

The input to the $\mathcal{MSSC}$ problem is a set system: a collection of sets $\mathcal{S} = \{S_1, S_2, \ldots, S_m\}$ over a universe $U = \{x_1, x_2, \ldots, x_N\}$ of *items*, where each $S_i \subseteq U$. A solution is an ordering π on the sets in $\mathcal{S}$, with an associated cost defined as follows. Let π be $S'_1, S'_2, \ldots, S'_m$. Each item in S'_1 pays a cost of 1, each item in $S'_2 - S'_1$ pays a cost of 2, and so on. The cost of π is the sum of the costs of all items. Formally, define the costs $c^\pi_x = \text{argmin}_i\{x \in S'_i\}$, for $x \in U$, and $\text{cost}(\pi) = \sum_{x \in U} c^\pi_x$. The $\mathcal{MSSC}$ problem is to find an ordering with the minimum cost. For a constant d, the d-$\mathcal{MSSC}$ problem is the special case of $\mathcal{MSSC}$ in which every set in the set system has at most d elements. Feige et al. [2004] proved the following hardness results for these problems.

THEOREM 4.2. [FEIQE ET AL. 2004].

(1) *For any $\epsilon > 0$, it is NP-hard to approximate the $\mathcal{MSSC}$ problem within a ratio of $(4 - \epsilon)$.*
(2) *For any $\epsilon > 0$, there exists a constant d such that it is NP-hard to approximate d-$\mathcal{MSSC}$ within a ratio of $(2 - \epsilon)$.*

Our hardness results for $\mathcal{UDT}$ and 2-$\mathcal{UDT}$ are obtained via approximation preserving reductions from the $\mathcal{MSSC}$ and the d-$\mathcal{MSSC}$ problems, respectively. The reduction from $\mathcal{MSSC}$ to $\mathcal{UDT}$ is easier and we present it first (in Section 4.2.1). The reduction from d-$\mathcal{MSSC}$ to 2-$\mathcal{UDT}$ is similar, but involves more technical details; it is presented in Section 4.2.2.

4.2.1. Hardness for the $\mathcal{UDT}$ Problem. Here, we prove the hardness result for the $\mathcal{UDT}$ problem by exhibiting a reduction from $\mathcal{MSSC}$.

THEOREM 4.3. *For any $\epsilon > 0$, it is NP-hard to approximate the $\mathcal{UDT}$ problem within a ratio of $(4 - \epsilon)$.*

PROOF. Given an $\mathcal{MSSC}$ instance $\mathcal{S} = \{S_1, S_2, \dots, S_m\}$ over a universe $U = \{x_1, x_2, \dots, x_N\}$, construct an $N \times m$ table $\mathcal{D}$ as follows. Each item x corresponds to an entity and each set S_i corresponds to an attribute a_i. For $1 \leq j \leq m$, $1 \leq i \leq n$, set the entry $x_i.a_j$ as follows: if $x_i \in S_j$ then set $x_i.a_j = i$, else set $x_i.a_j = 0$. Observe that any decision tree for $\mathcal{D}$ is *left-deep*: for any internal node u, except the branch labeled 0, every other branch out of u leads to a leaf node.

We claim that given an ordering π of $\mathcal{S}$, we can construct a decision tree $\mathcal{T}$ such that $|\mathcal{T}| = \text{cost}(\pi)$ and vice versa. Let $\pi = S'_1, S'_2, \dots, S'_m$ and $a'_1, a'_2, \dots, a'_m$ be the corresponding sequence of attributes. Construct a left-deep tree $\mathcal{T}$, in which the root node is labeled a'_1 and its 0^{th} child is labeled a'_2 and so on. In general, label the internal node in i^{th} level with a'_i. It can be seen that $\mathcal{T}$ is indeed a decision tree for $\mathcal{D}$ and that $|\mathcal{T}| = \text{cost}(\pi)$. The converse is shown via a similar construction. Given a decision tree $\mathcal{T}$, traverse the tree starting with the root node and always taking the branches labeled 0. Write down the sequence of sets corresponding to the internal nodes seen in this traversal and let π denote the sequence. Notice that the sets appearing in this sequence cover all elements of U and that $\text{cost}(\pi) = |T|$. (Some sets in $\mathcal{S}$ may not appear in this sequence. To be formally compliant with the definition of solutions, we append the missing sets in an arbitrary order). The claim, in conjunction with Theorem 4.2 (part 1), implies the following result. □

4.2.2. Hardness for the 2-$\mathcal{UDT}$ Problem. In this section, we shall prove the hardness result for the 2-$\mathcal{UDT}$ problem.

