

A Dual-Tree Algorithm for Fast k -means Clustering With Large k

Ryan R. Curtin
ryan@ratml.org

Abstract

k -means is a widely used clustering algorithm, but for k clusters and a dataset size of N , each iteration of Lloyd’s algorithm costs $O(kN)$ time. This is problematic because increasingly, applications of k -means involve both large N and large k , and there are no accelerated variants that handle this situation. To this end, we propose a dual-tree algorithm that gives the *exact* same results as standard k -means; when using cover trees, we bound the single-iteration runtime of the algorithm as $O(N + k \log k)$, under some assumptions. To our knowledge these are the first sub- $O(kN)$ bounds for exact Lloyd iterations. The algorithm performs competitively in practice, especially for large N and k , in low-dimensional settings. The algorithm is generic, so any type of tree (i.e. kd -tree, octree, etc.) may be used.

1 Introduction

Of all the clustering algorithms in use today, among the simplest and most popular is the venerated k -means clustering algorithm, usually implemented via Lloyd’s algorithm: given a dataset S , repeat the following two steps (a ‘Lloyd iteration’) until the centroids of each of the k clusters converge:

1. Assign points to the cluster with nearest centroid.
2. Recalculate the centroids for each cluster using the assignments of each point.

Clearly, a simple implementation of this algorithm will take $O(kN)$ time where $N = |S|$. But the number of iterations is not bounded unless the practitioner manually sets a maximum, and k -means is not guaranteed to converge to the global best clustering. Despite these shortcomings, in practice k -means tends to quickly converge to reasonable solutions. Even so, there is no shortage of techniques for improving the clusters k -means converges to: refinement of initial centroids [7] and weighted sampling of initial centroids [2] are just two of many popular existing strategies.

There are also a number of methods that use the triangle inequality to accelerate the runtime of a single iteration of k -means, generally focused only on the large N case. Algorithms of this sort include the work of

Pellegrino and Moore [31], Elkan [19], Hamerly [23], and Ding et al. [17].

In this paper, we describe a dual-tree k -means algorithm for the large k and large N case that outperforms current competing algorithms in that setting. This algorithm also has bounded single-iteration runtime under some assumptions (see Section 6). The algorithm, which is our main contribution, has several appealing aspects:

- **Empirical efficiency.** In the large k and large N setting for which this algorithm is designed, it outperforms all other alternatives, and scales better to larger datasets. The algorithm is especially efficient for datasets of low (intrinsic) dimensionality.
- **Runtime guarantees.** Using adaptive runtime analysis techniques, we bound the single-iteration runtime of our algorithm with respect to the intrinsic dimensionality of the centroids and data, when cover trees are used. This gives theoretical support for the use of our algorithm. In addition, the bound is dependent on the intrinsic dimensionality, *not* the extrinsic dimensionality.
- **Generalizability.** We develop our algorithm using a tree-independent dual-tree algorithm abstraction [14]; this means that our algorithm may be used with *any* type of valid tree. This includes not just kd -trees but also metric trees, cone trees, cover trees, octrees, and others.
- **Separation of concerns.** The abstraction we use to develop our algorithm allows us to focus on and formalize each of the pruning rules individually (Section 4). This aids understanding of the algorithm and eases insertion of future improvements and better pruning rules.

Section 2 shows the relevance of the large k case; then, in Section 3, we show that we can build a tree on the k clusters, and then a dual-tree algorithm [14] can be used to efficiently perform an exact single iteration of k -means clustering. Section 4 details the four pruning strategies used in our algorithm, and Section 5 introduces the algorithm itself. Sections 6

Algorithm	Setup	Worst-case	Memory
naive	n/a	$O(kN)$	$O(k + N)$
blacklist	$O(N \log N)$	$O(kN)$	$O(k \log N + N)$
elkan	n/a	$O(k^2 + kN)$	$O(k^2 + kN)$
hamerly	n/a	$O(k^2 + kN)$	$O(k + N)$
yinyang	$O(k^2 + kN)$	$O(kN)$	$O(kN)$
dualtree	$O(N \log N)$	$O(k \log k + N)$ ¹	$O(k + N)$

Table 1: Runtime and memory bounds for k -means algorithms.

and 7 show the theoretical and empirical results for the algorithm, and finally Section 8 concludes the paper and paints directions for future improvements.

2 Scaling k -means

Although the original publications on k -means only applied the algorithm to a maximum dataset size of 760 points, the half-century of relentless progress since then has seen dataset sizes scale into billions. Due to its simplicity, though, k -means has remained relevant, and is still applied in many large-scale applications.

In cases where N is large but k remains small, a good choice is a sampling algorithm, which will return an approximate clustering. One sampling technique, coresets, can produce good clusterings for N in the millions using a few thousand points [20]. However, for large k , the number of samples required to produce good clusterings can become prohibitive.

For large k , then, we turn to an alternative approach: accelerating exact Lloyd iterations. Existing techniques include the brute-force implementation, the *blacklist* algorithm [31], Elkan’s algorithm [19], and Hamerly’s algorithm [23], as well as the recent Yinyang k -means algorithm [17]. The blacklist algorithm builds a kd -tree on the dataset and, while the tree is traversed, blacklists individual clusters that cannot be the closest cluster (the *owner*) of any descendant points of a node. Elkan’s algorithm maintains an upper bound and a lower bound on the distance between each point and centroid; Hamerly’s algorithm is a memory-efficient sim-

plification of this technique. The Yinyang algorithm organizes the centroids into groups of about 10 (depending on parameters) using 5 iterations of k -means on the centroids followed by a single iteration of standard k -means on the points. Once groups are built, the Yinyang algorithm attempts to prune groups of centroids using rules similar to Elkan and Hamerly’s algorithms.

Of these algorithms, only Yinyang k -means considers centroids in groups at all, but it does not consider points in groups. On the other hand, the blacklist algorithm is the only one that builds a tree on the points and is able to assign multiple points to a single cluster at once. So, though each algorithm has its own useful region, none of the four we have considered here are well-suited to the case of large N **and** large k (see Table 2 for a rough sketch).

Table 1 shows setup costs, worst-case per-iteration runtimes, and memory usage of each of these algorithms as well as the proposed dual-tree algorithm¹. The expected runtime of the blacklist algorithm is, under some assumptions, $O(k + k \log N + N)$ per iteration. The expected runtime of Hamerly’s and Elkan’s algorithm is $O(k^2 + \alpha N)$ time, where α is the expected number of clusters visited by each point (in both Elkan and Hamerly’s results, α seems to be small).

However, none of these algorithms are specifically tailored to the large k case, and the large k case is common. Pelleg and Moore [31] report several hundred clusters in a subset of 800k objects from the SDSS dataset. Clusterings for n -body simulations on astronomical data often involve several thousand clusters [25]. Csurka et al. [10] extract vocabularies from image sets using k -means with $k \sim 1000$. Coates et al. [9] show that k -means can work surprisingly well for unsupervised feature learning for images, using k as large as 4000 on 50000 images. Also, in text mining, datasets may have up to 18000 unique labels [4]. Can and Ozkaran [8] suggest that the number of clusters in text data is directly related to the size of the vocabulary, suggesting $k \sim mN/t$ where m is the vocabulary size, n is the number of documents, and t is the number

	smaller k	larger k
smaller N	brute-force elkan	X
larger N	sampling hamerly pelleg-moore yinyang	dualtree

Table 2: Good candidate algorithm choices for different settings. Elkan’s algorithm is not suited to large N because of its memory requirements.

¹The dual-tree algorithm worst-case runtime bound also depends on some assumptions on dataset-dependent constants. This is detailed further in Section 6.

of nonzero entries in the term matrix. Thus, it is important to have an algorithm with favorable scaling properties for both large k and N .

3 Tree-based algorithms

The blacklist algorithm is an example of a *single-tree algorithm*: one tree (the *reference tree*) is built on the dataset, and then that tree is traversed. This approach is applicable to a surprising variety of other problems, too [5, 30, 16]. It is only natural to build a tree on the data points: tree building is a one-time $O(N \log N)$ cost and for large N , this cost is generally negligible compared to the cost of clustering.

The speedup of the blacklist algorithm comes from the hierarchical nature of trees: during the algorithm, we may rule out a cluster centroid for *many points at once*. The same reason is responsible for the impressive speedups obtained for other single-tree algorithms, such as nearest neighbor search [5, 28]. But for nearest neighbor search, the nearest neighbor is often required not just for a query point but instead a *query set*. This motivated the development of *dual-tree algorithms*, which also build a tree on the query set (the *query tree*) in order to share work across query points.

This general approach is applicable to k -means with large k : we may build a *query tree* on the data points, as well as a *reference tree* on the k cluster centroids. We may then use a dual-tree algorithm to rule out *many* centroids for *many* points at once.

