

Adapting ADtrees For Improved Performance on Large Datasets with High Arity Features

Robert Van Dam, Irene Langkilde-Geary and Dan Ventura

Computer Science Department
Brigham Young University
rvandam00@gmail.com, i.l.geary@gmail.com, ventura@cs.byu.edu

Abstract. The ADtree, a data structure useful for caching sufficient statistics, has been successfully adapted to grow lazily when memory is limited and to update sequentially with an incrementally updated dataset. However, even these modified forms of the ADtree still exhibit inefficiencies in terms of both space usage and query time, particularly on datasets with very high dimensionality and with high arity features. We propose four modifications to the ADtree, each of which can be used to improve size and query time under specific types of datasets and features. These modifications also provide an increased ability to precisely control how an ADtree is built and to tune its size given external memory or speed requirements.

Keywords: ADtree, Sufficient Statistics, Caching, Large Data Sets, High Arity Feature

1. Introduction

As fields such as machine learning and data mining continue to mature and are applied to new problems, the need arises to model and manipulate datasets that are substantially larger than those that have traditionally been attempted. The higher dimensionality of many of these datasets exceeds the capacity of current techniques for storing model parameters, but higher dimensionality is correlated with more relevant, meaningful, and interesting problems.

This explosion of data has become manifested in a variety of knowledge and information processing areas and recent advances have been demonstrated in

Received Oct 01, 2010

Revised Nov 17, 2011

Accepted Feb 25, 2012

algorithms for tackling variations of the problem in many domains. For classification tasks, training data must be compressed into a model for decision making. This can be challenging if the data cannot fit in memory, and one approach to handling this case is a piecewise training algorithm (for linear SVM in particular) on blocks of training data that produces a single model for classification (Yu, Hsieh, Chang and Lin, 2010). The problem of predicting rare (outlier) events in massive amounts of data has recently been applied in both the context of fraud detection, where condensed representations are employed (Koufakou, Secretan and Georgiopoulos, 2011) and of advertising where a hierarchical approach admits scalability (Agarwal, Agrawal, Khanna and Kota, 2010). Computing similarities or distances in very high-dimensional spaces is becoming common place as large repositories of multimedia data become available. Recent techniques for efficiently computing similarities of such data include preprocessing with a dimensionality reduction step and computing an approximate distance in low dimensional space (Kim, Chung, Lee and Kim, 2011) and clever indexing schemes that allow efficient bucketing of similar data into equivalence classes (Zhang and Alhajj, 2010). A related piece of work addresses the problem in a different way by considering the case when the distance function is the computational bottleneck and applies indexing of the distance function itself (by precomputing key values) to quickly acquire bounds for the real distance (Chen, Liu, Furuse, Yu and Ohbo, 2011). Processing and mining data streams is another instantiation of the massive data problem, and becomes even worse when users are allowed different sub-space views of the stream (Huang, Sun and Wang, 2010).

The problem we consider here is one of counting sufficient statistics, specifically in the context of statistical approaches to Natural Language Processing. Such approaches are well-known to require very large corpora to support statistics gathering. Further, many NLP approaches treat data as living in very high-dimensional spaces (because objects consists of many words) and many of these dimensions as having very high arity (because each word can take on many values). Unfortunately, most existing models suffer from practical (if not theoretical) problems of tractability with high dimensional (and high arity) data. The problem is particularly acute for statistical models that can make use of a full joint probability distribution, because these generative models require the ability to model the probability of any potential event or arbitrary tuple of feature/value combinations. Early efforts for addressing this counting problem include indexing structures such as the BSP tree (Fuchs, Kedem and Naylor, 1980; Gaede and Gunther, 1998), the kd-tree (Bentley, 1975; Gaede and Gunther, 1998) and the SE-tree (Rymon, 1993), but all of these demonstrate significant deficiencies for many problems that involve massive amounts of high dimensional, high arity data.

A caching data structure called the ADtree has been proposed that overcomes many of the time and space problems suffered by these earlier approaches (Moore and Lee, 1998); it is a generalization of the kd-tree that provides access to all (statistical) events while still improving memory usage. Additionally, the ADtree can benefit from both independence and correlations among its features in order to create a structure that is orders of magnitude smaller than its theoretical upper bound and capable of meeting memory requirements for higher dimensional datasets. ADtrees have proven useful in a variety of settings including for learning association rules (Anderson and Moore, 1998), for n -gram counting in sequential data (such as large natural language corpora) (Van Dam

and Ventura, 2007), and for incremental environments in which data is collected over time (Roure and Moore, 2006).

The ADtree provides a substantially improved mechanism over naïve methods for algorithms that require frequent counting over a dataset. Unfortunately, the problem of counting is fundamentally an exponential problem and many datasets are still much too large even for an ADtree. The structure that Moore and Lee presented was designed to be built in its entirety and then used by the algorithm, after which it could either be stored for future use or simply discarded. However, it was later noted that what are usually termed *static* ADtrees (those that are completely built before being used) often waste space (and therefore waste time in generation) on parts of the tree that will not be used by the algorithm for which the tree was built (Komarek and Moore, 2000). If these parts of the tree were easily identifiable *a priori*, then a simple modification could create the tree without including the unnecessary subtrees.

The unnecessary parts of the tree can often be the result of a complicated interaction between the specific machine learning algorithm and dataset (that is, changing either the algorithm or the dataset would change which parts of the tree are extraneous). Komarek and Moore’s solution was to create a dynamic adaptation of the ADtree that functions as a lazy caching structure that grows the tree only when presented with specific requests (queries) that are not already cached. Their adaptation again makes use of several clever techniques to avoid repeatedly iterating over the dataset as might otherwise seem necessary.

Static ADtrees provide a memory-efficient mechanism for utilizing large datasets with high dimensionality in machine learning algorithms. Dynamic ADtrees extend those capabilities even further. However, real-world sized datasets can still be difficult or impossible to manage as they can still generate an ADtree that does not fit into memory. Furthermore, although dynamic ADtrees can be applied to high arity features, they give extremely poor average access times and very minimal space savings on these attributes since one of the space-time tradeoffs used to reduce the size of the tree is no longer applicable for high arity features. Additionally, the trade-offs between the amortized runtime and re-usability of static ADtrees and the space-savings and flexibility of dynamic ADtrees remain relatively unexplored.

We present four modifications to the ADtree that address these issues and that improve performance on large, high-arity datasets: new indexing strategies (Van Dam, Langkilde-Geary and Ventura, 2008), static/dynamic tree hybridization, feature ordering strategies, and pruning techniques.

2. Background

The full ADtree contains two types of nodes that alternate along every path in the tree. *ADnodes* store the count of one conjunctive query. *Vary nodes* group ADnodes according to a single feature. The Vary node child of an ADnode for feature a_i has one child for each value v_j . These grandchildren ADnodes specialize the grandparent’s query Q by storing the counts of $Q \wedge a_i = v_j$. Such a full ADtree contains every combination of feature-value pairs and is not yet efficient in its memory usage.

The original ADtree included three approaches that can be combined to reduce its overall size. First, the tree is sparsified by removing all zero counts. Second, the ADnodes near the bottom of the tree are not expanded. Instead

they are replaced with “leaf lists” of indices into the dataset whenever the number of relevant rows (the count) drops below a pre-determined threshold. Third, certain counts are removed from the tree in such a way that they can be recovered from other counts still stored in the tree. The trick is for each Vary node to remove the child ADnode (called the MCV) that has the largest count among all its child nodes, providing the largest expected space-savings without sacrificing the ability to recover any counts. This last technique assumes that the trade-off of increasing average query time (due to recovering the implicitly stored counts) is reasonable in comparison to the expected space savings.

The dynamic tree follows the same basic structure; however, since it is built lazily, at any given point only a portion of the tree will have been expanded. The dynamic tree also contains some additional support information used to temporarily cache information needed for later expansion. Fully expanded portions of the tree no longer require these extra support nodes.

2.1. WSJ Dataset and Constraint-based Part-of-Speech Tagger Client

Because both the space and time requirements of the ADtree depend on the sequence of queries it is used to answer, testing our proposed modifications requires generating various query sequences. For this purpose, we use a client algorithm that performs part-of-speech tagging on a modified version of the Wall Street Journal (WSJ) corpus (*University of Pennsylvania Linguistic Data Consortium*, n.d.). In order to isolate the timing and memory usage of just the ADtree building code, the tagger was used only to generate query logs for each feature set. All of the experiments reported here were performed using only the query logs to generate the ADtrees, and each query log can be thought of as a unique “dataset”. The nature of the tagging task results in the generation of query logs that exhibit high variability and thus constitute an ideal platform for testing the ADtree because such variability in general requires fast access to all possible queries.

The WSJ corpus includes each word along with its Part-of-Speech (POS) tag as well as the syntactic structure of each sentence. The dataset used here is a modified version that maintains the word and POS tag as features but ignores the syntactic sentence structure. Four additional features were derived from the POS tags: category (simplified POS class), subcategory, tense, and number. Seven additional features were derived solely from the word itself: whether or not the word was “rare” in the dataset, capitalization style, orthographic representation (e.g. contractions), has-hyphen, and has-digit, the last 3 characters of the word, and if it contained one of a short list of derivational suffixes. Several of these features have been used in previous part-of-speech taggers (Toutanova and Manning, 2000). Each of these features breaks what would be a single tagging decision into a set of smaller, related decisions. In addition, three more features were added to the new dataset related to morphology. Those features were a base word form and the corresponding inflectional morphological rule (pattern) as well as whether the base stem was also “rare”.

This enhanced WSJ corpus is large enough, in terms of dimensionality and the arity of its features, that building a static ADtree is not feasible. Table 1 shows the number of unique values in the corpus for each feature. This corresponds to slightly less than 10^{23} possible events (or conjunctive queries). The “code”

Table 1. All available features, their arity (number of values) and how skewed is the Most Common Value (MCV). ‘Original Word’ (‘w’) and ‘Part of Speech’ (‘c’) come from the original WSJ corpus, the remaining features are derived from one or both of these (except for the two ‘Rare’ features, which are derived from the corpus as a whole). “MCV skew” measures the independent skew of each feature across the entire dataset. When calculated given a conditional context of other feature/value assignments, this measure can change dramatically (as can the MCV itself).

Code	Feature	Values	MCV Skew
w	Original Word	42136	5%
b	Word Stem	35698	5%
3	Last 3 Characters	4591	5%
p	Inflectional Suffix	255	94%
c	Part of Speech	45	14%
0	Simplified POS	22	32%
1	POS subcategory	15	38%
s	Derivational Suffix	10	81%
t	Tense	5	86%
C	Capitalization Style	5	86%
n	Number	3	53%
z	Orthographic Style	3	98%
h	Has Hyphen	2	99%
8	Has Digit	2	97%
o	Rare Word	2	95%
a	Rare Stem	2	94%

Table 2. Each code represents a subset of features (+ means that neighboring words were included). There is not a ‘+’ feature set corresponding to every plain feature set because the client algorithm could not complete enough sentences to create usable query logs.