THEOREM 4.4. *For any $\epsilon > 0$, it is NP-hard to approximate the 2-$\mathcal{UDT}$ problem within a ratio of $(2 - \epsilon)$.*

The proof is similar to that of Theorem 4.3. The reduction is from d-$\mathcal{MSSC}$ instances, for a suitable constant d. Observe that the entries in the table can only be 0 or 1 as opposed to the index of the elements in the previous construction. The required reduction is obtained by using $\lceil \log d \rceil$ auxiliary columns to identify elements of each set. The rest of the section is devoted to a formal proof.

We first present a general construction, using which we shall derive the theorem. Fix any integer $d > 0$ and we shall show a reduction from d-$\mathcal{MSSC}$ to 2-$\mathcal{UDT}$. Given a d-$\mathcal{MSSC}$ instance $\mathcal{S} = \{S_1, S_2, \dots, S_m\}$ over a universe of items $U = \{x_1, x_2, \dots, x_N\}$, construct a binary table $\mathcal{D}$ having N entities and $m(1 + \lceil \log d \rceil)$ attributes, as follows. Each item $x \in U$ corresponds to an entity x in $\mathcal{D}$. Each set S_i corresponds to $1 + \lceil \log d \rceil$ attributes: a *main* attribute named a_i and $\lceil \log d \rceil$ *auxiliary* attributes named $a_{i,1}, a_{i,2}, \ldots, a_{i,\lceil \log d \rceil}$. For filling the table $\mathcal{D}$, consider each set S_j. Order the items in S_j arbitrarily and let the ordering be $S_j = x'_1, x'_2, \ldots, x'_\ell$, where $\ell \leq d$. For each entity $x'_i \in S_j$, set $x'_i.a_j = 1$; write i as a $\lceil \log d \rceil$-bit string and fill the entries $x'_i.a_{j,1}, x'_i.a_{j,2}, \ldots, x'_i.a_{j,\lceil \log d \rceil}$ with these bits. For any entity $x \notin S_j$, set the value on all these $1 + \lceil \log d \rceil$ attributes to be 0. This completes the construction of $\mathcal{D}$. We make two claims connecting the solutions of the $\mathcal{MSSC}$ and the decision trees of $\mathcal{D}$.

Any decision tree T for $\mathcal{D}$ is a binary tree in which each internal node has two branches labeled 0 and 1; we call these the 0-branch and the 1-branch, and the corresponding children the 0-child and the 1-child, respectively. Let $\pi = S'_1, S'_2, \ldots, S'_m$ be an ordering of $\mathcal{S}$. We say that a set S'_i covers an entity x, if $x \in S'_i$ and $x \notin S'_j$, for all $j < i$.

LEMMA 4.5. *Given an ordering π of $\mathcal{S}$, we can construct a decision tree $\mathcal{T}$ of $\mathcal{D}$ such that $|\mathcal{T}| \leq \mathrm{cost}(\pi) + N\lceil \log d \rceil$. In particular, if π^* and $\mathcal{T}^*$ are the respective optimal solutions, then $|\mathcal{T}^*| \leq \mathrm{cost}(\pi^*) + N\lceil \log d \rceil$.*

PROOF. Let $\pi = S'_1, S'_2, \ldots, S'_m$ and let $a'_1, a'_2, \ldots, a'_m$ be the sequence of main attributes corresponding to these sets. We construct a tree $\mathcal{T}$ that would be "almost left-deep." We start the construction by making a'_1 the label of the root node. Notice that all the entities in S'_1 will follow the 1-branch and all the other entities will follow the 0-branch. For the former entities, the auxiliary attributes corresponding to S'_1 contain the index of the entities within S'_1. So, using these attributes we can identify each entity within S'_1 paying a cost of at most $\lceil \log d \rceil$. (Formally, construct a complete binary tree of depth $\lceil \log d \rceil$ with the auxiliary attributes as labels and assign the entities within S'_1 appropriately, and attach this tree to the 1-branch of the root node). We make a'_2 the label of the 0-child of the root node. The discussion for this node is similar to that of the root node. All entities that are covered by S'_2 will follow the 1-branch and are identified using the auxiliary attributes corresponding to S'_2. The remaining entities follow the 1-branch. In general, if we follow the 0-branches from the root node, and reach a level i, we will have a node labeled by a'_i. The entities covered by S'_i will follow the 1-branch of this node and they will be identified using a complete binary tree of depth at most $\lceil \log d \rceil$. For any entity x, if S'_i is the set covering x, then starting at the root node, x will follow the 0-branch until it reaches the node labeled a'_i, where it will follow the 1-branch and then get identified by the complete binary tree on the 1-branch. The entity incurs a cost of i for the former process and a cost of at most $\lceil \log d \rceil$ for the latter process. Thus $|\mathcal{T}| \leq \mathrm{cost}(\pi) + N\lceil \log d \rceil$. □