A recent result generalizes the class of dual-tree algorithms, simplifying their expression and development [14]. Any dual-tree algorithm can be decomposed into three parts: a type of space tree, a pruning dual-tree traversal, and a point-to-point `BaseCase()` function and node-to-node `Score()` function that determines when pruning is possible². So, given any type of tree (such as a kd -tree) and a pruning dual-tree traversal (such as a dual depth-first traversal [11]), the algorithm will call the `Score()` function with a query tree node \mathcal{N}_q and a reference tree node \mathcal{N}_r . The `Score()` function will determine if the node combination may be pruned, in which case there will be no further recursion down those subtrees. If the node combination is not pruned, then the `BaseCase()` function will be called with each combination of query points held in \mathcal{N}_q and reference points held in \mathcal{N}_r . Because there are already numerous types of trees and traversals [14, 11], to create a dual-tree k -means algorithm that can perform the assignments step we only need to develop two functions:

- `Score(\mathcal{N}_q , \mathcal{N}_r)`, which determines if any descendant centroids in \mathcal{N}_r could own any descendant

points of \mathcal{N}_q , and prunes the combination if not;

- `BaseCase(p_q , c_r)`, which calculates if c_r is the closest centroid to p_q .

The two types of trees we will explicitly consider in this paper are the kd -tree and the cover tree [6], but it should be remembered that the algorithm as provided is sufficiently general to work with any other type of tree. Therefore, we standardize notation for trees; see Table 3 and [14]. It is important to note that the set \mathcal{P}_i is not equivalent to \mathcal{D}_i^p . Lastly, we say that a centroid c owns a point p if c is the closest centroid to p .

\mathcal{T}_i	a tree
\mathcal{N}_i	a node
\mathcal{C}_i	the set of child nodes of \mathcal{N}_i
\mathcal{P}_i	the set of points held in \mathcal{N}_i
\mathcal{D}_i^p	the set of descendant points of \mathcal{N}_i

Table 3: Notation for trees, from [14].

4 Pruning strategies

All of the existing accelerated k -means algorithms operate by avoiding unnecessary work via the use of pruning strategies. Thus, we will pursue four pruning strategies, each based on or related to earlier work [31, 19, 23].

These pruning strategies are meant to be used during the dual-tree traversal, for which we have built a query tree \mathcal{T}_q on the points and a reference tree \mathcal{T}_r on the centroids. Therefore, these pruning strategies consider not just combinations of single points and centroid p_q and c_i , but the combination of sets of points and sets of centroids, represented by a query tree node \mathcal{N}_q and a centroid tree node \mathcal{N}_r . This allows us to prune many centroids for many points simultaneously.

Strategy one. When visiting a particular combination $(\mathcal{N}_q, \mathcal{N}_r)$ (with \mathcal{N}_q holding points and \mathcal{N}_r holding centroids), the combination should be pruned if every descendant centroid in \mathcal{N}_r can be shown to own none of the points in \mathcal{N}_q . If we have cached an upper bound $\text{ub}(\mathcal{N}_q)$ on the distance between any descendant point of \mathcal{N}_q and its nearest cluster centroid that satisfies

$$(4.1) \quad \text{ub}(\mathcal{N}_q) \geq \max_{p_q \in \mathcal{D}_q^p} d(p_q, c_q)$$

where c_q is the cluster centroid nearest to point p_q , then the node \mathcal{N}_r can contain no centroids that own any descendant points of \mathcal{N}_q if

$$(4.2) \quad d_{\min}(\mathcal{N}_q, \mathcal{N}_r) > \text{ub}(\mathcal{N}_q).$$

This relation bears similarity to the pruning rules for nearest neighbor search [14] and max-kernel search [15]. Figure 1a shows a situation where \mathcal{N}_r can be pruned; in this case, ball-shaped tree nodes are used, and the upper bound $\text{ub}(\mathcal{N}_q)$ is set to $d_{\max}(\mathcal{N}_q, \mathcal{N}_{r2})$.

²Precise definitions can be found in the original paper [14].

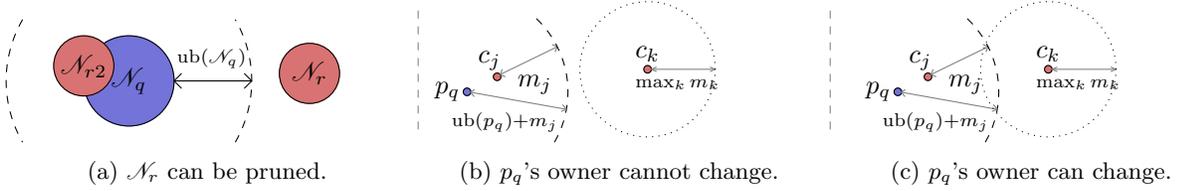


Figure 1: Different pruning situations.

Strategy two. The recursion down a particular branch of the query tree should terminate early if we can determine that only one cluster can possibly own all of the descendant points of that branch. This is related to the first strategy. If we have been caching the number of pruned centroids (call this $\text{pruned}(\mathcal{N}_q)$), as well as the identity of any arbitrary non-pruned centroid (call this $\text{closest}(\mathcal{N}_q)$), then if $\text{pruned}(\mathcal{N}_q) = k - 1$, we may conclude that the centroid $\text{closest}(\mathcal{N}_q)$ is the owner of all descendant points of \mathcal{N}_q , and there is no need for further recursion in \mathcal{N}_q .

Strategy three. The traversal should not visit nodes whose owner could not have possibly changed between iterations; that is, the tree should be coalesced to include only nodes whose owners may have changed. There are two easy ways to use the triangle inequality to show that the owner of a point cannot change between iterations. Figures 1b and 1c show the first: we have a point p_q with owner c_j and second-closest centroid c_k . Between iterations, each centroid will move when it is recalculated; define the distance that centroid c_i has moved as m_i . Then we bound the distances for the next iteration: $d(p_q, c_j) + m_j$ is an upper bound on the distance from p_q to its owner next iteration, and $d(p_q, c_k) - \max_i m_i$ is a lower bound on the distance from p_q to its second closest centroid next iteration. We may use these bounds to conclude that if

$$(4.3) \quad d(p_q, c_j) + m_j < d(p_q, c_k) - \max_i m_i,$$

then the owner of p_q next iteration must be c_j . Generalizing from individual points p_q to tree nodes \mathcal{N}_q is easy. This pruning strategy can only be used when all descendant points of \mathcal{N}_q are owned by a single centroid, and in order to perform the prune, we need to establish a lower bound on the distance between any descendant point of the node \mathcal{N}_q and the second closest centroid. Call this bound $\text{lb}(\mathcal{N}_q)$. Remember that $\text{ub}(\mathcal{N}_q)$ provides an upper bound on the distance between any descendant point of \mathcal{N}_q and its nearest centroid. Then, if all descendant points of \mathcal{N}_q are owned by some cluster c_j in one iteration, and

$$(4.4) \quad \text{ub}(\mathcal{N}_q) + m_j < \text{lb}(\mathcal{N}_q) - \max_i m_i,$$

then \mathcal{N}_q is owned by cluster c_j in the next iteration. Implementationally, it is convenient to have $\text{lb}(\mathcal{N}_q)$ store a lower bound on the distance between any descendant

point of \mathcal{N}_q and the nearest pruned centroid. Then, if \mathcal{N}_r is entirely owned by one cluster, all other centroids are pruned, and $\text{lb}(\mathcal{N}_q)$ holds the necessary lower bound for pruning according to the rule above.

The second way to use the triangle inequality to show that an owner cannot change depends on the distances between centroids. Suppose that p_q is owned by c_j at the current iteration; then, if

$$(4.5) \quad d(p_q, c_j) - m_j < 2 \left(\min_{c_i \in C, c_i \neq c_j} d(c_i, c_j) \right)$$

then c_j will own p_q next iteration [19]. We may adapt this rule to tree nodes \mathcal{N}_q in the same way as the previous rule; if \mathcal{N}_q is owned by cluster c_j during this iteration and

$$(4.6) \quad \text{ub}(\mathcal{N}_q) + m_j < 2 \left(\min_{c_i \in C, c_i \neq c_j} d(c_i, c_j) \right)$$

then \mathcal{N}_q is owned by cluster c_j in the next iteration. Note that the above rules do work with individual points p_q instead of nodes \mathcal{N}_q if we have a valid upper bound $\text{ub}(p_q)$ and a valid lower bound $\text{lb}(p_q)$. Any nodes or points that satisfy the above conditions do not need to be visited during the next iteration, and can be removed from the tree for that next iteration.