Feature Sets	Event Space
wc	1.90×10^6
wcCh8	3.79×10^7
wc01tn	9.39×10^9
wc01tns3	4.31×10^{14}
wbcp	1.73×10^{13}
wbcpCh8	3.45×10^{14}
wbcps3	7.92×10^{17}
wbc01tnps3	3.92×10^{21}
wbc01tnzps3Ch8ao	9.41×10^{23}
wc+	6.82×10^{18}
wcCh8+	1.36×10^{20}
wc01tn+	6.75×10^{21}
wc01tns3+	3.10×10^{26}
wbcp+	6.21×10^{25}
wbc01tnps3+	3.07×10^{29}
wbc01tnzps3Ch8ao+	7.37×10^{31}

character for each feature is used as shorthand to represent different groupings of features (see Table 2).

Several different groupings of features were chosen, in part because they represent natural feature groupings for the tagger client application but also as a mechanism for distinguishing how the proposed modifications interact with features of different arities. Table 2 lists several different feature sets by a short-

hand “name”, which consists of one character for each feature in the set using the single character “codes” in Table 1. Feature sets that were also combined with the word and POS tag for both the left and right words (adding four extra features to each group) are distinguished in the results using a ‘+’ at the end of the name. Each of these different feature sets produces a unique query log and thus a unique dataset for testing the performance of the ADtree and our modifications.

3. ADtree Modifications

We propose four different modifications to the ADtree to facilitate its use for large, high-arity datasets, with each modification depending upon slightly varying data and usage characteristics. The situations in which each modification can be applied are detailed both in this section and in Section 4. It is worth noting that the ADtree used in most of these cases is a dynamic version of the tree. However, many of these improvements can also be used in a static tree (although that sometimes implies a change in the expected space-time tradeoff). The applicability of each modification to a static ADtree will be discussed as part of each of the following sections.

3.1. Improved Space-Time Tradeoffs for High Arity Features

It is noteworthy that the majority of previously published results using ADtrees have used datasets with relatively low arity features. Although there is no technical limit to the arity of features used in an ADtree, there are practical limitations. In particular, the space-saving technique of removing all MCV ADnodes makes the assumption of a reasonable space-time trade-off that is more easily violated with high arity features. Dynamic trees can suffer additional problems when removing MCV nodes even for low arity features. Below we will present two modifications to the ADtree that decrease the overall space usage when using high arity features while still maintaining reasonable build/query times.

The key problem is that removing the MCV ADnode for high arity features can dramatically increase query time and only slightly decreases space usage. If a query involves the MCV of feature a_i , it is necessary to sum at least $n_i - 1$ values (where n_i is the arity of a_i). Given a query Q that specifies values for q features, the worst case scenario could require summing over $\prod_{i=1}^q (n_i - 1) + 1$ values due to recursive MCV “collisions”. Although this worst case is rare (since data sparsity and/or correlation tend to cause Vary nodes lower in the tree to have fewer ADnodes), it helps to illustrate the potential for longer query times when the n_i are large. For instance, a query involving 3 features with 10 values each has a worst case of summing over 9^3 values. If the 3 features had 10,000 values, this becomes 9999^3 .

Figure 1 illustrates the “best” case for a feature of 10,000 values in which every sibling node is a leaf node (or at least a leaf list). If a query that encounters that MCV ADnode involves further features in the missing subtree then it will need to traverse the subtrees (if they exist) of all 9999 siblings. This increases the potential for further MCV “collisions”, eventually leading to the worst case scenario where every path involves one MCV ADnode for every feature in the query.

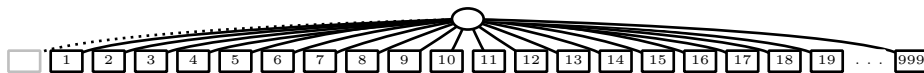


Fig. 1. Iconic representation of a vary node for a high arity feature ignoring any further subtrees. This effectively illustrates the “best case” scenario with respect to the base problem of summing over so many counts in order to retrieve a single number. A real-world situation could be far worse if the subqueries had to descend subtrees and potentially encounter more missing MCV nodes. (Although node size is not represented graphically, the only assumption is that the MCV is not highly skewed).

Although there are many more queries that don’t require summing, higher arity increases the possibility for very slow queries. Additionally, since the more common values have a higher prior probability of being queried (at least for some client algorithms), the most expensive queries are the most likely to occur (or more likely to occur often).

The original justification for excluding MCV ADnodes was that it provided significant space savings in exchange for a minor increase in query time. For binary features, there is no significant increase in query time and each Vary node’s subtree is cut at least in half. For a feature with 10 values, the Vary node’s subtree is reduced by at least 10% in exchange for summing over at least 9 values. With 10,000 values, the subtree is reduced by at least 0.01% in exchange for summing over at least 9999 values. The space savings becomes increasingly insignificant and conversely the query time becomes increasingly significant with higher arity features.

There is a further complication when using a dynamic tree. If a query is made that involves at least one MCV then all $n_i - 1$ of its siblings must be expanded. If those nodes are never directly queried then they are simply wasted space. (In the previous example of 3 features with 10,000 values there are 9999^3 unique paths consisting of nearly 10^{12} nodes¹). Although not as dramatic for lower arity features, this is true for any arity above two². Therefore, if the algorithm does not query all combinations exhaustively then excluding all MCV ADnodes can actually create a larger tree than if they had been included.

3.1.1. Solution 1: Complete Vary Nodes and RatioMCVs

The first solution is to establish a threshold for high arities, above which the MCV ADnode is included. Note that this does not entail making a list of features for which the MCV will always be included. The “arity” used to determine whether or not to include the MCV is the number of children of the parent Vary node. This means that the exclusion of MCVs becomes context sensitive and a given feature may have its MCV included in one part of the tree and excluded in another (just as the actual MCV is context dependent). The threshold, however, is kept constant throughout the tree. A Vary node that has an intact MCV ADnode child will be referred to as a Complete Vary Node.

Although an arity threshold should help provide an improved balance of space and time, there are still some cases where using a simple threshold might

¹ There are $9999^3 + 2 \cdot 9999^2 + 2 \cdot 9999 = 999,900,009,999$ total nodes (AD and Vary) needed.

² With binary features, it is always redundant to include both nodes so we keep the smaller (non-MCV) one.

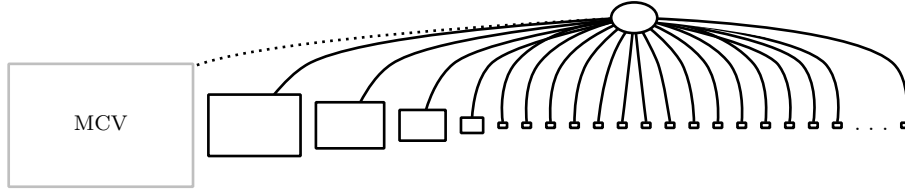


Fig. 2. Vary node for a feature with skewed MCV (high proportional count). Relative node size represents the relative counts at each node (as well as the *a priori* potential size of the corresponding subtrees). Although it is still necessary to sum over many nodes to recover some queries, the MCV subtree is so large and the potential for further MCV collisions in the smaller subtrees so small that the trade-off is more reasonable (contrast Figure 1).

incorrectly indicate the need to include the MCV. In particular, if a feature is highly skewed due to a large MCV (given the context in the tree) then the MCV may represent a significant portion of a Vary node’s subtree. In such a case, the subtree of the MCV may be so large that including it would be too expensive, even if it has a large number of siblings. For instance, including the MCV ADnode shown in Figure 2 based solely on arity while ignoring its disproportionately large count could be worse than leaving it out (as it would have been by default).

Indeed, a strong skew towards the MCV implies smaller counts for its sibling nodes and therefore shorter subtrees. This eliminates the possibility of the worst case scenarios noted above and mitigates the need to include the MCV ADnode. We therefore introduce the concept of a RatioMCV that is an MCV that has a count above a parameterized threshold ratio of the total of the grandparent ADnode. Then, any MCV that is also a RatioMCV will be removed, regardless of the number of sibling values. The combination of arity threshold and RatioMCV should appropriately handle all possible situations since the thresholds can be adjusted to accommodate different datasets and feature arities.

3.1.2. Solution 2: Clump nodes

As was noted, excluding MCV ADnodes can interact negatively with dynamic expansion of the tree, counter-intuitively increasing its size. Although the consequences of this interaction are worse for higher arity features, the problem can occur with any number of values above 2. Since the usefulness of excluding the MCV ADnodes is well documented (except as noted here), some other solution is needed.

One feasible solution stems from the fact that this problem does not occur with binary features. A query involving the MCV of a binary feature follows the same path as a similar query switching only the MCV value to its complement. If all of the sibling values of the MCV were “clumped” together as a single “other” value, then any query involving the MCV could be found using this Clump node and performing the appropriate subtraction. This eliminates the need to calculate a sum and does not expand any nodes not explicitly queried (see Figure 3). Furthermore, Clump nodes can be built on demand just like any other node in the tree.

There is a slight catch when using the Clump node, however. A Clump node cannot be used to answer queries involving any of the “clumped” values. Therefore, a query involving such a value will require the expansion of the corresponding node, resulting in some redundancy in the tree. One expensive way to deal

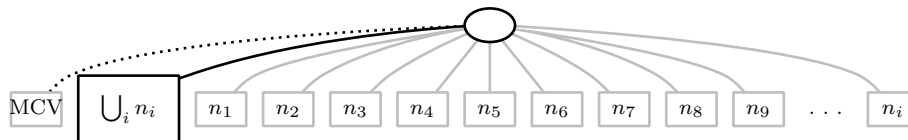


Fig. 3. Here the issue of skew is ignored and so the nodes are represented as uniform in size. However, the Clump node represents both the sum of the counts of the other nodes as well as the union of the rows (represented here) for that node’s partition of the dataset.

with this would be to delete the Clump node and create a new one without the newly expanded value. Using this method, the time needed to query the MCV will slowly increase until the Clump node is completely removed. The approach used in this paper takes the alternative approach of temporarily permitting the redundancy until some percentage of the “other” values are expanded. At that point, the Clump node is deleted.

3.2. Static and Dynamic ADtree Hybridization

Static trees provide excellent amortized conjunctive counting query retrieval time, particularly if they are reused multiple times. However, if a dynamic tree is required due to memory limitations, the tree must be regenerated on each run of the client algorithm.

Even if the series of queries produced by different runs of an application are not identical, it is likely that the ADtrees from each run share some structure. One reasonable solution is to hybridize the benefits of the static and dynamic trees by generating some of the shared portions of the tree in advance. This can reduce the time needed to dynamically generate the tree without wasting space on unneeded portions of the tree. In general, the first and second levels of the tree are the most likely to be needed by an arbitrary algorithm. These levels represent all possible single and paired feature queries and constitute partial computation of higher-order queries. So it would generally be beneficial to statically precompute these portions of the tree. Additionally, this hybridization could be parameterized to statically generate only a subset of features or values.

In many respects, this hybridization technique is a way to memoize a query-based analog to “currying”, a function generation technique used in some higher-order functional languages. The simplest form of currying takes an n -argument function and a value for one of those arguments and returns a new, $(n - 1)$ -argument function in which all occurrences of the first argument have been replaced with the given value. This can be generalized to replacing k of the original function’s arguments, where $k < n$, by repeatedly currying the result of the previous currying. The hybrid ADtree, essentially “curries” the query function for all values of each feature independently. The second level can then be generated by taking the results of the first level and currying each query/node with all possible values of a second, distinct feature³. The result is an efficiently stored,

³ Fortunately, the actual 2-level tree generation is more efficient than the currying analogy implies since all ADnodes (values) of a particular feature can be generated together in a single step.

memoized table of all (or a parameterized subset of) 2-feature full and partial queries.