LEMMA 4.6. *Given a decision tree $\mathcal{T}$ for $\mathcal{D}$, we can construct an ordering π for $\mathcal{S}$ such that* $\mathrm{cost}(\pi) \leq |\mathcal{T}| + N$.

PROOF. Starting from the root node traverse the tree always taking the 0-branch until an entity x^* (a leaf node) is reached. Let $b'_1, b'_2, \ldots, b'_r$ be the sequence of attributes seen in this traversal. Construct an ordering π by writing down the sets corresponding to these attributes (each attribute b'_i can either be a main attribute or an auxiliary attribute; in either case, we write down the corresponding set). The sequence may not include all sets. However, except for x^*, all the other entities are covered by the sets in the sequence π. We deal with x^* by appending to π any set that includes x^* (and to be formally compliant with the definition of $\mathcal{MSSC}$ solutions, the sets not listed in π are appended in any arbitrary order). Notice that $\mathrm{cost}(\pi) \leq |\mathcal{T}| + N$ (the extra N is included for handling the cost of covering x^*). □

For convenience, we switch over to average costs instead of absolute costs. If π is an ordering of an $\mathcal{MSSC}$ instance having N items, we define $\mathrm{cost}_a(\pi) = \mathrm{cost}(\pi)/N$. Similarly, if T is a decision tree of a table $\mathcal{D}$ having N entities, we define $|T|_a = |T|/N$.

An important property of d-$\mathcal{MSSC}$ is that every set can cover at most d (a constant) number of items and so, in any solution, the average cost is at least $N/2d$. The following proposition formalizes the claim.

PROPOSITION 4.7. *For any d, for any d-$\mathcal{MSSC}$ instance having N items, the optimal ordering π^* satisfies*

$$\mathrm{cost}_a(\pi^*) \geq \frac{N}{2d}.$$

LEMMA 4.8. *Suppose for some constant α, there exists an α-approximation for the 2-$\mathcal{UDT}$ problem. Let $\delta > 0$ and d be any constants. Then, there exists an algorithm for the d-$\mathcal{MSSC}$ problem and some constant N_0 such that the algorithm achieves an approximation factor of $(1+\delta)\alpha$ on all instances whose universe contains at least N_0 items.*

PROOF. Given a d-$\mathcal{MSSC}$ instance $\mathcal{M}$ over a universe containing N elements, construct a binary table $\mathcal{D}$ using the (reduction) procedure described before. Use the α-approximation algorithm to obtain a solution $\mathcal{T}$ for $\mathcal{D}$. Apply Lemma 4.6 to transform $\mathcal{T}$ into a solution π for $\mathcal{M}$. Let $\mathcal{T}^*$ and π^* denote the optimal solution for $\mathcal{D}$ and $\mathcal{M}$, respectively. By Lemma 4.5 and 4.6, we can relate $\mathrm{cost}(\pi)$ and $\mathrm{cost}(\pi^*)$ as follows.

$$\begin{aligned}\mathrm{cost}_a(\pi) &\leq |\mathcal{T}|_a + 1\\ &\leq \alpha|\mathcal{T}^*|_a + 1\\ &\leq \alpha \cdot \mathrm{cost}_a(\pi^*) + \alpha\lceil\log d\rceil + 1\end{aligned}$$