Strategy four. The traversal should use bounding information from previous iterations; for instance, $\text{ub}(\mathcal{N}_q)$ should not be reset to ∞ at the beginning of each iteration. Between iterations, we may update $\text{ub}(\mathcal{N}_q)$, $\text{ub}(p_q)$, $\text{lb}(\mathcal{N}_q)$, and $\text{lb}(p_q)$ according to the following rules:

$$\begin{aligned} \text{ub}(\mathcal{N}_q) &\leftarrow \begin{cases} \text{ub}(\mathcal{N}_q) + m_j & \text{if } \mathcal{N}_q \text{ is} \\ & \text{owned by a single cluster } c_j \\ \text{ub}(\mathcal{N}_q) + \max_i m_i & \text{if } \mathcal{N}_q \text{ is} \\ & \text{not owned by a single cluster,} \end{cases} \\ \text{ub}(p_q) &\leftarrow \text{ub}(p_q) + m_j, \\ \text{lb}(\mathcal{N}_q) &\leftarrow \text{lb}(\mathcal{N}_q) - \max_i m_i, \\ \text{lb}(p_q) &\leftarrow \text{lb}(p_q) - \max_i m_i. \end{aligned}$$

Special handling is required when descendant points of \mathcal{N}_q are not owned by a single centroid (Equation 4.7). It is also true that for a child node \mathcal{N}_c of \mathcal{N}_q , $\text{ub}(\mathcal{N}_q)$ is a valid upper bound for \mathcal{N}_c and $\text{lb}(\mathcal{N}_q)$ is a valid lower bound for \mathcal{N}_c : that is, the upper and lower bounds may be taken from a parent, and they are still valid.

5 The dual-tree k -means algorithm

These four pruning strategies lead to a high-level k -means algorithm, described in Algorithm 1. During the course of this algorithm, to implement each of our pruning strategies, we will need to maintain the following quantities:

- $\text{ub}(\mathcal{N}_q)$: an upper bound on the distance between any descendant point of a node \mathcal{N}_q and the nearest centroid to that point.
- $\text{lb}(\mathcal{N}_q)$: a lower bound on the distance between any descendant point of a node \mathcal{N}_q and the nearest pruned centroid.
- $\text{pruned}(\mathcal{N}_q)$: the number of centroids pruned during traversal for \mathcal{N}_q .
- $\text{closest}(\mathcal{N}_q)$: if $\text{pruned}(\mathcal{N}_q) = k - 1$, this holds the owner of all descendant points of \mathcal{N}_q .
- $\text{canchange}(\mathcal{N}_q)$: whether or not \mathcal{N}_q can change owners next iteration.
- $\text{ub}(p_q)$: an upper bound on the distance between point p_q and its nearest centroid.
- $\text{lb}(p_q)$: a lower bound on the distance between point p_q and its second nearest centroid.
- $\text{closest}(p_q)$: the closest centroid to p_q (this is also the owner of p_q).
- $\text{canchange}(p_q)$: whether or not p_q can change owners next iteration.

At the start of the algorithm, each upper bound is initialized to ∞ , each lower bound is initialized to ∞ , $\text{pruned}(\cdot)$ is initialized to 0 for each node, and $\text{closest}(\cdot)$ is initialized to an invalid centroid for each node and point. $\text{canchange}(\cdot)$ is set to `true` for each node and point. Thus line 6 does nothing on the first iteration.

First, consider the dual-tree algorithm called on line 9. As detailed earlier, we can describe a dual-tree algorithm as a combination of tree type, traversal, and point-to-point `BaseCase()` and node-to-node `Score()` functions. Thus, we need only present `BaseCase()` (Algorithm 2) and `Score()` (Algorithm 3)³.

The `BaseCase()` function is simple: given a point p_q and a centroid c_r , the distance $d(p_q, c_r)$ is calculated; $\text{ub}(p_q)$, $\text{lb}(p_q)$, and $\text{closest}(p_q)$ are updated if needed.

`Score()` is more complex. The first stanza (lines 4–6) takes the values of $\text{pruned}(\cdot)$ and $\text{lb}(\cdot)$ from the parent node of \mathcal{N}_q ; this is necessary to prevent $\text{pruned}(\cdot)$ from undercounting. Next, we prune if the owner of \mathcal{N}_q is already known (line 7). If the minimum distance

³In these algorithms, we assume that any point present in a node \mathcal{N}_i will also be present in at least one child $\mathcal{N}_c \in \mathcal{C}_i$. It is possible to fully generalize to any tree type, but the exposition is significantly more complex, and our assumption covers most standard tree types anyway.

Algorithm 1 High-level outline of dual-tree k -means.

- 1: **Input:** dataset $S \in \mathcal{R}^{N \times d}$, initial centroids $C \in \mathcal{R}^{k \times d}$.
 - 2: **Output:** converged centroids C .
 - 3: $\mathcal{T} \leftarrow$ a tree built on S
 - 4: **while** centroids C not converged **do**
 - 5: {Remove nodes in the tree if possible.}
 - 6: $\mathcal{T} \leftarrow \text{CoalesceNodes}(\mathcal{T})$
 - 7: $\mathcal{T}_c \leftarrow$ a tree built on C
 - 8: {Call dual-tree algorithm.}
 - 9: Perform a dual-tree recursion with \mathcal{T} , \mathcal{T}_c , `BaseCase()`, and `Score()`.
 - 10: {Restore the tree to its non-coalesced form.}
 - 11: $\mathcal{T} \leftarrow \text{DecoalesceNodes}(\mathcal{T})$
 - 12: {Update centroids and bounding information.}
 - 13: $C \leftarrow \text{UpdateCentroids}(\mathcal{T})$
 - 14: $\mathcal{T} \leftarrow \text{UpdateTree}(\mathcal{T})$
 - 15: **end while**
 - 16: **return** C
-

between any descendant point of \mathcal{N}_q and any descendant centroid of \mathcal{N}_r is greater than $\text{ub}(\mathcal{N}_q)$, then we may prune the combination (line 16). In that case we may also improve the lower bound (line 14). Note the special handling in line 15: our definition of tree allows points to be held in more than one node; thus, we must avoid double-counting clusters that we prune.⁴ If the node combination cannot be pruned in this way, an attempt is made to update the upper bound (lines 17–20). Instead of using $d_{\max}(\mathcal{N}_q, \mathcal{N}_r)$, we may use a tighter upper bound: select any descendant centroid c from \mathcal{N}_r and use $d_{\max}(\mathcal{N}_q, c)$. This still provides a valid upper bound, and in practice is generally smaller than $d_{\max}(\mathcal{N}_q, \mathcal{N}_r)$. We simply set $\text{closest}(\mathcal{N}_q)$ to c (line 20); $\text{closest}(\mathcal{N}_q)$ only holds the owner of \mathcal{N}_q if all centroids except one are pruned—in which case the owner *must* be c .

Thus, at the end of the dual-tree algorithm, we know the owner of every node (if it exists) via $\text{closest}(\cdot)$ and $\text{pruned}(\cdot)$, and we know the owner of every point via $\text{closest}(\cdot)$. A simple algorithm to do this is given here as Algorithm 4 (`UpdateCentroids()`); it is a depth-first recursion through the tree that terminates a branch when a node is owned by a single cluster.

Next is updating the bounds in the tree and determining if nodes and points can change owners next iteration; this work is encapsulated in the `UpdateTree()` algorithm, which is an implementation of strategies 3 and 4 (see the appendix for details). Once `UpdateTree()`

⁴For trees like the kd -tree and the metric tree, which do not hold points in more than one node, no special handling is required: we will never prune a cluster twice for a given query node \mathcal{N}_q .

Algorithm 2 BaseCase() for dual-tree k -means.

```
1: Input: query point  $p_q$ , reference centroid  $c_r$ 
2: Output: distance between  $p_q$  and  $c_r$ 
3: if  $d(p_q, c_r) < \text{ub}(p_q)$  then
4:    $\text{lb}(p_q) \leftarrow \text{ub}(p_q)$ 
5:    $\text{ub}(p_q) \leftarrow d(p_q, c_r)$ 
6:    $\text{closest}(p_q) \leftarrow c_r$ 
7: else if  $d(p_q, c_r) < \text{lb}(p_q)$  then
8:    $\text{lb}(p_q) \leftarrow d(p_q, c_r)$ 
9: end if
10: return  $d(p_q, c_r)$ 
```