3.3. Feature Ordering Strategies

The original ADtrees maintain a predetermined fixed order on features that determines the path that a query should follow through the tree. Any conjunctive query is first sorted according to this ordering and then the tree is traversed in that same order. This fixed order is generally a naive ordering such as the original order from the dataset or even a simple alphabetical ordering (assuming descriptive feature names exist). A more tailored ordering (to be used with either static or dynamic ADtrees) would determine the relative ordering of features that minimizes the overall size of the tree. Maximum benefit would come from a context-specific ordering because of the increased flexibility to use different relative orderings of features that best minimize the tree given the context.

A context specific ordering requires the use of a heuristic to determine, at each level of the tree, which feature is next along a particular path. A well-defined heuristic is necessary that can be used both for generating the tree as well as for determining the path that should be followed for any given query. Unfortunately, it is not known whether there is a generally optimal heuristic or even whether the same heuristic would work equally well for any two distinct datasets or applications. In particular, a metric that selects for strongest correlation may miss an opportunity to exploit independence and vice versa.

Besides the canonical and alphabetical orderings, several simple heuristics can be designed based on some of the characteristics of the ADtree. The two simplest characteristics on which to base an ordering are arity of the features and size (count) of the MCV of each feature. A third, and somewhat less obvious measure on which to base an ordering is the overall entropy of the values of a feature. Entropy in this context can be defined as

$$H(a_i) = -\sum_{j=1}^{n_i} p(v_j) \cdot \log_2(p(v_j)) \quad (1)$$

where the a_i is the i -th feature in the dataset, n_i is the arity of a_i and $v_1 \cdots v_{n_i}$ are the possible values of a_i . The value of $p(v_j)$, the conditional probability that v_j occurs given the current context in the tree, is estimated using the observed counts such that

$$p(v_j) = \text{count}(Q \wedge a_i = v_j) / \text{count}(Q) \quad (2)$$

where Q is the grandparent query (in the case that Q is the empty query, $\text{count}(Q)$ is defined to be the total number of rows in the dataset.) In essence, entropy is a measure of the context-specific distribution of values of a_i given Q . A uniform distribution will have maximal entropy, while a single value (and therefore terminating MCV) will have an entropy of 0.

In general, these three heuristics are not uncorrelated. For instance, (relatively) large MCV counts imply smaller sibling counts and therefore lower overall entropy. High arity features automatically have a higher potential maximum entropy (since the entropy of a perfectly uniform distribution is $\log_2(n_i)$) and so an ordering based on largest arity will be identical to one based on highest entropy given a set of independently and uniformly distributed features. However, such a dataset is not likely to be considered sufficiently interesting to perform an analysis on, and so it is reasonable to assume that an entropy-based ordering will

be more tailored to the specific distributions of a given dataset. That does not, however, imply that such a tailored ordering is inherently “better” in terms of the resulting size of the tree. In fact, it is not readily obvious in which “direction” to use each ordering (that is, will a lower or higher entropy, a larger or smaller MCV result in a smaller tree). The usefulness of each ordering will generally be very data dependent.

Superficially, ordering based strictly on arity essentially trades depth for width. If the largest arity features are ordered first, the levels of the tree after the first (which is always essentially fixed in size) will likely be very wide because so many feature/value combinations will occur in the first few levels. This has the fortunate side-effect that it partitions the fixed rows of the dataset quickly and therefore presumably, the tree will not grow very deep as strong correlations lead to terminal MCVs and smaller counts lead to terminal leaf-lists. Furthermore, strong correlations among the high arity features can greatly reduce the width of the tree because fewer combinations of the first few feature/value pairs are encountered. In the reverse ordering, the top levels of the tree are smaller because there are many fewer value combinations. However, if the reverse ordering is used and there is a strong correlation between the lowest arity features and the highest, there will be many occurrences at the lower levels of features with a single possible value determined by those correlations. As a result, the tree could be much smaller since it started out “narrow” and many paths terminated prematurely.

Similar analysis of the orderings based on MCV counts and entropy reveal that it could be very difficult (likely impossible) to establish a single, universally-optimal ordering. It may not even be possible to determine an optimal ordering for a given dataset unless the correlations among all features is known *a priori* (which implies a prior analysis of the dataset not unlike that accomplished by the process of building the ADtree). As a result, it may be necessary in many cases to try some or all of the orderings empirically. We will provide results of all five orderings (canonical, alphabetical, by MCV count, by feature arity, and by entropy) performed forward and reversed on the different subsets of features of the enhanced WSJ dataset with the expectation that varying degrees of correlation among the features in each set satisfactorily gauges each ordering’s effectiveness on different types of datasets.

3.4. Pruning in Limited Space Scenarios

Komarek and Moore noted that dynamic ADtrees could be extended beyond their implementation to provide both internal and external pruning mechanisms. External pruning could be as simple as providing an interface to permit the client application to specify to the ADtree which subtrees to prune. Internal pruning could be based on the raw internally stored counts⁴ or on a priority queue based on frequency of use as suggested by Komarek and Moore. It is not unreasonable to imagine a system that utilizes both mechanisms by pruning using default rules while also permitting the client to submit requests to prune specific subtrees.

⁴ Without additional knowledge of the client algorithm, the most reasonable prior probability that a node will be used is simply the Maximum Likelihood Estimate taken directly from the counts, which are readily available.

The fundamental problem of pruning nodes from the ADtree is to simultaneously maximize the freed memory gained from the pruning of a subtree while minimizing the risk of pruning some portion of the tree that will need to be regrown. It is also possible that the low-level implementation of the tree could associate a time cost with pruning (due to extra garbage collection for instance). This formulation of maximizing payoff versus minimizing risk provides some immediate insight into certain classes of subtrees that it would generally be disadvantageous to prune. Leaf nodes/lists provide a very minimal reward for freed space (assuming an implementation using only indices, not duplicated data rows). Nodes at the top level that have high counts would generally provide the largest reward but also imply the biggest risk since they (or a subtree) have a high probability of being needed again later.

On the other hand, if a high arity feature occurs early in the feature ordering then its subtrees will often be larger than subtrees of other features at that level because they have more features on which to split. However, the high arity feature’s lower count values have a lower likelihood of being needed again (and less penalty for regeneration). Therefore, any reasonable internal pruning mechanism would generally expect to avoid pruning the most expensive subtrees and would instead focus on infrequently used mid-level subtrees or top-level subtrees of high arity features with lower counts.

The results given in Section 4.4 represent a simple compromise between internal and external pruning. Given the above specifications for what kind of pruning is likely to be generally applicable, implementing the pruning itself is very straightforward. In this case, the pruning algorithm chosen is reasonably general but also noticeably biased by the particular dataset being used. The algorithm removes all ADnodes (values) with a count less than a pre-determined threshold (1000 was chosen somewhat arbitrarily) in the first level for high arity features. Since the majority of the high arity features for the enhanced WSJ are based on words, their distributions are generally similar to a power law and so the algorithm basically prunes the “long tail”. The difficulty in using this or any pruning algorithm really lies in determining an opportune moment in the processing to perform the pruning. If used as a purely internal mechanism, the easiest and most reasonable approach is to prune if and when memory usage exceeds a predetermined (or possibly parameterized) threshold, resulting in essentially an inverted-greedy pruning algorithm (greedy in the sense of using as much memory as possible). However, this is not likely to correspond to a “convenient” time with respect to the client application. The aforementioned compromise used in this case was to provide the client application with the ability to initiate pruning externally (although not with any external parameters as suggested above). Since the tagger application treats sentences as a whole, while a sentence is being tagged, the words of the sentence are repeatedly used in queries. Therefore, the completion of a sentence (or a set of sentences) seems the most natural event to trigger pruning in the tree. Other client algorithms might also have natural boundaries where pruning is least likely to remove needed counts. The results discussed here were obtained by pruning after every 100 sentences were finished tagging. With feature sets for which the tree grows only gradually, this is likely to provide a reasonable tradeoff between the relatively expensive operation of pruning and risk of running out of memory. With feature sets that cause the tree to grow more rapidly, this 100 sentence gap between prunings may not happen frequently enough to avoid known memory limitations.

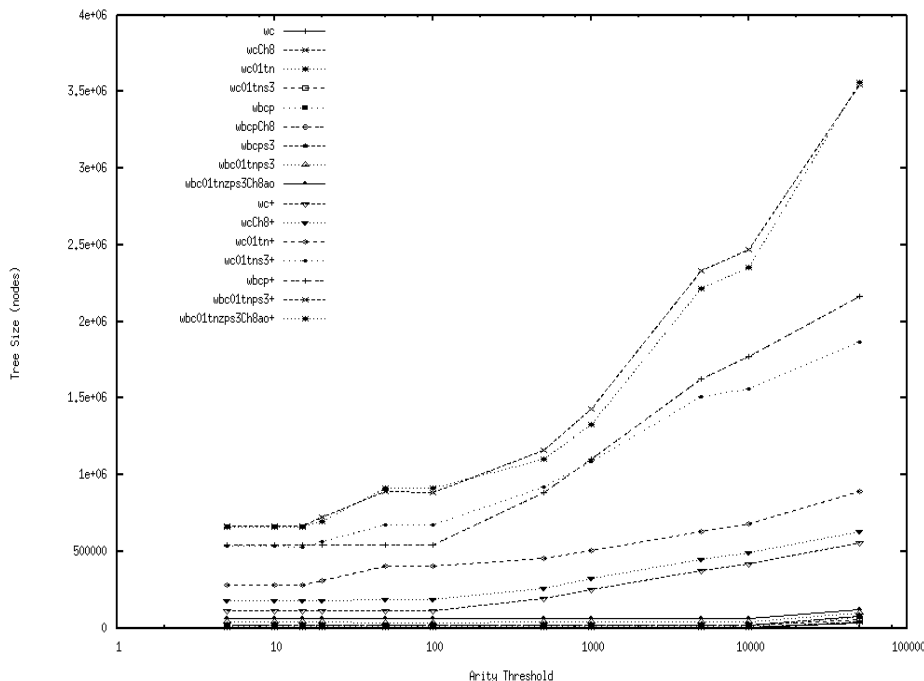


Fig. 4. Results of varying the arity threshold from 5 to 50000 on tree size (number of nodes) for each of the 16 feature sets. Notice the x -axis uses a logarithmic scale. The individual lines are less important than the overall trend. The large scale of the y -axis disguises the fact that most of the feature sets show a small upward curve at the low end of the x -axis.

4. Validation and Analysis of Results

Here we analyze the results of ADtrees for each of the datasets discussed in Section 2.1. Note that there are a few reported results for which specific datasets are either marked or excluded. Where the results are marked (such as with hash marks), this should be understood to mean that experiment terminated early, resulting in a slightly larger margin of error than the others. On the other hand, some experiments completed so few sentences (if any) that the margin of error was considered to be too large to report. In those cases, there may be a gap in a line graph or a missing bar in a bar graph. This does not necessarily imply that the missing dataset would always be too large or too difficult to use in conjunction with ADtrees and/or the specific modification being analyzed.