We shall choose a suitable N_0 such that $(\alpha\lceil\log d\rceil + 1) \leq \alpha\delta(\mathrm{cost}_a(\pi^*))$ or equivalently, $\mathrm{cost}_a(\pi^*) \geq (\alpha\lceil\log d\rceil + 1)/(\alpha\delta)$. This would imply that the algorithm achieves an approximation factor of $\alpha(1+\delta)$ on all instances having at least N_0 items. This task is accomplished by applying Proposition 4.7, which says that $\mathrm{cost}_a(\pi^*) \geq N/2d$. So, we fix

$$N_0 = \left\lceil \frac{2d(\alpha\lceil\log d\rceil + 1)}{\alpha\delta} \right\rceil. \qquad \square$$

We now prove Theorem 4.4. Suppose there exists an α-approximation algorithm for the 2-$\mathcal{UDT}$ problem for some constant $\alpha < 2$. Choose $\delta > 0$ such that $(1+\delta)\alpha < 2$ and let $\beta = (1+\delta)\alpha$. Invoke Theorem 4.2 to obtain a constant d such that it is NP-hard to approximate d-$\mathcal{MSSC}$ within a factor of β. Now by Lemma 4.8, there exists an algorithm for the d-$\mathcal{MSSC}$ problem that has an approximation ratio of β on all instances over a universe of size at least N_0. For instances having a smaller universe, we can perform an exhaustive search in polynomial time, since N_0 is a constant. This means that NP = P. We have proved the theorem.

5. RAMSEY NUMBERS AND ERDÖS'S CONJECTURE

In this section, we take a closer look at our approximation ratio and discuss its connection to a Ramsey-theoretic conjecture by Erdös. We presented an algorithm for the $\mathcal{WDT}$ problem having an approximation ratio of $O(r_K \log N)$. Let us now focus on bounds for the inverse Ramsey numbers r_n, for $n \geq 1$.

Recall that for any k, $R_k \geq \frac{3^k+1}{2}$ [Nesetril and Rosenfeld 2001; Schur 1916]. From this we get that for any n, $r_n \leq 2 + 0.64\log n$. Notice that any improvement in the upper bound of r_n would automatically improve our approximation ratio. Better upper bounds are known for r_n (see Nesetril and Rosenfeld [2001], Exoo [1994], and Chung and Grinstead [1983]); but they improve the preceding bound only by constant factors. We observe that the upper bound for r_n cannot be improved significantly because of the following result: $R_k \le 1 + k!e$ [West 2001], which implies $r_n = \Omega(\log n / \log\log n)$.

Observe that our approximation ratio actually involves $\widetilde{r}_n$, rather than r_n. Therefore, one can try to derive a better upper bound on $\widetilde{r}_n$. Unfortunately, we show that $\widetilde{r}_n = \Omega(\log n / \log\log n)$. The claim is implied by the following theorem which can be proved based on an argument similar to the one used to obtain the same bound for R_k.

THEOREM 5.1. *For any k, $\widetilde{R}_k \le 1 + k!e$.*

PROOF. Let $n = \widetilde{R}_k - 1$. By the definition of $\widetilde{R}_k$, $\widetilde{G}_n$ can be directed Ramsey colored using only k colors. Let $\widetilde{\tau}$ be such a coloring. Pick any vertex from $\widetilde{G}_n$, say the vertex u. We first claim that $\widetilde{\tau}$ can be transformed to be symmetric with respect to u, meaning we can modify $\widetilde{\tau}$ in such a way that for any other vertex x, the edge (x, u) gets the same color as (u, x). This is accomplished by considering each vertex x and (locally) relabeling the colors assigned to its outgoing edges such that (x, u) gets the same color as (u, x). This does not increase the number of colors. From now on, it is assumed that we have modified $\widetilde{\tau}$ in the preceding manner.

There are $n-1$ outgoing edges from u, which are colored using k colors. So, there must exist a color class c having at least $(n-1)/k$ edges, that is, there should exist a color c such that the set of vertices $V' = \{x | \widetilde{\tau}(u, x) = c\}$ satisfies the inequality $|V'| \ge (n-1)/k$. The main observation is that for any $x, y \in V'$, the edge (x, y) cannot have c as its color. The observation can be seen as follows. We have $\widetilde{\tau}(u, x) = \widetilde{\tau}(u, y) = c$, and so, $\widetilde{\tau}(x, y)$ should be different from $\widetilde{\tau}(x, u)$, by the definition of directed Ramsey colorings. On the other hand, $\widetilde{\tau}(x, u) = c$, because of the transformation that we performed. Therefore $\widetilde{\tau}(x, y) \ne c$. To summarize, we have argued that the color c is not assigned to any edge in the subgraph induced by V'. Therefore, only $k-1$ colors are used in for the edges of the previous subgraph. It follows that $|V'| \le \widetilde{R}_{k-1} - 1$. Putting together the lower and upper bounds on $|V'|$, we get that $(n-1)/k \le |V'| \le \widetilde{R}_{k-1} - 1$. Hence, $n \le k(R_{k-1} - 1) + 1$. Since $n = \widetilde{R}_k - 1$, we have established the following recurrence relation on the directed Ramsey numbers. We have