Algorithm 3 Score() for dual-tree k -means.

```
1: Input: query node  $\mathcal{N}_q$ , reference node  $\mathcal{N}_r$ 
2: Output: score for node combination  $(\mathcal{N}_q, \mathcal{N}_r)$ , or
    $\infty$  if the combination can be pruned
3: {Update the number of pruned nodes, if needed.}
4: if  $\mathcal{N}_q$  not yet visited and is not the root node then
5:    $\text{pruned}(\mathcal{N}_q) \leftarrow \text{parent}(\mathcal{N}_q)$ 
6:    $\text{lb}(\mathcal{N}_q) \leftarrow \text{lb}(\text{parent}(\mathcal{N}_q))$ 
7: end if
8: if  $\text{pruned}(\mathcal{N}_q) = k - 1$  then return  $\infty$ 
9:  $s \leftarrow d_{\min}(\mathcal{N}_q, \mathcal{N}_r)$ 
10:  $c \leftarrow$  any descendant cluster centroid of  $\mathcal{N}_r$ 
11: if  $d_{\min}(\mathcal{N}_q, \mathcal{N}_r) > \text{ub}(\mathcal{N}_q)$  then
12:   {This cluster node owns no descendant points.}
13:   if  $d_{\min}(\mathcal{N}_q, \mathcal{N}_r) < \text{lb}(\mathcal{N}_q)$  then
14:     {Improve the lower bound for pruned nodes.}
15:      $\text{lb}(\mathcal{N}_q) \leftarrow d_{\min}(\mathcal{N}_q, \mathcal{N}_r)$ 
16:   end if
17:    $\text{pruned}(\mathcal{N}_q) += |\mathcal{D}_r^p \setminus \{\text{clusters not pruned}\}|$ 
18:    $s \leftarrow \infty$ 
19: else if  $d_{\max}(\mathcal{N}_q, c) < \text{ub}(\mathcal{N}_q)$  then
20:   {We may improve the upper bound.}
21:    $\text{ub}(\mathcal{N}_q) \leftarrow d_{\max}(\mathcal{N}_q, \mathcal{N}_r)$ ,  $\text{closest}(\mathcal{N}_q) \leftarrow c$ 
22: end if
23: {Check if all clusters (except one) are pruned.}
24: if  $\text{pruned}(\mathcal{N}_q) = k - 1$  then return  $\infty$ 
25: return  $s$ 
```

sets the correct value of $\text{canchange}(\cdot)$ for every point and node, we coalesce the tree for the next iteration with the `CoalesceTree()` function. Coalescing the tree is straightforward: we simply remove any nodes from the tree where $\text{canchange}(\cdot)$ is **false**. This leaves a smaller tree with no nodes where $\text{canchange}(\cdot)$ is **false**. Decoalescing the tree (`DecoalesceTree()`) is done by restoring the tree to its original state. See the appendix for more details.

Algorithm 4 UpdateCentroids().

```
1: Input: tree  $\mathcal{T}$  built on dataset  $S$ 
2: Output: new centroids  $C$ 
3:  $C := \{c_0, \dots, c_{k-1}\} \leftarrow \mathbf{0}^{k \times d}$ ;  $n = \mathbf{0}^k$ 
4: { $s$  is a stack.}
5:  $s \leftarrow \{\text{root}(\mathcal{T})\}$ 
6: while  $|s| > 0$  do
7:    $\mathcal{N}_i \leftarrow s.\text{pop}()$ 
8:   if  $\text{pruned}(\mathcal{N}_i) = k - 1$  then
9:     {The node is entirely owned by a cluster.}
10:     $j \leftarrow$  index of  $\text{closest}(\mathcal{N}_i)$ 
11:     $c_j \leftarrow c_j + |\mathcal{D}_i^p| \text{centroid}(\mathcal{N}_i)$ ;  $n_j \leftarrow n_j + |\mathcal{D}_i^p|$ 
12:   else
13:     {The node is not entirely owned by a cluster.}
14:     if  $|\mathcal{C}_i| > 0$  then  $s.\text{push}(\mathcal{C}_i)$ 
15:     else
16:       for  $p_i \in \mathcal{P}_i$  not yet considered
17:          $j \leftarrow$  index of  $\text{closest}(p_i)$ 
18:          $c_j \leftarrow c_j + p_i$ ;  $n_j \leftarrow n_j + 1$ 
19:       end if
20:     end while
21:   for  $c_i \in C$ , if  $n_i > 0$  then  $c_i \leftarrow c_i/n_i$ 
22:   return  $C$ 
```

6 Theoretical results

Space constraints allow us to only provide proof sketches for the first two theorems here. Detailed proofs are given in the appendix.

THEOREM 6.1. *A single iteration of dual-tree k -means as given in Algorithm 1 will produce exactly the same results as the brute-force $O(kN)$ implementation.*

Proof. (Sketch.) First, we show that the dual-tree algorithm (line 9) produces correct results for $\text{ub}(\cdot)$, $\text{lb}(\cdot)$, $\text{pruned}(\cdot)$, and $\text{closest}(\cdot)$ for every point and node. Next, we show that `UpdateTree()` maintains the correctness of those quantities and only sets $\text{canchange}(\cdot)$ to **false** when the node or point cannot change owner. It is easily shown that `CoalesceTree()` and `DecoalesceTree()` do not affect the results of the dual-tree algorithm because the only nodes and points removed are those where $\text{canchange}(\cdot) = \text{false}$. Lastly, we show that `UpdateCentroids()` produces centroids correctly.

Next, we consider the runtime of the algorithm. Our results are with respect to the *expansion constant* c_k of the centroids [6], which is a measure of intrinsic dimension. c_{qk} is a related quantity: the largest expansion constant of C plus any point in the dataset. Our results also depend on the imbalance of the tree $i_t(\mathcal{T})$, which in practice generally scales linearly in N [13]. As with the other theoretical results, more detail on each of these quantities is available in the appendix.

THEOREM 6.2. *When cover trees are used, a single iteration of dual-tree k -means as in Algorithm 1 can be performed in $O(c_k^4 c_{qk}^5 (N + i_t(\mathcal{T})) + c_k^9 k \log k)$ time.*

Proof. (Sketch.) Cover trees have $O(N)$ nodes [6], so all of the steps of the algorithm other than the tree building and dual-tree algorithm can be done in $O(N)$ time. Building a tree on the centroids takes $O(c_k^6 k \log k)$ time, where c_k is the expansion constant of the centroids. Using recent results [13], we may show that our pruning rules are at least as tight as nearest neighbor search; thus the dual-tree algorithm (line 11) may be performed in $O(c_{kr}^9 (N + i_t(\mathcal{T})))$ time. Also, we must perform nearest neighbor search on the centroids, which costs $O(c_k^9 (k + i_t(\mathcal{T})))$ time. This gives a total per-iteration runtime of $O(c_{kr}^9 (N + i_t(\mathcal{T})) + c_k^6 k \log k + c_k^9 i_t(\mathcal{T}_k))$.

This result holds intuitively. By building a tree on the centroids, we are able to prune many centroids at once, and as a result the amortized cost of finding the nearest centroid to a point is $O(1)$. This meshes with earlier theoretical results [6, 13, 33] and earlier empirical results [22, 21] that suggest that an answer can be obtained for a single query point in $O(1)$ time. Note that this worst-case bound depends on the expansion constant of the centroids, c_k , and the related quantity c_{qk} . If the intrinsic dimension of the centroids is low—that is, if the centroids are distributed favorably—the dual-tree algorithm will be more efficient.

However, this bound is generally quite loose in practice. First, runtime bounds for cover trees are known to be loose [13]. Second, this particular bound does not consider the effect of coalescing the tree. In any given iteration, especially toward the end of the k -means clustering, most points will have `change(·) = false` and thus the coalesced tree will be far smaller than the full tree built on all N points.

THEOREM 6.3. *Algorithm 1 uses no more than $O(N + k)$ memory when cover trees are used.*

Proof. This proof is straightforward. A cover tree on N points takes $O(N)$ space. So the trees and associated bounds take $O(N)$ and $O(k)$ space. Also, the dataset and centroids take $O(N)$ and $O(k)$ space.

7 Experiments

The next thing to consider is the empirical performance of the algorithm. We use the publicly available `kmeans` program in `mlpack` [12]; in our experiments, we run it as follows:

```
$ kmeans -i dataset.csv -I centroids.csv -c
    $k -v -e -a $algorithm
```

Dataset	N	d	tree build time	
			kd -tree	cover tree
cloud	2048	10	0.001s	0.005s
cup98b	95413	56	1.640s	32.41s
birch3	100000	2	0.037s	2.125s
phy	150000	78	4.138s	22.99s
power	2075259	7	7.342s	1388s
lcdm	6000000	3	4.345s	6214s

Table 4: Dataset information.

where `$k` is the number of clusters and `$algorithm` is the algorithm to be used. Each algorithm is implemented in C++. For the `yinyang` algorithm, we use the authors’ implementation. We use a variety of k values on mostly real-world datasets; details are shown in Table 4 [27, 34, 29]. The table also contains the time taken to build a kd -tree (for `blacklist` and `dualtree-kd`) and a cover tree (for `dualtree-ct`). Cover trees are far more complex to build than kd -trees; this explains the long cover tree build time. Even so, the tree only needs to be built once during the k -means run. If results are required for multiple values of k —such as in the X-means algorithm [32]—then the tree built on the points may be re-used.

Clusters were initialized using the Bradley-Fayyad refined start procedure [7]; however, this was too slow for the very large datasets, so in those cases points were randomly sampled as the initial centroids. k -means was then run until convergence on each dataset. These simulations were performed on a modest consumer desktop with an Intel i5 with 16GB RAM, using `mlpack`’s benchmarking system [18].

Average runtime per iteration results are shown in Table 5. The amount of work that is being pruned away is somewhat unclear from the runtime results, because the `elkan` and `hamerly` algorithms access points linearly and thus benefit from cache effects; this is not true of the tree-based algorithms. Therefore, the average number of distance calculations per iteration are also included in the results.