4.1. High Arity Features

The first set of experiments was designed to determine how the tree size varied with respect to the arity threshold. Eleven threshold values were chosen between 5 and 50000. Figure 4 shows the tree size (number of nodes) for each feature set as the arity threshold is varied. Additionally, it is important to note that since the highest arity feature has only slightly more than 40,000 values, an arity threshold of 50,000 is actually equivalent to an unmodified ADtree. One detail

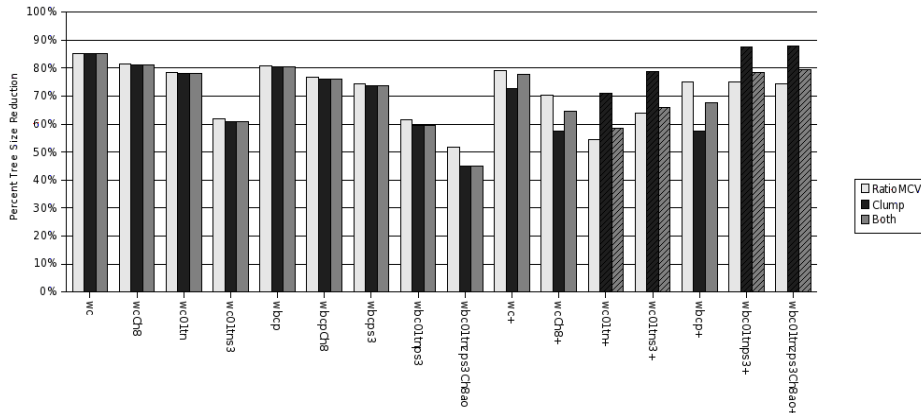


Fig. 5. Percent tree size reduction relative to the baseline (higher is better). Bars with hash marks did not finish constructing the tree due to memory complications.

obscured by the scale of the graph is that several of the feature sets exhibit a slight upturn at the low end of arity threshold. In fact, the lowest point is usually between threshold values of 20 - 50. This is of course a function of this dataset and the correlations among features and could vary greatly from one dataset to another. Additionally, the optimal tree size for a feature set does not always correspond to the lowest average query time. The remaining experiments use a threshold of 100 (when applicable) since most of the feature sets are still close to their optimal size at that threshold. It is, however, readily apparent from the graph that almost any choice of threshold will generally be an improvement over the baseline.

Once a reasonable value for the arity threshold was established, both the MCV ratio threshold and the Clump node deletion threshold were investigated. Each of these two thresholds correspond to a ratio that can vary between 50-99%. The MCV ratio threshold establishes how skewed the MCV of a high arity feature has to be before it will be included despite the feature's arity. The Clump node deletion threshold establishes how much (partial) redundancy is permitted before the Clump node is removed. Interestingly, it was not possible to establish a correlation between the exact value for either of these thresholds and their effectiveness. Both the size of the tree and average query time were found to be very robust with respect to these two ratio thresholds.

For this reason, it was decided to arbitrarily use a ratio threshold of 50%, meaning that the MCV is at least as large as the sum of all the other possible values. Similarly, the Clump node deletion threshold was also set at 50%. Three variations of the ADtree were built: using just the Complete Vary Node + RatioMCV modification, using just the Clump node modification, and using both modifications together. Then ADtrees were built for each combination of parameters (including the feature set) and the tree size (node count), elapsed runtime and cumulative memory usage were all measured. Elapsed runtime is a fairly reasonable estimate for average query time because each experiment represents only the time required to build the tree using the query logs recorded from the POS tagger. Cumulative memory usage is primarily dominated by the amount of memory used in temporary data structures needed to perform calcu-

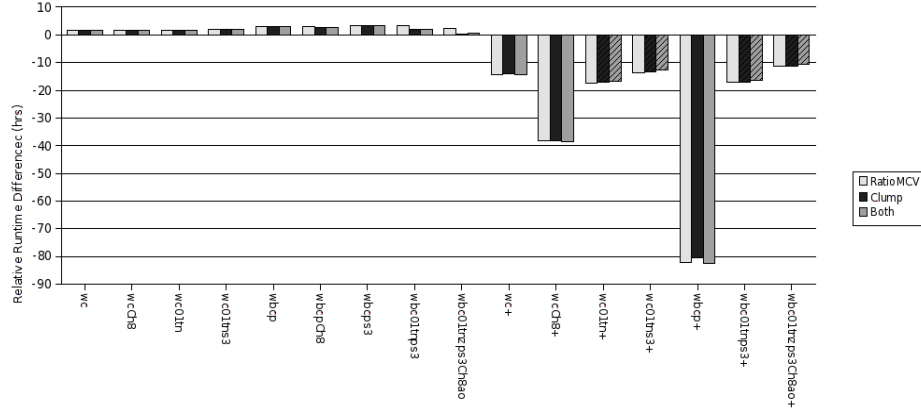


Fig. 6. Relative difference in runtime (in hours) from the baseline (lower is better). Bars with hash marks did not finish constructing the tree due to memory complications.

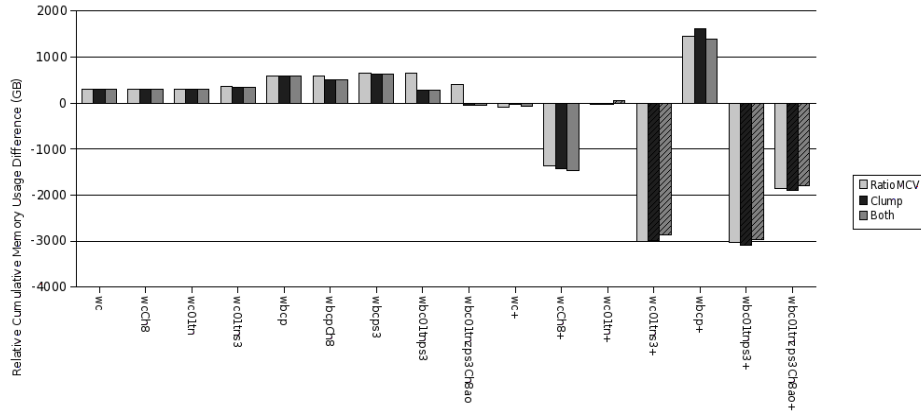


Fig. 7. Difference in cumulative memory usage (in GB) from the baseline (lower is better). Bars with hash marks did not finish constructing the tree due to memory complications.

lations while dynamically generating the tree. The results of these experiments are summarized in Figures 5-7.

In general, it is clear that both modifications achieve the desired results of reducing tree size without substantially increasing build/query time. In particular, the smaller feature sets (those not including neighboring words) obtained on average better reduction in tree size but take a few hours longer to construct the tree and use more cumulative memory. The feature sets that included the neighboring words on the other hand generally performed faster while still obtaining reasonable tree size reductions. Although there is some variation among feature sets, it is not unreasonable to conclude that either modification successfully mitigates the unintended consequences of using MCVs in an ADtree with high arity features.

Even though each modification reduces the tree size, the use of Complete Vary Nodes (and RatioMCVs) appears to achieve equal or better results in terms of

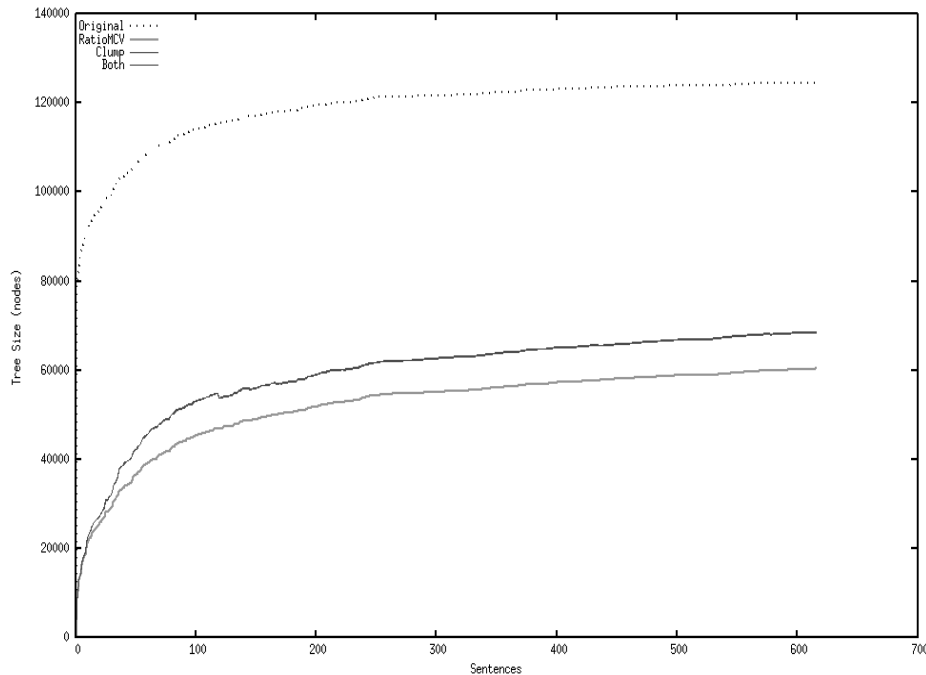


Fig. 8. Example of how tree size varies as sentences are tagged using the feature set `wbc01tnzps3Ch8ao`. Tree growth tapers off towards the right as fewer unique words (and correspondingly unique value assignments) are encountered. Although using a clump node does slightly increase tree size, it is still significantly less than the unmodified tree.

tree size but takes on average slightly longer to do so. This is most evident in Figures 8 and 9 as the performance corresponding to Complete Vary Node and RatioMCV with respect to the number of sentences completed slowly separates from that of the Clump node. However, there is a potential advantage to be gained by this difference in tree size versus elapsed time. Even though the original motivation for using both of these modifications was to improve performance on a dynamically generated ADtree, the RatioMCV modification can be applied just as easily to a statically built ADtree, resulting in the same tree size reduction. In that case, the cost of building the tree is a one time cost that can be amortized over all future uses of that tree. Furthermore, generating a static tree can be optimized in ways that can't be applied to a dynamic tree, potentially eliminating the extra time cost of using Complete Vary Nodes with RatioMCVs. Because the Clump node always represents 100% redundant information in a static tree, it is less easily justified (though still potentially useful if space is available).

It is also apparent from the results that combining both modifications together in the same ADtree only rarely gives any advantage over using whichever modification proved most effective for the given feature set. However, since it is not necessarily obvious before hand which modification will provide the best performance, using both seems to be an effective way to “hedge your bet” as to which modification will perform best (assuming that a dynamic tree will be used).

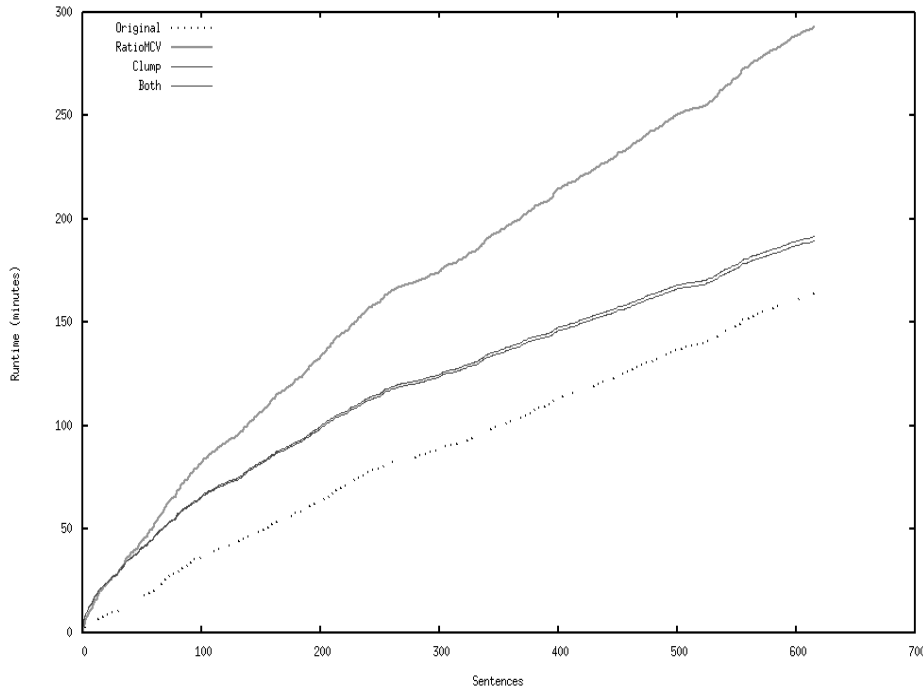


Fig. 9. Example of how much runtime elapses as sentences are tagged using the feature set `wbc01tnzps3Ch8ao`. The extra computation time needed for each modification is clear, particularly for the Complete Vary Node and RatioMCV.