$$\widetilde{R}_k \le 2 + k(\widetilde{R}_{k-1} - 1),$$

with the boundary condition being $\widetilde{R}_1 = 3$. By solving the recurrence relation, we get that for any $k \ge 1$,

$$\widetilde{R}_k \le 2 + k! \sum_{i=0}^{k-1} \frac{1}{i!}.$$

The RHS is at most $(1 + k!e)$. The theorem is proved. □

Notice that there is a gap in the upper and lower bounds for R_k. Erdös conjectured that for some constant α, for all k, $R_k \le \alpha^k$. This is equivalent to $r_n = \Omega(\log n)$.

We discuss the implication of our results in possibly proving the conjecture. The idea is to show that, in terms of worst-case performance factors, the rounded greedy algorithm performs poorly! We observe that a lower bound of $\Omega(\log K \log N)$ on the approximation ratio for the rounded greedy algorithm would imply the conjecture. More explicitly, we note that the following hypothesis implies the conjecture.

Hypothesis. There exists a constant $\beta > 0$ such that for any K, there exists a K-$\mathcal{WDT}$ table $\mathcal{D}$ and a weight function $w(\cdot)$ on which the tree $\mathcal{T}$ produced by the rounded greedy algorithm is such that $w(\mathcal{T}) \ge (\beta \log K \log N) w(\mathcal{T}^*)$, where $\mathcal{T}^*$ is the optimal solution.

A result by Garey and Graham [1974] could be a starting point for constructing such instances. They analyzed the worst-case performance of the greedy procedure for the 2-$\mathcal{UDT}$ problem and by constructing counter-examples, obtained a lower bound of $\Omega(\log N/\log\log N)$ for the approximation ratio of the procedure.

One can also attempt to prove the conjecture under the assumption NP $\neq$ P by showing that it is NP-hard to approximate the K-$\mathcal{WDT}$ within a factor of $\Omega(\log K \log N)$. More precisely, exhibit a constant $c > 0$ and show that for all $K \geq 2$, it is NP-hard to approximate the K-$\mathcal{WDT}$ problem within a factor of $c \log K \log N$. However, as mentioned in the Introduction, extending the $O(\log N)$-approximation algorithm for the $\mathcal{UDT}$ problem by Chakaravarthy et al. [2009] to the weighted case will rule out the this approach.

6. CONCLUSION AND OPEN PROBLEMS

We studied the problem of constructing good decision trees for entity identification, in the general setup where attributes are multivalued and the entities are associated with probabilities. We designed an algorithm and proved an approximation ratio involving Ramsey numbers, and also presented hardness results.

There are several interesting open questions. An obvious avenue is to bridge the gap between the approximation ratio and the hardness factor for 2-$\mathcal{UDT}$, K-$\mathcal{UDT}$, and $\mathcal{WDT}$.

The directed Ramsey numbers $\widetilde{r}_n$ introduced in this article pose challenging open problems: Is $\widetilde{r}_n = r_n$, for all n? Is $\widetilde{r}_n = O(\log n/\log\log n)$? Proving the second statement in the affirmative would improve our approximation ratios. If both the statements are shown true then the conjecture by Erdös would be disproved! Finally, it would be interesting, if the conjecture can be proved using the approach suggested.

ACKNOWLEDGMENTS

We thank A. Guillory for useful discussions and the anonymous referees for helpful comments.

REFERENCES

Adler, M. and Heeringa, B. 2008. Approximating optimal binary decision trees. In *Proceedings of the 11th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems*. Lecture Notes in Computer Science, vol. 5171. Springer, Berlin, 1–9.