It is immediately clear that for large datasets, `dualtree-kd` is fastest, and `dualtree-ct` is almost as fast. `elkan`, because it holds kN bounds, is able to prune away a huge amount of work and is very fast for small datasets; however, maintaining all of these bounds becomes prohibitive with large k and the algorithm exhausts all available memory. `blacklist` has the same issue: on the largest datasets, with the largest k values, the space required to maintain all the blacklists is too much. This is also true of `yinyang`, which must maintain bounds between each point and each group of centroids. For large k , this burden becomes too much and the algorithm fails. `hamerly` and the dual-tree algorithms, on the other hand, are the best-behaved

dataset	k	iter.	avg. per-iteration runtime (distance calculations)					
			elkan	hamerly	yinyang	blacklist	dualtree-kd	dualtree-ct
cloud	3	8	1.50e-4s (867)	1.11e-4s (1.01k)	1.11e-1s (2.00k)	4.68e-5s (302)	1.27e-4s (278)	2.77e-4s (443)
cloud	10	14	2.09e-4s (1.52k)	1.92e-4s (4.32k)	7.66e-2s (9.55k)	1.55e-4s (2.02k)	3.69e-4s (1.72k)	5.36e-4s (2.90k)
cloud	50	19	5.87e-4s (2.57k)	5.30e-4s (21.8k)	9.66e-3s (15.6k)	8.20e-4s (12.6k)	1.23e-3s (5.02k)	1.09e-3s (9.84k)
cup98b	50	224	0.0445s (25.9k)	0.0557s (962k)	0.0465s (313k)	0.0409s (277k)	0.0955s (254k)	0.1089s (436k)
cup98b	250	168	0.1972s (96.8k)	0.4448s (8.40M)	0.1417s (898k)	0.2033s (1.36M)	0.4585s (1.38M)	0.3237s (2.73M)
cup98b	750	116	1.1719s (373k)	1.8778s (36.2M)	0.2653s (1.26M)	0.6365s (4.11M)	1.2847s (4.16M)	0.8056s (81.4M)
birch3	50	129	0.0194s (24.2k)	0.0093s (566k)	0.0378s (399k)	0.0030s (42.7k)	0.0082s (37.4k)	0.0378s (67.9k)
birch3	250	812	0.0895s (42.8k)	0.0314s (2.59M)	0.0711s (239k)	0.0164s (165k)	0.0183s (79.7k)	0.0485s (140k)
birch3	750	373	0.3253s (292k)	0.0972s (8.58M)	0.1423s (476k)	0.0554s (450k)	0.02989s (126k)	0.0581s (235k)
phy	50	34	0.0668s (82.3k)	0.1064s (1.38M)	0.1072s (808k)	0.0081s (33.0k)	0.02689s (67.8k)	0.0945s (188k)
phy	250	38	0.1627s (121k)	0.4634s (6.83M)	0.2469s (2.39M)	0.0249s (104k)	0.0398s (90.4k)	0.1023s (168k)
phy	750	35	0.7760s (410k)	2.9192s (43.8M)	0.6418s (5.61M)	0.2478s (1.19M)	0.2939s (1.10M)	0.3330s (1.84M)
power	25	4	0.3872s (2.98M)	0.2880s (12.9M)	1.1257s (33.5M)	0.0301s (216k)	0.0950s (87.4k)	0.6658s (179k)
power	250	101	2.6532s (425k)	0.1868s (7.83M)	1.2684s (10.3M)	0.1504s (1.13M)	0.1354s (192k)	0.6405s (263k)
power	1000	870	<i>out of memory</i>	6.2407s (389M)	4.4261s (9.41M)	0.6657s (2.98M)	0.4115s (1.57M)	1.1799s (4.81M)
power	5000	504	<i>out of memory</i>	29.816s (1.87B)	22.7550s (58.6M)	4.1597s (11.7M)	1.0580s (3.85M)	1.7070s (12.3M)
power	15000	301	<i>out of memory</i>	111.74s (6.99B)	<i>out of memory</i>	<i>out of memory</i>	2.3708s (8.65M)	2.9472s (30.9M)
lcdm	500	507	<i>out of memory</i>	6.4084s (536M)	8.8926s (44.5M)	0.9347s (4.20M)	0.7574s (3.68M)	2.9428s (7.03M)
lcdm	1000	537	<i>out of memory</i>	16.071s (1.31B)	18.004s (74.7M)	2.0345s (5.93M)	0.9827s (5.11M)	3.3482s (10.0M)
lcdm	5000	218	<i>out of memory</i>	64.895s (5.38B)	<i>out of memory</i>	12.909s (16.2M)	1.8972s (8.54M)	3.9110s (19.0M)
lcdm	20000	108	<i>out of memory</i>	298.55s (24.7B)	<i>out of memory</i>	<i>out of memory</i>	4.1911s (17.8M)	5.5771s (43.2M)

Table 5: Empirical results for k -means.

with memory usage and do not have any issues with large N or large k ; however, `hamerly` is very slow on large datasets because it is not able to prune many points at once.

Similar to the observations about the `blacklist` algorithm, the tree-based approaches are less effective in higher dimensions [31]. This is an important point: the performance of tree-based approaches suffer in high dimensions in part because the bound $d_{\min}(\cdot, \cdot)$ generally becomes looser as dimension increases. This is partly because the volume of nodes in high dimensions is much higher; consider that a ball has volume that is exponential in the dimension.

Even so, in our results, we see speedup in reasonable dimensions (for example, the `phy` dataset has 78 dimensions). Further, because our algorithm is tree-independent, we may use tree structures tailored to high-dimensional data [3]—including ones that have not yet been developed. From our results we believe as a rule of thumb that the dual-tree k -means algorithm can be effective up to a hundred dimensions or more.

Another clear observation is that when k is scaled on a single dataset, the `dualtree-kd` and `dualtree-ct` algorithms nearly always scale better (in terms of runtime) than the other algorithms. These results show that our algorithm satisfies its original goals: to be able to scale effectively to large k and N .

8 Conclusion and future directions

Using four pruning strategies, we have developed a flexible, tree-independent dual-tree k -means algorithm that is the best-performing algorithm for large k and large datasets in small-to-medium dimensions. It is theoretic-

ally favorable, has a small memory footprint, and may be used in conjunction with initial point selection and approximation schemes for additional speedup.

There are still interesting future directions to pursue, though. The first direction is parallelism: because our dual-tree algorithm is agnostic to the type of traversal used, we may use a parallel traversal [14], such as an adapted version of a recent parallel dual-tree algorithm [26]. The second direction is kernel k -means and other spectral clustering techniques: our algorithm may be merged with the ideas of [15] to perform kernel k -means. The third direction is theoretical. Recently, more general notions of intrinsic dimensionality have been proposed [24, 1]; these may enable tighter and more descriptive runtime bounds. Our work thus provides a useful and fast k -means algorithm and also opens promising avenues to further accelerated clustering algorithms.

References

- [1] L. Amsaleg, O. Chelly, T. Furon, S. Girard, M.E. Houle, K. Kawarabayashi, and M. Nett. Estimating local intrinsic dimensionality. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 29–38, 2015.
- [2] D. Arthur and S. Vassilvitskii. k -means++: The advantages of careful seeding. In *Proceedings of the 18th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1027–1035, 2007.
- [3] S. Arya, D.M. Mount, N.S. Netanyahu, R. Silverman, and A.Y. Wu. An optimal algorithm for approximate nearest neighbor searching in fixed dimensions. *Journal of the ACM (JACM)*, 45(6):891–923, 1998.
- [4] S. Bengio, J. Weston, and D. Grangier. Label embedding trees for large multi-class tasks. In *Advances in*