From the cumulative memory usage metric, it is clear that the implementation of ADtrees used here is not optimal in memory utilization, particularly with respect to the temporary data structures that are built and then discarded very frequently in the course of dynamically building the tree. It may be possible to achieve nearly universal build/query time improvements (rather than the mixed results achieved here) by minimizing the amount of temporary memory used and therefore the amount of time spent doing garbage collection.

4.2. Hybridization

Because of the improvements achieved by the use of Complete Vary Nodes and RatioMCVs, the experiments using the hybrid version of the tree used Complete Vary Nodes and a RatioMCV threshold of 0.5. Correspondingly, Figures 10, 12, and 14 compare the results of creating a hybrid tree in combination with a Complete Vary Node/RatioMCV tree to the results in Figures 5-7 for just the Complete Vary Node/RatioMCV. As in the previous graphs, the results are relative to a baseline that uses none of these modifications. Additionally, Figures 11, 13, and 15 show just the improvement (or lack thereof) of just the hybrid tree using the Complete Vary Node/RatioMCV results as a new baseline.

In order to understand these results, it is important to understand that a single static tree was built in advance. This tree consists of the first two levels

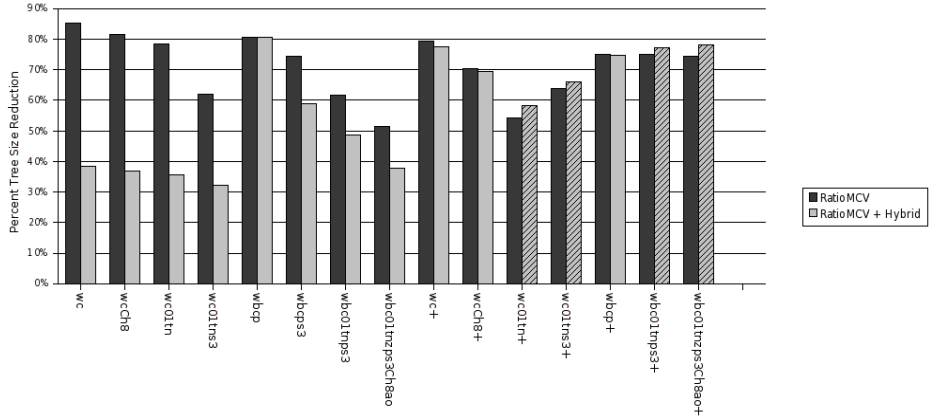


Fig. 10. Percent tree size reduction relative to the baseline (higher is better). Because the same partial static tree is used for each feature set, the smaller feature sets actually perform much worse. Bars with hash marks did not finish all of the queries in the respective query logs due to memory complications.

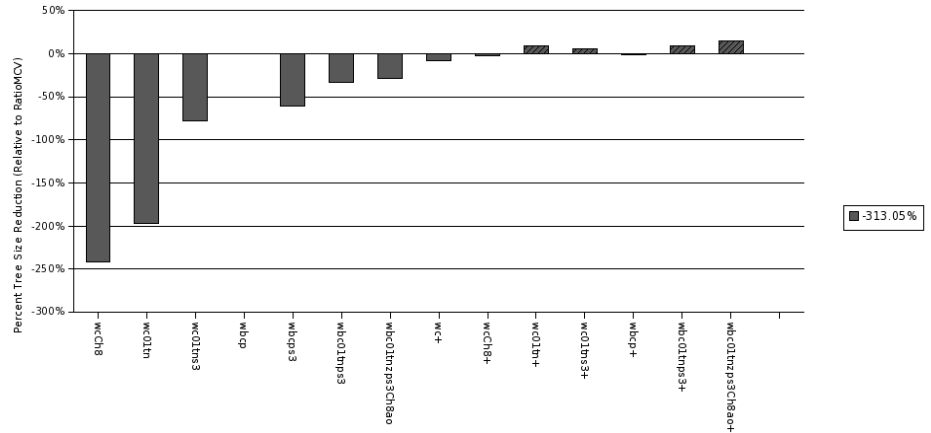


Fig. 11. Percent tree size reduction relative to the RatioMCV results in Figure 10 (higher is better). Here the poorer performance (larger tree size) for smaller feature sets is very clear. Note that since this graph does not use the baseline as a comparison point, the scale of this graph is not relevant to that of Figure 10.

of a tree using all of the available features (including the neighboring features). In other words, it is the first two levels of the 'wbc01tnzps3Ch8ao+' tree. This means that for all other feature sets the tree starts out containing paths that will never be used. The result is that the hybrid, although still greatly outperforming the baseline in terms of tree size, is actually larger than the RatioMCV version of the tree. This difference is most pronounced for the smallest feature subset (wc) and reduces to near zero for larger subsets (particularly those involving the neighboring words).

On the other hand, the runtime of the hybrid trees is unaffected by the excess tree nodes and all but one of the feature sets improved over both the original

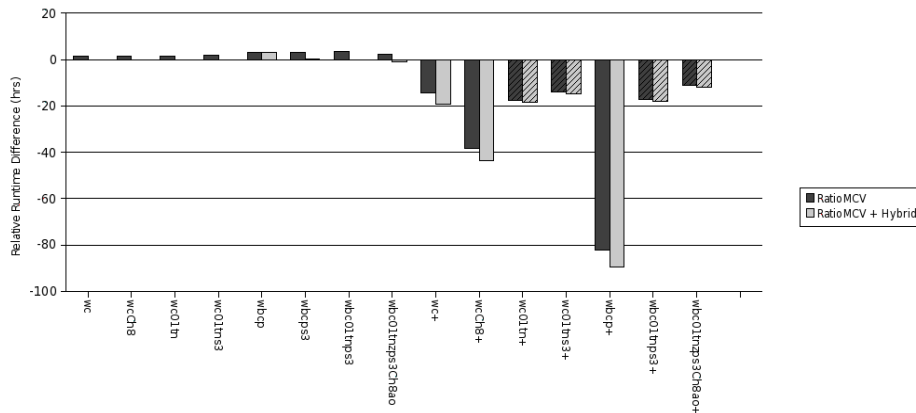


Fig. 12. Relative difference in runtime (in hours) from the baseline (lower is better). The improvement in runtime is not always large but it is consistent across (almost) all feature sets. It also consistently larger for larger feature sets since the partial static amortizes a larger percentage of the work done constructing the tree. Bars with hash marks did not finish all of the queries in the respective query logs due to memory complications.

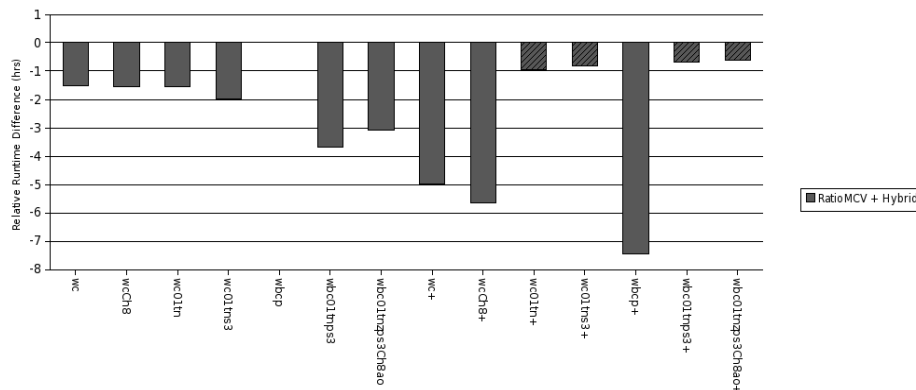


Fig. 13. Relative difference in runtime (in hours) relative to the RatioMCV results in Figure 12 (lower is better). The apparently poorer performance of those bars marked with hash marks is due to early termination caused by memory management bugs. It is reasonable to infer that the trend of increased improvement for larger feature sets would have continued without these complications.

and the RatioMCV baselines (with the one outlier performing the same as the RatioMCV baseline). Additionally, the cumulative memory usage substantially improves over the RatioMCV baseline in almost every case, in most cases at least canceling the extra memory used by the RatioMCV tree. It is also interesting to add back in the one-time cost of building the static portion of the hybrid tree. Building a static tree is generally more efficient in terms of both time and temporary data structures than dynamically building the same set of nodes. However, the cost of building the tree will still generally be greater than the improvements seen in some of the least improved feature sets (particularly those farthest right in the graph that use neighboring words). Therefore, if the tree

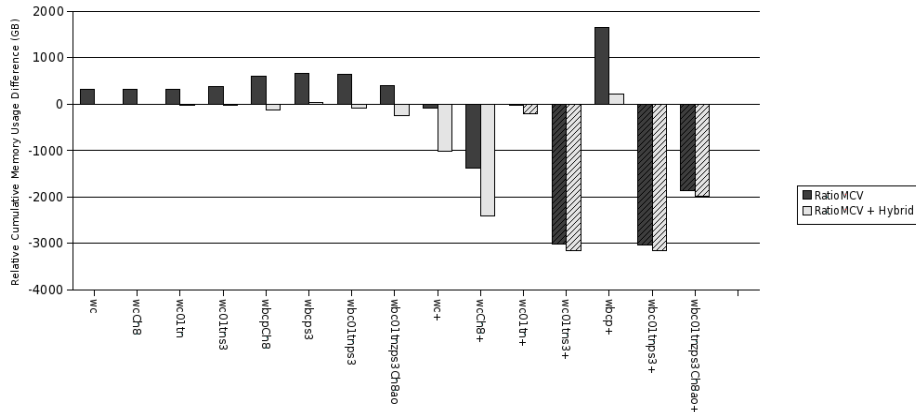


Fig. 14. Difference in cumulative memory usage (in GB) from the baseline (lower is better). Bars with hash marks did not finish all of the queries in the respective query logs due to memory complications.