Chakaravarthy, V., Pandit, V., Roy, S., Awasthi, P., and Mohania, M. 2007. Decision trees for entity identification: Approximation algorithms and hardness results. In *Proceedings of the 26th ACM Symposium on Principles of Database Systems*. ACM, New York, 53–62.

Chakaravarthy, V., Pandit, V., Roy, S., and Sabharwal, Y. 2009. Approximating decision trees with multiway branches. In *Proceedings of the 36th International Colloquium on Automata, Languages and Programming*. Lecture Notes in Computer Science, vol. 5555. Springer, Berlin.

Chung, F. and Grinstead, C. 1983. A survey of bounds for classical Ramsey numbers. *J. Graph Theory 7*, 25–37.

Dasgupta, S. 2005. Analysis of a greedy active learning strategy. In *Proceedings of the 17th Annual Conference on Neural Information Processing Systems*. MIT Press, Cambridge, MA, 337–344.

Exoo, G. 1994. A lower bound for Schur numbers and multicolor Ramsey numbers. *Electron. J. Combin. 1,* R8.

Feige, U., Lovász, L., and Tetali, P. 2004. Approximating min sum set cover. *Algorithmica 40,* 4, 219–234.

Garey, M. 1970. Optimal binary decision trees for diagnostic identification problems. Ph.D. thesis, University of Wisconsin, Madison.

Garey, M. 1972. Optimal binary identification procedures. *SIAM J. Appl. Math. 23,* 2, 173–186.

Garey, M. and Graham, R. 1974. Performance bounds on the splitting algorithm for binary testing. *Acta Inf. 3*, 347–355.

Graham, R., Rothschild, B., and Spencer, J. 1990. *Ramsey Theory*. John Wiley & Sons, New York.

Heeringa, B. 2006. Improving access to organized information. Ph.D. thesis, University of Massachusetts, Amherst.

Heeringa, B. and Adler, M. 2005. Approximating optimal decision trees. Tech. rep. TR 05-25, University of Massachusetts, Amherst.

HYAFIL, L. AND RIVEST, R. 1976. Constructing optimal binary decision trees is NP-complete. *Inf. Process. Lett. 5,* 1, 15–17.

KOSARAJU, S., PRZYTYCKA, M., AND BORGSTROM, R. 1999. On an optimal split tree problem. In *Proceedings of the 5th International Workshop on Algorithms and Data Structures*. Lecture Notes in Computer Science, vol. 1272. Springer, Berlin, 69–92.

MORET, B. 1982. Decision trees and diagrams. *ACM Comput. Surv. 14,* 4, 593–623.

MURTHY, S. 1998. Automatic construction of decision trees from data: A multi-disciplinary survey. *Data Mining Knowl. Discov. 2,* 4, 345–389.

NESETRIL, J. AND ROSENFELD, M. 2001. I. Schur, C.E. Shannon and Ramsey numbers, a short story. *Discr. Math. 229,* 1-3, 185–195.

PANKHURST, R. 1970. A computer program for generating diagnostic keys. *Comput. J. 13,* 2, 145–151.

RADZISZOWSKI, S. 1994. Small Ramsey numbers. *Electron. J. Combin. 1,* 7.

RAZ, R. AND SAFRA, S. 1997. A sub-constant error-probability low-degree test, and a sub-constant error-probability PCP characterization of NP. In *Proceedings of the 29th ACM Symposium on Theory of Computing*. ACM, New York.

REYNOLDS, A., DICKS, J., ROBERTS, I., WESSELINK, J., IGLESIA, B., ROBERT, V., BOEKHOUT, T., AND RAYWARD-SMITH, V. 2003. Algorithms for identification key generation and optimization with application to yeast identification. In *Proceedings of EvoWorkshops*. Lecture Notes in Computer Science, vol. 2611. Springer, Berlin, 107–118.

SCHUR, I. 1916. Uber die kongruenz $x^m + y^m \equiv z^m \pmod{p}$. *Jber. Deustsch. Math. Verein 25*, 114–117.

WEST, D. 2001. *Introduction to Graph Theory*. Prentice Hall.

WIJTZES, T., BRUGGEMAN, M., NOUT, M., AND ZWIETERING, M. 1997. A computer system for identification of lactic acid bacteria. *Int. J. Food Microbiol. 38,* 1, 65–70.

Received February 2008; revised December 2008; accepted April 2009