- Neural Information Processing Systems 23 (NIPS '10)*, volume 23, page 3, 2010.
- [5] J.L. Bentley. Multidimensional binary search trees used for associative searching. *Communications of the ACM*, 18(9):509–517, 1975.
 - [6] A. Beygelzimer, S.M. Kakade, and J. Langford. Cover trees for nearest neighbor. In *Proceedings of the 23rd International Conference on Machine Learning (ICML '06)*, pages 97–104, 2006.
 - [7] P.S. Bradley and U.M. Fayyad. Refining initial points for k-means clustering. In *Proceedings of the 15th International Conference on Machine Learning (ICML '98)*, pages 91–99, 1998.
 - [8] F. Can and E.A. Ozkarahan. Concepts and effectiveness of the cover-coefficient-based clustering methodology for text databases. *ACM Transactions on Database Systems*, 15(4):483–517, December 1990.
 - [9] A. Coates, A.Y. Ng, and H. Lee. An analysis of single-layer networks in unsupervised feature learning. In *Proceedings of AISTATS*, pages 215–223, 2011.
 - [10] G. Csurka, C. Dance, L. Fan, J. Willamowski, and C. Bray. Visual categorization with bags of keypoints. In *Workshop on Statistical Learning in Computer Vision, ECCV*, volume 1, pages 1–16, 2004.
 - [11] R.R. Curtin. *Improving dual-tree algorithms*. PhD thesis, Georgia Institute of Technology, August 2015.
 - [12] R.R. Curtin, J.R. Cline, N.P. Slagle, W.B. March, P. Ram, N.A. Mehta, and A.G. Gray. MLPACK: A scalable C++ machine learning library. *Journal of Machine Learning Research*, 14:801–805, 2013.
 - [13] R.R. Curtin, D. Lee, W.B. March, and P. Ram. Plug-and-play dual-tree algorithm runtime analysis. *arXiv preprint arXiv:1501.05222*, 2015.
 - [14] R.R. Curtin, W.B. March, P. Ram, D.V. Anderson, A.G. Gray, and C.L. Isbell Jr. Tree-independent dual-tree algorithms. In *Proceedings of the 30th International Conference on Machine Learning (ICML '13)*, pages 1435–1443, 2013.
 - [15] R.R. Curtin and P. Ram. Dual-tree fast exact max-kernel search. *Statistical Analysis and Data Mining*, 7(4):229–253, 2014.
 - [16] R.R. Curtin, P. Ram, and A.G. Gray. Fast exact max-kernel search. In *Proceedings of SIAM International Conference on Data Mining 2013 (SDM '13)*, pages 1–9, 2013.
 - [17] Y. Ding, Y. Zhao, X. Shen, M. Musuvathi, and T. Mytkowicz. Yinyang k-means: A drop-in replacement of the classic k-means with consistent speedup. In *Proceedings of The 32nd International Conference on Machine Learning (ICML '15)*, pages 579–587, 2015.
 - [18] M. Edel, A. Soni, and R.R. Curtin. An automatic benchmarking system. In *Proceedings of the NIPS 2014 Workshop on Software Engineering for Machine Learning*, 2014.
 - [19] C. Elkan. Using the triangle inequality to accelerate k-means. In *Proceedings of the 20th International Conference on Machine Learning (ICML '03)*, volume 3, pages 147–153, 2003.
 - [20] G. Frahling and C. Sohler. A fast k -means implementation using coresets. *International Journal of Computational Geometry & Applications*, 18(06):605–625, 2008.
 - [21] A.G. Gray and A.W. Moore. ‘N-Body’ problems in statistical learning. In *Advances in Neural Information Processing Systems 14 (NIPS 2001)*, volume 4, pages 521–527, 2001.
 - [22] A.G. Gray and A.W. Moore. Nonparametric density estimation: Toward computational tractability. In *SIAM International Conference on Data Mining (SDM)*, pages 203–211, 2003.
 - [23] G. Hamerly. Making k -means even faster. In *Proceedings of the 2010 SIAM International Conference on Data Mining*, pages 130–140, 2010.
 - [24] M.E. Houle. Dimensionality, discriminability, density and distance distributions. In *2013 IEEE 13th International Conference on Data Mining Workshops (ICDMW)*, pages 468–473, 2013.
 - [25] Y.C. Kwon, D. Nunley, J.P. Gardner, M. Balazinska, B. Howe, and S. Loebman. Scalable clustering algorithm for n-body simulations in a shared-nothing cluster. In *Scientific and Statistical Database Management*, pages 132–150. Springer, 2010.
 - [26] D. Lee, R.W. Vuduc, and A.G. Gray. A distributed kernel summation framework for general-dimension machine learning. In *Proceedings of the 2012 SIAM International Conference on Data Mining (SDM '12)*, pages 391–402, 2012.
 - [27] M. Lichman. UCI machine learning repository, 2013.
 - [28] T. Liu, A.W. Moore, K. Yang, and A.G. Gray. An investigation of practical approximate nearest neighbor algorithms. In *Advances in Neural Information Processing Systems 18 (NIPS '04)*, pages 825–832, 2004.
 - [29] R. Lupton, J.E. Gunn, Z. Ivezić, G.R. Knapp, and S. Kent. The SDSS imaging pipelines. In *Astronomical Data Analysis Software and Systems X*, volume 238, page 269, 2001.
 - [30] A.W. Moore. Very fast em-based mixture model clustering using multiresolution kd-trees. *Advances in Neural Information Processing Systems*, pages 543–549, 1999.
 - [31] D. Pelleg and A.W. Moore. Accelerating exact k-means algorithms with geometric reasoning. In *Proceedings of KDD '99*, pages 277–281. ACM, 1999.
 - [32] D. Pelleg and A.W. Moore. X-means: Extending k -means with efficient estimation of the number of clusters. In *Proceedings of the Seventeenth International Conference on Machine Learning (ICML '00)*, pages 727–734, 2000.
 - [33] P. Ram, D. Lee, W.B. March, and A.G. Gray. Linear-time algorithms for pairwise statistical problems. *Advances in Neural Information Processing Systems 22 (NIPS 2009)*, pages 1527–1535, 2009.
 - [34] T. Zhang, R. Ramakrishnan, and M. Livny. BIRCH: A new data clustering algorithm and its applications. *Data Mining and Knowledge Discovery*, 1(2):141–182, 1997.

A Dual-Tree Algorithm for Fast k -means Clustering With Large k : Supplementary Material

Ryan R. Curtin
ryan@ratml.org

Unfortunately, space constraints prevent in-depth explanation of each of the points in the main paper. This supplementary material is meant to clarify all of the parts of the dual-tree k -means algorithm that space did not permit in the main paper.

1 Updating the tree

In addition to updating the centroids, the bounding information contained within the tree must be updated according to pruning strategies 3 and 4. Unfortunately, this yields a particularly complex recursive algorithm, given in Algorithm 1.

The first if statement (lines 4–10) catches the case where the parent cannot change owner next iteration; in this case, the parent’s upper bound and lower bound can be taken as valid bounds. In addition, the upper and lower bounds are adjusted to account for cluster movement between iterations, so that the bounds are valid for next iteration.

If the node \mathcal{N}_i has an owner, the algorithm then attempts to use the pruning rules established in Equations 4 and 6 in the main paper, to determine if the owner of \mathcal{N}_i can change next iteration. If not, `canchange(\mathcal{N}_i)` is set to `false` (line 18). On the other hand, if the pruning check fails, the upper bound is tightened and the pruning check is performed a second time. It is worth noting that $d_{\max}(\mathcal{N}_i, c_j)$ may not actually be less than the current value of `ub(\mathcal{N}_i)`, which is why the min is necessary.

After recursing into the children of \mathcal{N}_i , if \mathcal{N}_i could have an owner change, each point is individually checked using the same approach (lines 31–45). However, there is a slight difference: if a point’s owner can change, the upper and lower bounds must be set to ∞ (lines 44–45). This is only necessary with points; `BaseCase()` does not take bounding information from previous iterations into account, because no work can be avoided in that way.

Then, we may set `canchange(\mathcal{N}_i)` to `false` if every point in \mathcal{N}_i and every child of \mathcal{N}_i cannot change owners (and the points and nodes do not necessarily have to have the same owner). Otherwise, we must set

`pruned(\mathcal{N}_i)` to 0 for the next iteration.

2 Coalescing the tree

After `UpdateTree()` is called, the tree must be coalesced to remove any nodes where `canchange(\cdot) = false`. This can be accomplished via a single pass over the tree. A simple implementation is given in Algorithm 2. `DecoalesceTree()` may be implemented by simply restoring a pristine copy of the tree which was cached right before `CoalesceTree()` is called.

3 Runtime bound proof

We can use adaptive algorithm analysis techniques in order to bound the running time of Algorithm 1 in the main paper, based on [2] and [1]. This analysis depends on the *expansion constant*, which is a measure of intrinsic dimension defined below, originally from [6].

DEFINITION 3.1. *Let $B_S(p, \Delta)$ be the set of points in S within a closed ball of radius Δ around some $p \in S$ with respect to a metric d :*

$$(3.1) \quad B_S(p, \Delta) = \{r \in S : d(p, r) \leq \Delta\}.$$

*Then, the **expansion constant** of S with respect to the metric d is the smallest $c \geq 2$ such that*

$$(3.2) \quad |B_S(p, 2\Delta)| \leq c|B_S(p, \Delta)| \quad \forall p \in S, \quad \forall \Delta > 0.$$

The expansion constant is a bound on the number of points which fall into balls of increasing sizes. A low expansion constant generally means that search tasks like nearest neighbor search can be performed quickly with trees, whereas a high expansion constant implies a difficult dataset. Thus, if we assume a bounded expansion constant like in previous theoretical works [1, 8, 6, 3, 2], we may assemble a runtime bound that reflects the difficulty of the dataset.