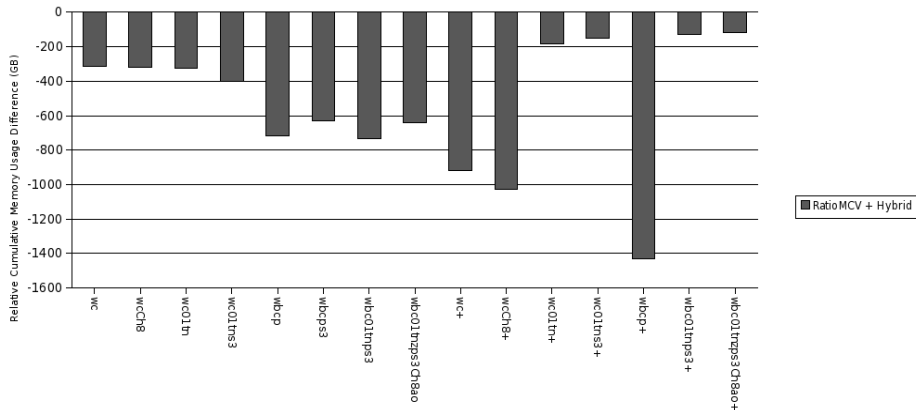


Fig. 15. Difference in cumulative memory usage (in GB) relative to the RatioMCV results in Figure 14 (lower is better). Again, the bars with hash marks presumably would have continued the trends if all queries had been completed.

is to be built only once, building a hybrid tree is not useful for all datasets. However, if that one-time cost can be amortized over many builds of the same or similar trees, both space and time savings will generally be achievable in most situations and for most problems/datasets.

Regarding the issue of using one only static tree with all features, obviously it would also be possible to generate a static two-level tree for use in a hybrid using exactly the features with which the tree would be built. For instance, one such static tree could have been pre-built for all of the feature sets experimented with here. However, this is not representative of many real-world scenarios. It is often the case that at the outset of solving a new research problem, the exact set of features that will be used is still unknown. If it is known that a dynamic ADtree will be needed (due to expected memory constraints), a hybrid tree

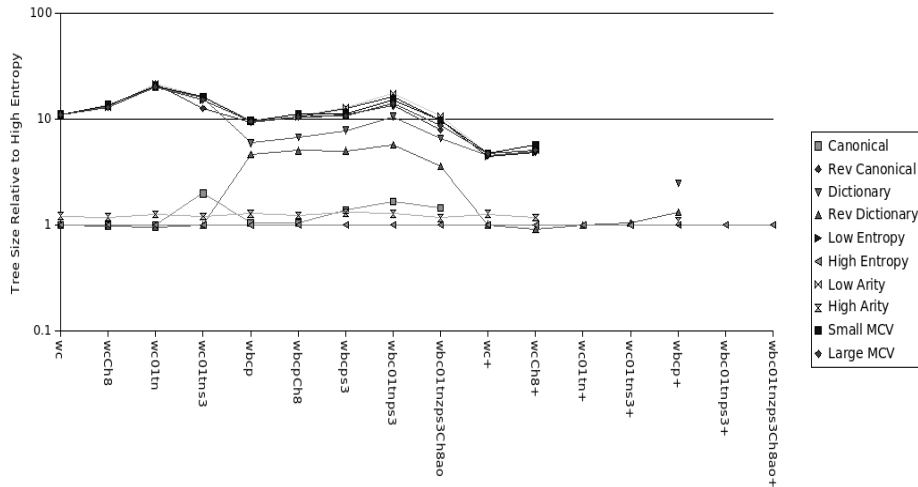


Fig. 16. Tree size of each ordering strategy relative to High Entropy feature ordering (lower is better). A value of 10 is 10 times worse than High Entropy. The ordering strategies tend to cluster into two groups, with the notable exception of Reverse Dictionary and to a lesser extent Dictionary (over only a portion of the feature sets).

will almost always be a reasonable compromise. However, the question of which features should be included in the static portion of the tree would still be an open question. Rather than generate one static tree for each possible subset of features, a single static tree with all possible features could be generated and used throughout the solution development stages. Then, once the set of features has been fixed, a new static tree could be built to exactly those specifications. The results obtained here show that using a hybrid with unnecessary features, though clearly not optimal, will still provide a significant improvement in runtime over a simple dynamic approach without exhausting available memory resources. Then, once the feature set has been determined, the excess memory use can be eliminated by simply generating a single, new two-level tree.

4.3. Ordering

Throughout the experiments presented here, it has been assumed that the variety of subsets of features of the enhanced WSJ corpus could serve as a reasonable substitute for comparing results across a wider range of datasets. The results for Complete Vary Nodes, RatioMCVs, Clump nodes, and hybrid trees indicated that this was not an unreasonable assumption to make, with interesting and illustrative differences among the various feature sets. Unfortunately, the results of trying different ordering strategies bring to the forefront a particular weakness in this assumption. As can be seen in Figures 16 and 17, the tree sizes tend to cluster into relatively clean groups. Those two groups are almost entirely defined by where in the ordering each strategy places the features ‘word’ (w) and ‘base’ (b). More specifically, the two main clusters are determined by the placement of feature ‘word’ (w) and instances where alphabetical (dictionary) and its reverse break from the clusters is determined by placement of the feature ‘base’ (b).

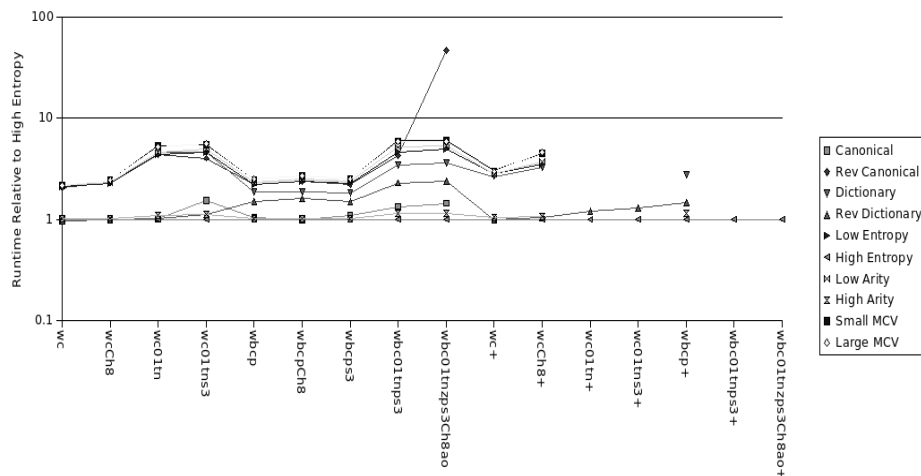


Fig. 17. Runtime in minutes of each ordering strategy relative to High Entropy feature ordering (lower is better). A value of 10 is 10 times worse than High Entropy. Again, the orderings tend towards two main groups (although Dictionary and Reverse Dictionary again show some deviation from this.)

There are likely two main reasons why ‘word’ (and to a lesser extent ‘base’) so clearly dominate these results. The first is that all the features for the central word (that is, excluding the neighboring words) are very strongly correlated to word. Therefore, if ‘word’ occurs high in the tree, there should be significantly less branching in those paths following ‘word’. Similarly, since ‘base’ is the morphological stem of ‘word’, in many ways its behavior in the tree mimics that of ‘word’ (although the relative ordering of ‘word’ and ‘base’ appears to be the primary cause of the distinct paths of the dictionary and reverse dictionary orderings). The second reason is simply that ‘word’ and ‘base’ are the two highest arity features and so they quickly subdivide the dataset such that, as was anticipated, very little further branching occurs before the leaf-list threshold is reached.

In general, the four orderings that appear to perform the best (both in terms of tree size and runtime) are high entropy, high arity, canonical, and reverse dictionary (alphabetical) in roughly that order. Figures 18-21 show the data from Figures 16-17 transposed so that trends with respect to each feature set can be better observed. Figures 18 and 20 present the ordering strategies sorted such that pairs of forward and reverse strategies are adjacent to one another. Figures 19 and 21 sort the ordering strategies according to the tree size of the feature set ‘wbc01tnzp3ch8ao’⁵ (even though the runtimes are then not exactly in order) making the difference between the two clusters of ordering schemes very explicit.

Interestingly, both strategies based on the count of MCVs faired poorly, whereas for all other pairs, one of each pair performed well. Although this is certainly due in part to relative placement of ‘word’ and ‘base’ in the orderings, it may also be likely that the absolute size of the MCV is not a useful indicator.

⁵ ‘wbc01tnzp3ch8ao’ was chosen because it is the largest set for which all the orderings finished.

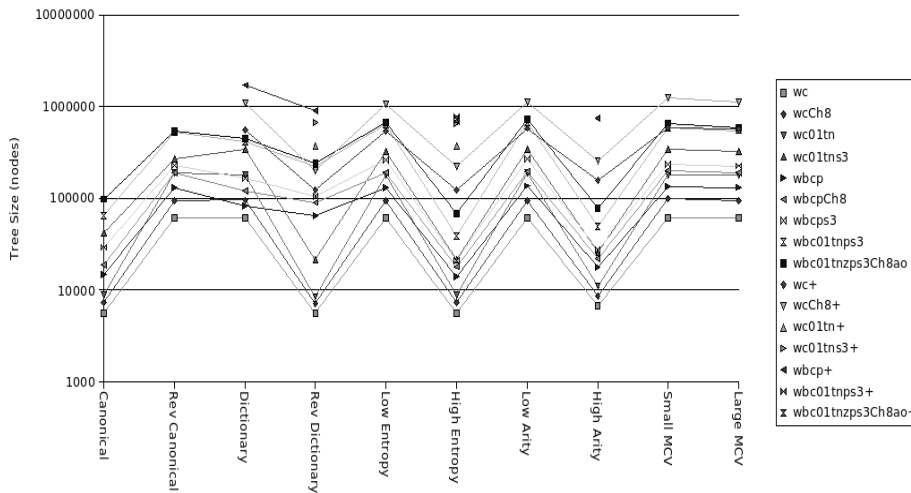


Fig. 18. Actual tree size of each ordering strategy (lower is better). The *x*-axis is sorted such that pairs of forward and reverse ordering strategies are next to each other. Note the log scale for the *y*-axis. Most of the pairs have one ordering which clearly outperforms the other. Only the orderings related to MCV both perform poorly.

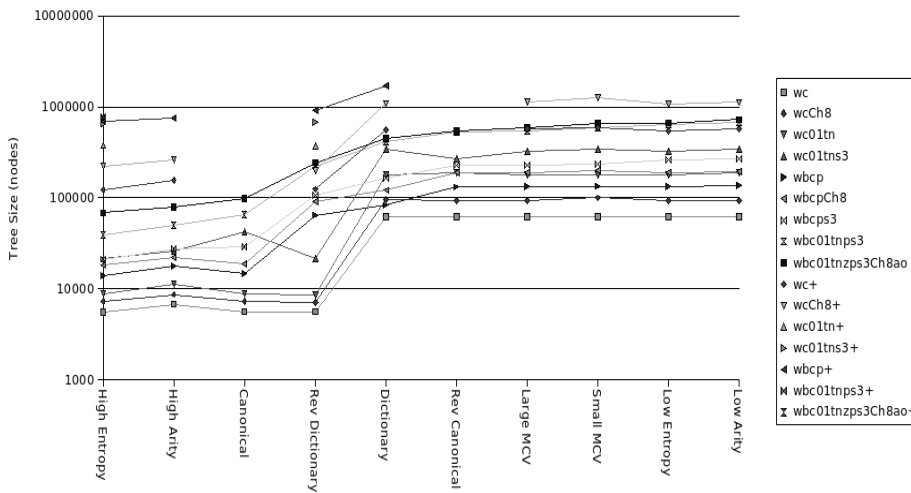


Fig. 19. Actual tree size of each ordering strategy (lower is better). The *x*-axis is sorted according to the tree size obtained by `wbc01tnzps3ch8ao`. This sorting more clearly illustrates the two basic groups of orderings. Note the log scale for the *y*-axis.