Our theoretical analysis will concern the cover tree in particular. The cover tree is a complex data structure with appealing theoretical properties. We will only

Algorithm 1 UpdateTree() for dual-tree k -means.

```
1: Input: node  $\mathcal{N}_i$ ,  $\text{ub}(\cdot)$ ,  $\text{lb}(\cdot)$ ,  $\text{pruned}(\cdot)$ ,  $\text{closest}(\cdot)$ ,  $\text{canchange}(\cdot)$ , centroid movements  $m$ 
2: Output: updated  $\text{ub}(\cdot)$ ,  $\text{lb}(\cdot)$ ,  $\text{pruned}(\cdot)$ ,  $\text{canchange}(\cdot)$ 

3:  $\text{canchange}(\mathcal{N}_i) \leftarrow \text{true}$ 
4: if  $\mathcal{N}_i$  has a parent and  $\text{canchange}(\text{parent}(\mathcal{N}_i)) = \text{false}$  then
5:   {Use the parent's bounds.}
6:    $\text{closest}(\mathcal{N}_i) \leftarrow \text{closest}(\text{parent}(\mathcal{N}_i))$ 
7:    $j \leftarrow \text{index of closest}(\mathcal{N}_i)$ ,  $\text{canchange}(\mathcal{N}_i) \leftarrow \text{false}$ 
8:    $\text{ub}(\mathcal{N}_i) \leftarrow \text{ub}(\mathcal{N}_i) + m_j$ ,  $\text{lb}(\mathcal{N}_i) \leftarrow \text{lb}(\mathcal{N}_i) + \max_i m_i$ 
9: else if  $\text{pruned}(\mathcal{N}_i) = k - 1$  then
10:  { $\mathcal{N}_i$  is owned by a single cluster. Can that owner change next iteration?}
11:   $j \leftarrow \text{index of closest}(\mathcal{N}_i)$ 
12:   $\text{ub}(\mathcal{N}_i) \leftarrow \text{ub}(\mathcal{N}_i) + m_j$ ,  $\text{lb}(\mathcal{N}_i) \leftarrow \max(\text{lb}(\mathcal{N}_i) - \max_i m_i, \min_{k \neq j} d(c_k, c_j)/2)$ 
13:  if  $\text{ub}(\mathcal{N}_i) < \text{lb}(\mathcal{N}_i)$  then
14:    {The owner cannot change next iteration.}
15:     $\text{canchange}(\mathcal{N}_i) \leftarrow \text{false}$ 
16:  else
17:    {Tighten the upper bound and try to prune again.}
18:     $\text{ub}(\mathcal{N}_i) \leftarrow \min(\text{ub}(\mathcal{N}_i), d_{\max}(\mathcal{N}_i, c_j))$ 
19:    if  $\text{ub}(\mathcal{N}_i) < \text{lb}(\mathcal{N}_i)$  then  $\text{canchange}(\mathcal{N}_i) \leftarrow \text{false}$ 
20:  end if
21: else
22:   $j \leftarrow \text{index of closest}(\mathcal{N}_i)$ 
23:   $\text{ub}(\mathcal{N}_i) \leftarrow \text{ub}(\mathcal{N}_i) + m_j$ ,  $\text{lb}(\mathcal{N}_i) \leftarrow \text{lb}(\mathcal{N}_i) - \max_k m_k$ 
24: end if
25: {Recurse into each child.}
26: for each child  $\mathcal{N}_c$  of  $\mathcal{N}_i$ , call UpdateTree( $\mathcal{N}_c$ )
27: {Try to determine points whose owner cannot change if  $\mathcal{N}_i$  can change owners.}
28: if  $\text{canchange}(\mathcal{N}_i) = \text{true}$  then
29:   for  $p_i \in \mathcal{P}_i$  do
30:      $j \leftarrow \text{index of closest}(p_i)$ 
31:      $\text{ub}(p_i) \leftarrow \text{ub}(p_i) + m_j$ ,  $\text{lb}(p_i) \leftarrow \min(\text{lb}(p_i) - \max_k m_k, \min_{k \neq j} d(c_k, c_j)/2)$ 
32:     if  $\text{ub}(p_i) < \text{lb}(p_i)$  then
33:        $\text{canchange}(p_i) \leftarrow \text{false}$ 
34:     else
35:       {Tighten the upper bound and try again.}
36:        $\text{ub}(p_i) \leftarrow \min(\text{ub}(p_i), d(p_i, c_j))$ 
37:       if  $\text{ub}(p_i) < \text{lb}(p_i)$  then
38:          $\text{canchange}(p_i) \leftarrow \text{false}$ 
39:       else
40:         {Point cannot be pruned.}
41:          $\text{ub}(p_i) \leftarrow \infty$ ,  $\text{lb}(p_i) \leftarrow \infty$ 
42:       end if
43:     end if
44:   end for
45: else
46:   for  $p_i \in \mathcal{P}_i$  where  $\text{canchange}(p_i) = \text{false}$  do
47:     {Maintain upper and lower bounds for points whose owner cannot change.}
48:      $j \leftarrow \text{index of closest}(p_i)$ 
49:      $\text{ub}(p_i) \leftarrow \text{ub}(p_i) + m_j$ ,  $\text{lb}(p_i) \leftarrow \text{lb}(p_i) - \max_k m_k$ 
50:   end for
51: end if
52: if  $\text{canchange}(\cdot) = \text{false}$  for all children  $\mathcal{N}_c$  of  $\mathcal{N}_i$  and all points  $p_i \in \mathcal{P}_i$  then  $\text{canchange}(\mathcal{N}_i) \leftarrow \text{false}$ 
53: if  $\text{canchange}(\mathcal{N}_i) = \text{true}$  then  $\text{pruned}(\mathcal{N}_i) \leftarrow 0$ 
```

Algorithm 2 CoalesceTree() for dual-tree k -means.

```
1: Input: tree  $\mathcal{T}$ 
2: Output: coalesced tree  $\mathcal{T}$ 

3: {A depth-first recursion to hide nodes where
   canchange( $\cdot$ ) is false.}
4:  $s \leftarrow \{\text{root}(\mathcal{T})\}$ 
5: while  $|s| > 0$  do
6:    $\mathcal{N}_i \leftarrow s.\text{pop}()$ 

7:   {Special handling is required for leaf nodes and
   the root node.}
8:   if  $|\mathcal{C}_i| = 0$  then
9:     continue
10:  else if  $\mathcal{N}_i$  is the root node then
11:    for  $\mathcal{N}_c \in \mathcal{C}_i$  do
12:       $s.\text{push}(\mathcal{N}_c)$ 
13:    end for
14:  end if

15:  {See if children can be removed.}
16:  for  $\mathcal{N}_c \in \mathcal{C}_i$  do
17:    if canchange( $\mathcal{N}_c$ ) = false then
18:      remove child  $\mathcal{N}_c$ 
19:    else
20:       $s.\text{push}(\mathcal{N}_c)$ 
21:    end if
22:  end for

23:  {If only one child is left, then this node is unnec-
   essary.}
24:  if  $|\mathcal{C}_i| = 1$  then
25:    add child to parent( $\mathcal{N}_i$ )
26:    remove  $\mathcal{N}_i$  from parent( $\mathcal{N}_i$ )'s children
27:  end if
28: end while

29: return  $\mathcal{T}$ 
```

summarize the relevant properties here. Interested readers should consult the original cover tree paper [1] and later analyses [8, 2] for a complete understanding.

A cover tree is a leveled tree; that is, each cover tree node \mathcal{N}_i is associated with an integer scale s_i . The node with largest scale is the root of the tree; each node's scale is greater than its children's. Each node \mathcal{N}_i holds one point p_i , and every descendant point of \mathcal{N}_i is contained in the ball centered at p_i with radius 2^{s_i+1} . Further, every cover tree satisfies the following three invariants [1]:

- (*Nesting.*) When a point p_i is held in a node at some scale s_i , then each smaller scale will also have a node containing p_i .
- (*Covering tree.*) For every point p_i held in a node

\mathcal{N}_i at scale s_i , there exists a node with point p_j and scale $s_j + 1$ which is the parent of \mathcal{N}_i , and $d(p_i, p_j) < 2^{s_i+1}$.

- (*Separation.*) Given distinct nodes \mathcal{N}_i holding p_i and \mathcal{N}_j holding p_j both at scale s_i , $d(p_i, p_j) > 2^{s_i}$.

A useful result shows there are $O(N)$ points in a cover tree (Theorem 1, [1]). Another measure of importance of a cover tree is the *cover tree imbalance*, which aims to capture how well the data is distributed throughout the cover tree. For instance, consider a tree where the root, with scale s_r , has two nodes; one node corresponds to a single point and has scale $-\infty$, and the other node has scale $s_r - 1$ and contains every other point in the dataset as a descendant. This is very imbalanced, and a tree with many situations like this will not perform well for search tasks. Below, we reiterate the definition of cover tree imbalance from [2].