In particular, it is likely that relative MCV size (meaning, what percentage is the MCV among its siblings) could perform somewhat better than an absolute comparison. On the other hand, the absolute and relative MCV size metrics would be identical at the first level of the tree (at which point all rows in the dataset are still relevant) and at least in cases observed here, the first level ordering appears to generally dominate the overall tree size.

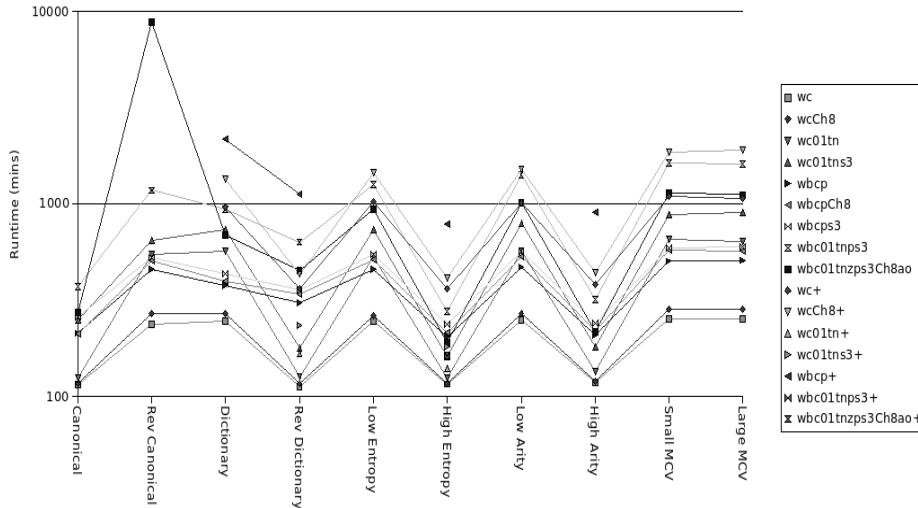


Fig. 20. Actual runtime of each ordering strategy (lower is better). The x -axis is sorted such that pairs of forward and reverse ordering strategies are next to each other. The alternating pattern of ordering pairs repeats as with tree size. Note the log scale for the y -axis.

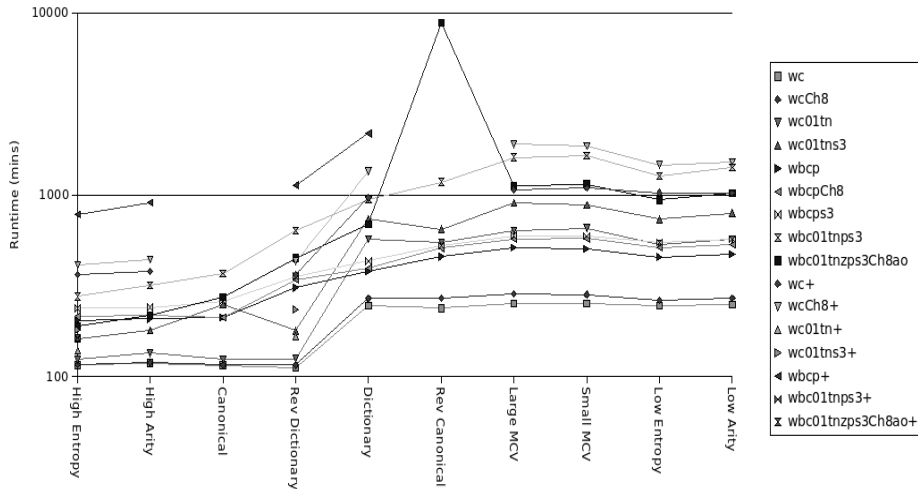


Fig. 21. Actual runtime of each ordering strategy (lower is better). The groupings remain approximately the same as with the tree size. The x -axis is sorted the same as in Figure 19. Note the log scale for the y -axis.

Another interesting trend, which can't as easily be discerned from the figures, occurs among those datapoints that have been left out. As was explained previously, the trees for each set of features are built using a log of queries as generated by the original tagger, rather than re-running the tagger in each instance. This establishes an upper limit for each feature set in terms of number of sentences, which no amount of modifying the tree can surpass. The cases where the

combination of feature set and ordering strategy did not finish all of the available sentence taggings have been removed from the graphs, resulting in some gaps in the lines. Although these cases are incomplete and therefore less reliable, there are some trends that indicate that for some of the largest feature sets, the ordering strategies that performed well actually swap with their reversed pair. Although further research would have to be performed to determine whether these trends are meaningful, the possibility still exists that highest entropy is not always the best ordering as it appears to be from the current results.

4.4. Pruning

The purpose of performing pruning within a dynamic ADtree is to try to reduce the overall tree size without significantly increasing build time. The particular pruning strategy, therefore, has to balance the reward of pruning with the risk of wasting time re-generating parts of the tree. The pruning strategy chosen in this case removes the low count ADnodes of high arity features after every 100 sentences are tagged. Figures 23 - 28 summarize the results in a format similar to that used for the hybrid tree. There are fewer feature sets represented here because several of the larger feature sets had query logs that covered less than 100 sentences and therefore no pruning occurred.

As was expected, pruning contributed significant tree size reductions to all feature sets, ranging from 5-25% on top of the already reduced RatioMCV trees. Interestingly, pruning appears to have contributed little to no additional runtime or cumulative memory usage. The largest runtime increase ('wbcp+') was 1.65 hours or approximately 2%, and the largest memory usage increase ('wcCh8+') was 8.7 GB or about 0.6%. Since pruning requires no additional data structures, no significant increase in either memory usage or runtime is expected. However, it is not unreasonable to consider the possibility that extra garbage collection could cause a slower overall runtime for some implementations. Unfortunately, for the purposes of measuring memory usage, garbage collection was artificially triggered after each sentence was tagged and this likely masked any extra garbage collection triggered by pruning. Furthermore, since this particular pruning strategy prunes the tree at most 13 times (there are slightly more than 1300 sentences), it is not anticipated that garbage collection could cause a significant increase in runtime.

Additionally, the values shown in Figures 23 - 28 represent the values obtained after the last sentence was tagged. However, since this is not necessarily immediately after the occurrence of a pruning, these numbers represent approximately "average" performance over the RatioMCV baseline. As can be seen in Figure 22, after each pruning the tree size continues to grow. This continued growth occurs due to a combination of regrowing some pruned nodes and adding previously unqueried nodes. The higher the rate of increase of the regrowth (relative to the unpruned rate of increase), the more nodes are being regenerated. Some sections appear to have more replacement growth than others but overall it appears that despite this pruning strategy's overall risk averse design, nodes are being removed that are later needed again. However, there are sufficient nodes that are never used again that the pruned tree always remains below the unpruned tree.

On the other hand, in some situations in which memory is severely limited, this pruning strategy might be considered insufficient. In that case, since the

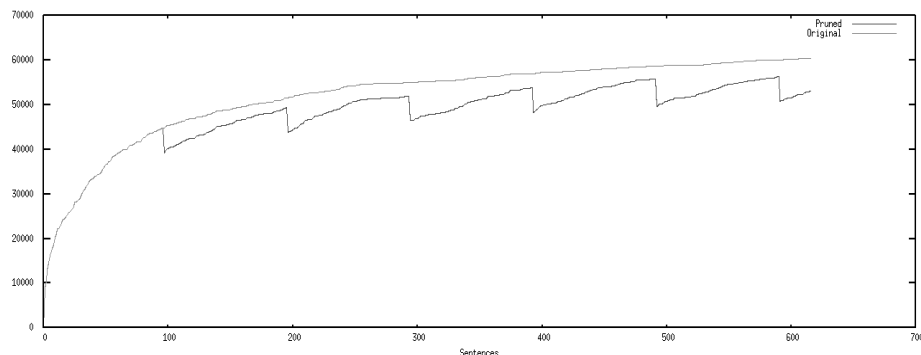


Fig. 22. Example of how tree size varies as sentences are tagged using the feature set `wbc01tnzps3Ch8ao` with and without pruning. The pruning strategy only removes a small percentage of the tree, some of which is re-built over the course of subsequent queries.

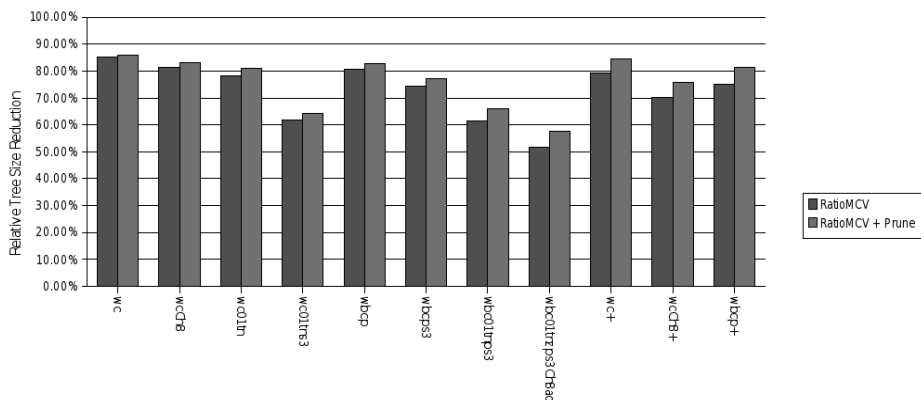


Fig. 23. Percent tree size reduction relative to the baseline when using pruning compared the results for Complete Vary Nodes and RatioMCVs (higher is better). Tree size is slightly better for all feature sets and performs the best for larger feature sets.

pruning here seems to perform without using significantly more time (given the caveat about garbage collection above), a very aggressive pruning approach (such as pruning the tree to a predetermined size) could still be usable despite the expected increase in the space/time tradeoff.

5. Conclusion

Despite the generally efficient design of the ADtree, the preceding results clearly illustrate that the modifications presented provide both generally applicable improvements as well as improvements that can be targeted at the specific characteristics of a given dataset. In particular, the addition of Complete Vary Nodes, RatioMCVs, and Clump nodes provide substantial improvements for datasets containing high arity attributes. The hybridization technique and pruning strategy presented here can be readily applied to any application of ADtrees in which

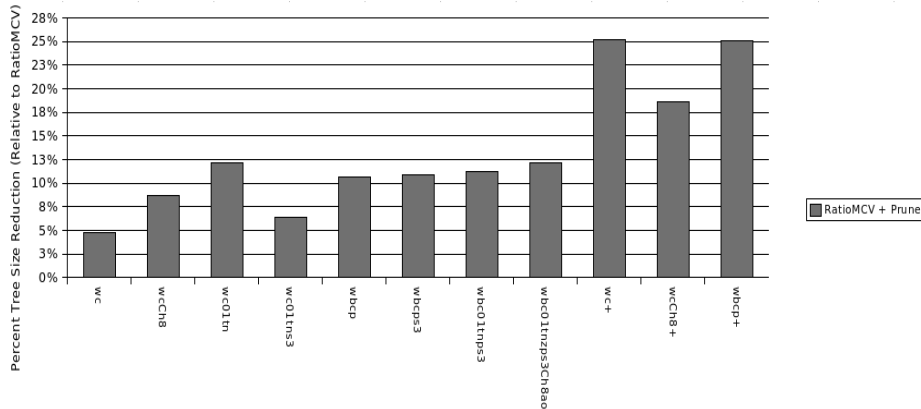


Fig. 24. Percent tree size reduction relative to the RatioMCV results in Figure 23 (higher is better). Here the trend of improved performance as feature set size increases is fairly clear.

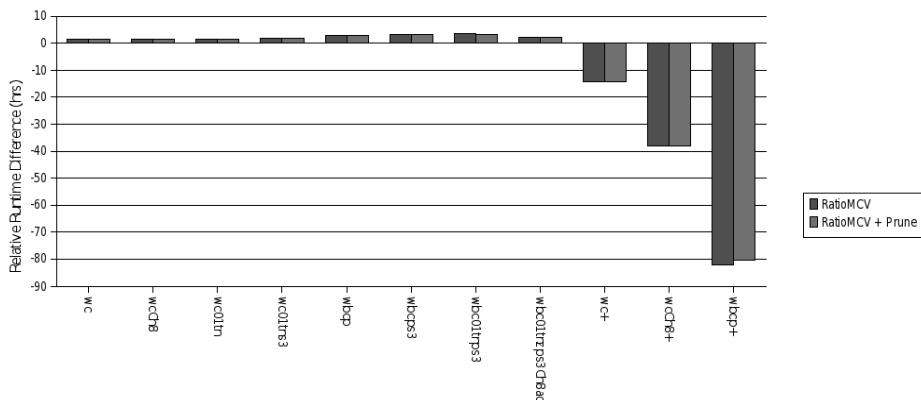


Fig. 25. Relative difference in runtime (in hours) from the baseline (lower is better). Pruning is a fairly inexpensive operation timewise and so the results are not significantly different.

using a dynamic tree is expected to be necessary. The ordering strategy based on highest entropy appears in these initial results to be generally applicable. Although this might not turn out to be true for all datasets, it certainly provides a beginning from which more advanced and/or more data specific ordering strategies can be developed.

Although there are several ways in which each of these modifications might individually be explored and developed further, in general, the results presented here would be complemented by further study involving a wider range of datasets. Then, need for more variety in the datasets is most pronounced in the results obtained from the tested ordering strategies, but the ability to duplicate the positive results on additional datasets would help to firmly establish the efficacy of each modification. Additionally, by testing on more, distinct datasets, it would be possible to develop a better sense of what data characteristics influence each technique. For instance, it could be very informative to determine

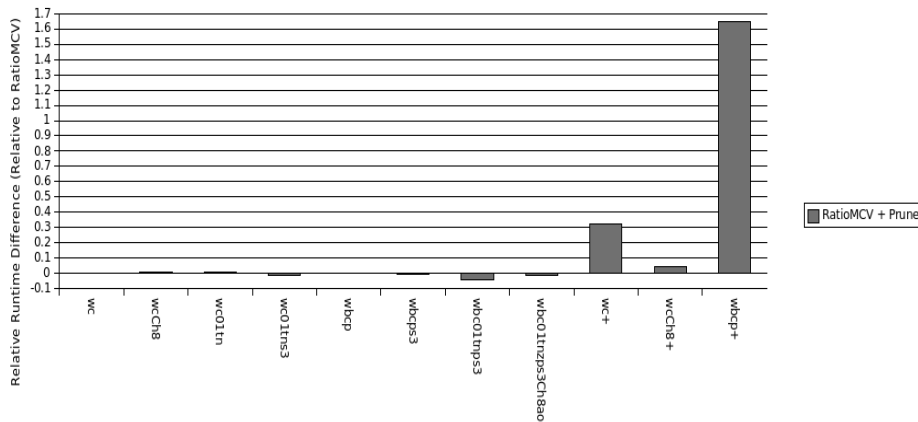


Fig. 26. Relative difference in runtime (in hours) relative to the RatioMCV results in Figure 25 (lower is better).

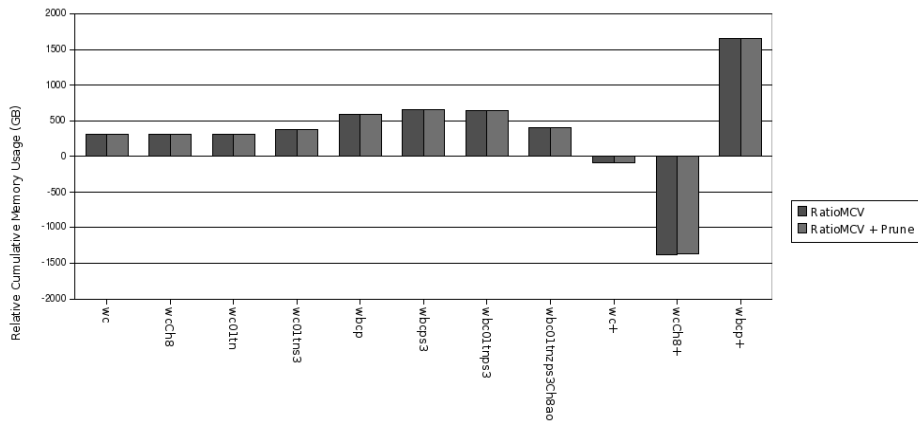


Fig. 27. Difference in cumulative memory usage (in GB) from the baseline (lower is better).

if the improvements for high arity attributes are maintained outside of the domain of natural language. It would also be interesting to explore a wider variety of pruning strategies in order to help determine how often and how much of the tree is reasonable to prune, given different memory and time requirements and/or restrictions.

References

- Agarwal, D., Agrawal, R., Khanna, R. and Kota, N. (2010), Estimating rates of rare events with multiple hierarchies through scalable log-linear models, *in* 'Proceedings of the 16th ACM SIGKDD Conference on Knowledge Discovery and Data Mining', pp. 213–222.
- Anderson, B. and Moore, A. (1998), Adtrees for fast counting and for fast learning of association rules, *in* 'Proceedings of the Fourth International Conference on Knowledge Discovery in Data Mining', AAAI Press, pp. 134–138.

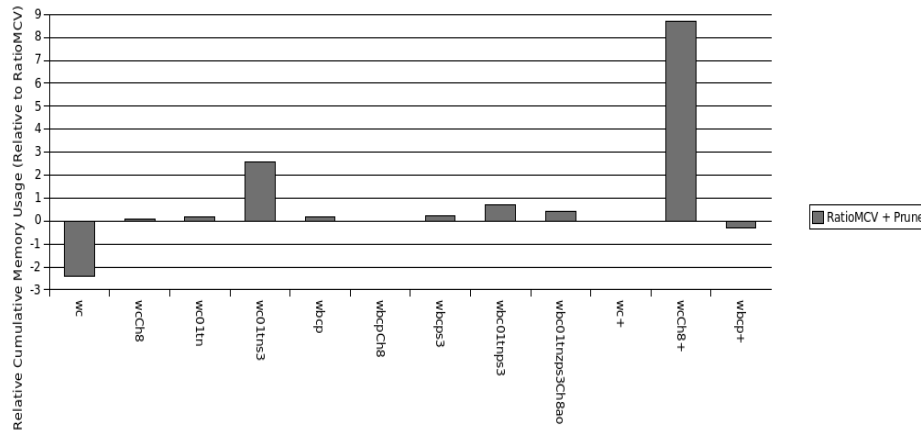


Fig. 28. Difference in cumulative memory usage (in GB) relative to the RatioMCV results in Figure 27 (lower is better).

- Bentley, J. (1975), ‘Multidimensional binary search trees used for associative searching’, *Communications of the Association for Computing Machinery* **18**(9), 509–517.
- Chen, H., Liu, J., Furuse, K., Yu, J. X. and Ohbo, N. (2011), ‘Indexing expensive functions for efficient multi-dimensional similarity search’, *Knowledge and Information Systems* **27**(2), 165–192.
- Fuchs, H., Kedem, Z. and Naylor, B. (1980), On visible surface generation by *a priori* tree structures, in ‘International Conference on Computer Graphics and Interactive Techniques’, pp. 124–133.
- Gaede, V. and Gunther, O. (1998), ‘Multidimensional access methods’, *Association for Computing Machinery Computing Survey* **30**(2), 170–231.
- Huang, Z., Sun, S. and Wang, W. (2010), ‘Efficient mining of skyline objects in subspaces over data streams’, *Knowledge and Information Systems* **22**(2), 159–183.
- Kim, Y., Chung, C.-W., Lee, S.-L. and Kim, D.-H. (2011), ‘Distance approximation techniques to reduce the dimensionality for multimedia databases’, *Knowledge and Information Systems* **28**(1), 227–248.
- Komarek, P. and Moore, A. (2000), A dynamic adaptation of ad-trees for efficient machine learning on large data sets, in ‘Proceedings of the International Conference on Machine Learning (ICML)’, pp. 495–502.
- Koufakou, A., Secretan, J. and Georgiopoulos, M. (2011), ‘Non-derivable itemsets for fast outlier detection in large high-dimensional categorical data’, *Knowledge and Information Systems* **29**(3), 697–725.
- Moore, A. and Lee, M. S. (1998), ‘Cached sufficient statistics for efficient machine learning with large datasets’, *Journal of Artificial Intelligence Research* **8**, 67–91.
- Roure, J. and Moore, A. (2006), Sequential update of adtrees, in ‘Proceedings of the 23rd International Conference on Machine Learning’, pp. 769–776.
- Rymon, R. (1993), An se-tree based characterization of the induction problem, in ‘International Conference on Machine Learning’, pp. 268–275.
- Toutanova, K. and Manning, C. (2000), Enriching the knowledge sources used in a maximum entropy part-of-speech tagger, in ‘Proceedings of the Joint Conference on Empirical Methods in Natural Language Processing and Very Large Corpora’, pp. 63–70.
- University of Pennsylvania Linguistic Data Consortium* (n.d.).
URL: <http://www ldc.upenn.edu/>
- Van Dam, R., Langkilde-Geary, I. and Ventura, D. (2008), Adapting adtrees for high arity features, in ‘Proceedings of the Association for the Advancement of Artificial Intelligence’, pp. 708–713.
- Van Dam, R. and Ventura, D. (2007), Adtrees for sequential data and n-gram counting, in ‘Proceedings of the IEEE Conference on Systems, Man, and Cybernetics’, pp. 492–497.
- Yu, H.-F., Hsieh, C.-J., Chang, K.-W. and Lin, C.-J. (2010), Large linear classification when



data cannot fit in memory, in 'Proceedings of the 16th ACM SIGKDD Conference on Knowledge Discovery and Data Mining', pp. 833–842.

Zhang, M. and Alhaji, R. (2010), 'Effectiveness of naq-tree as index structure for similarity search in high-dimensional metric space', *Knowledge and Information Systems* **22**(1), 159–183.

Robert Van Dam received a B.A. degree in Mathematics and a B.S. degree in Computer Science from the University of Rochester in 2005 and a M.S. degree from Brigham Young University in 2008. His research interests include natural language processing, fraud detection, large datasets, data structure optimization, and sentiment analysis.

Irene Langkilde-Geary received her PhD from the University of Southern California. She has worked as an Assistant Professor at Brigham Young University in Provo, Utah, as a consultant for Southwest Research Institute in Layton, Utah,



and as a Research Fellow at the University of Brighton, in the U.K. She currently conducts independent research in machine learning and natural language processing while mothering five young children.

Dan Ventura is a Professor of Computer Science at Brigham Young University, where he has been since 2001. His research interests include artificial intelligence, machine learning, neural networks, computational creativity and



the application of intelligent systems to real-world domains, including problems in medicine, chemistry and business.