DEFINITION 3.2. *The cover node imbalance $i_n(\mathcal{N}_i)$ for a cover tree node \mathcal{N}_i with scale s_i in the cover tree \mathcal{T} is defined as the cumulative number of missing levels between the node and its parent \mathcal{N}_p (which has scale s_p). If the node is a leaf child (that is, $s_i = -\infty$), then number of missing levels is defined as the difference between s_p and $s_{\min} - 1$ where s_{\min} is the smallest scale of a non-leaf node in \mathcal{T} . If \mathcal{N}_i is the root of the tree, then the cover node imbalance is 0. Explicitly written, this calculation is*

$$(3.3) \quad i_n(\mathcal{N}_i) = \begin{cases} s_p - s_i - 1 & \text{if } \mathcal{N}_i \text{ is not a} \\ & \text{leaf and not} \\ & \text{the root node} \\ \max(s_p - s_{\min} - 1, 0) & \text{if } \mathcal{N}_i \text{ is a leaf} \\ 0 & \text{if } \mathcal{N}_i \text{ is the root.} \end{cases}$$

This simple definition of cover node imbalance is easy to calculate, and using it, we can generalize to a measure of imbalance for the full tree.

DEFINITION 3.3. *The cover tree imbalance $i_t(\mathcal{T})$ for a cover tree \mathcal{T} is defined as the cumulative number of missing levels in the tree. This can be expressed as a function of cover node imbalances easily:*

$$(3.4) \quad i_t(\mathcal{T}) = \sum_{\mathcal{N}_i \in \mathcal{T}} i_n(\mathcal{N}_i).$$

Bounding $i_t(\mathcal{T})$ is non-trivial, but empirical results suggest that imbalance scales linearly with the size

of the dataset, when the expansion constant is well-behaved. A bound on $i_t(\mathcal{T})$ is still an open problem at the time of this writing.

With these terms introduced, we may introduce a slightly adapted result from [2], which bounds the running time of nearest neighbor search.

THEOREM 3.1. (Theorem 2, [2].) *Using cover trees, the standard cover tree pruning dual-tree traversal, and the nearest neighbor search `BaseCase()` and `Score()` as given in Algorithms 2 and 3 of [2], respectively, and also given a reference set S_r with expansion constant c_r , and a query set S_q , where the range of pairwise distances in S_r is completely contained in the range of pairwise distances in S_q , the running time of nearest neighbor search is bounded by $O(c_r^4 c_{qr}^5 (N + i_t(\mathcal{T}_q)))$, where $c_{qr} = \max((\max_{p_q \in S_q} c'_r), c_r)$, where c'_r is the expansion constant of the set $S_r \cup \{p_q\}$.*

Now, we may adapt this result slightly.

THEOREM 3.2. *The dual-tree k -means algorithm with `BaseCase()` as in Algorithm 2 in the main paper and `Score()` as in Algorithm 3 in the main paper, with a point set S_q that has expansion constant c_q and size N , and k centroids C with expansion constant c_k , takes no more than $O(c_k^4 c_{qk}^5 (N + i_t(\mathcal{T}_q)))$ time.*

Proof. Both `Score()` and `BaseCase()` for dual-tree k -means can be performed in $O(1)$ time. In addition, the pruning of `Score()` for dual-tree k -means is at least as tight as `Score()` for nearest neighbor search: the pruning rule in Equation 2 in the main paper is equivalent to the pruning rule for nearest neighbor search. Therefore, dual-tree k -means can visit no more nodes than nearest neighbor search would with query set S_q and reference set C . Lastly, note that the range of pairwise distances of C will be entirely contained in the range of pairwise distances in S_q , to see that we can use the result of Theorem 3.1. Adapting that result, then, yields the statement of the algorithm.

The expansion constant of the centroids, c_k , may be understood as the intrinsic dimensionality of the centroids C . During each iteration, the centroids change, so those iterations that have centroids with high intrinsic dimensionality cannot be bounded as tightly. More general measures of intrinsic dimensionality, such as those recently proposed by Houle [5], may make the connection between c_q and c_k clear.

Next, we turn to bounding the entire algorithm.

THEOREM 3.3. *A single iteration of the dual-tree k -means algorithm on a dataset S_q using the cover tree \mathcal{T} , the standard cover tree pruning dual-tree traversal,*

`BaseCase()` as given in Algorithm 2 in the main paper, `Score()` as given in Algorithm 3 in the main paper, will take no more than

$$(3.5) \quad O(c_k^4 c_{qk}^5 (N + i_t(\mathcal{T})) + c_k^9 k \log k)$$

time, where c_k is the expansion constant of the centroids, c_{qk} is defined as in Theorem 3.2, and $i_t(\mathcal{T})$ is the imbalance of the tree as defined in Definition 3.3.

Proof. Consider each of the steps of the algorithm individually:

- `CoalesceNodes()` can be performed in a single pass of the cover tree \mathcal{N} , which takes $O(N)$ time.
- Building a tree on the centroids (\mathcal{T}_c) takes $O(c_k^6 k \log k)$ time due to the result for cover tree construction time [1].
- The dual-tree algorithm takes $O(c_k^4 c_{qk}^5 (N + i_t(\mathcal{T})))$ time due to Theorem 3.2.
- `DecoalesceNodes()` can be performed in a single pass of the cover tree \mathcal{N} , which takes $O(N)$ time.
- `UpdateCentroids()` can be performed in a single pass of the cover tree \mathcal{N} , so it also takes $O(N)$ time.
- `UpdateTree()` depends on the calculation of how much each centroid has moved; this costs $O(k)$ time. In addition, we must find the nearest centroid of every centroid; this is nearest neighbor search, and we may use the runtime bound for monochromatic nearest neighbor search for cover trees from [8], so this costs $O(c_k^9 k)$ time. Lastly, the actual tree update visits each node once and iterates over each point in the node. Cover tree nodes only hold one point, so each visit costs $O(1)$ time, and with $O(N)$ nodes, the entire update process costs $O(N)$ time. When we consider the preprocessing cost too, the total cost of `UpdateTree()` per iteration is $O(c_k^9 k + N)$.

We may combine these into a final result:

$$(3.6) \quad O(N) + O(c_k^6 k \log k) + O(c_k^4 c_{qk}^5 (N + i_t(\mathcal{T}))) + O(N) + O(N) + O(c_k^9 k + N)$$

and after simplification, we get the statement of the theorem:

$$(3.7) \quad O(c_k^4 c_{qk}^5 (N + i_t(\mathcal{T})) + c_k^9 k \log k).$$

Therefore, we see that under some assumptions on the data, we can bound the runtime of the dual-tree k -means algorithm to something tighter than $O(kN)$ per iteration. As expected, we are able to amortize the cost of k across all N nodes, giving amortized $O(1)$ search for the nearest centroid per point in the dataset. This is similar to the results for nearest neighbor search, which obtain amortized $O(1)$ search for a single query point. Also similar to the results for nearest neighbor search is that the search time may, in the worst case, degenerate to $O(kN + k^2)$ when the assumptions on the dataset are not satisfied. However, empirical results [9, 4, 7, 1] show that well-behaved datasets are common in the real world, and thus degeneracy of the search time is uncommon.

Comparing this bound with the bounds for other algorithms is somewhat difficult; first, none of the other algorithms have bounds which are adaptive to the characteristics of the dataset. It is possible that the blacklist algorithm could be refactored to use the cover tree, but even if that was done it is not completely clear how the running time could be bounded. How to apply the expansion constant to an analysis of Hamerly’s algorithm and Elkan’s algorithm is also unclear at the time of this writing.

Lastly, the bound we have shown above is potentially loose. We have reduced dual-tree k -means to the problem of nearest neighbor search, but our pruning rules are tighter. Dual-tree nearest neighbor search assumes that every query node will be visited (this is where the $O(N)$ in the bound comes from), but dual-tree k -means can prune a query node entirely if all but one cluster is pruned (Strategy 2). These bounds do not take this pruning strategy into account, and they also do not consider the fact that coalescing the tree can greatly reduce its size. These would be interesting directions for future theoretical work.

References

- [1] A. Beygelzimer, S.M. Kakade, and J. Langford. Cover trees for nearest neighbor. In *Proceedings of the 23rd International Conference on Machine Learning (ICML '06)*, pages 97–104, 2006.
- [2] R.R. Curtin, D. Lee, W.B. March, and P. Ram. Plug-and-play dual-tree algorithm runtime analysis. *arXiv preprint arXiv:1501.05222*, 2015.
- [3] R.R. Curtin and P. Ram. Dual-tree fast exact max-kernel search. *Statistical Analysis and Data Mining*, 7(4):229–253, 2014.
- [4] A.G. Gray and A.W. Moore. ‘N-Body’ problems in statistical learning. In *Advances in Neural Information Processing Systems 14 (NIPS 2001)*, volume 4, pages 521–527, 2001.
- [5] M.E. Houle. Dimensionality, discriminability, density and distance distributions. In *2013 IEEE 13th International Conference on Data Mining Workshops (ICDMW)*, pages 468–473, 2013.
- [6] D.R. Karger and M. Ruhl. Finding nearest neighbors in growth-restricted metrics. In *Proceedings of the Thirty-Fourth Annual ACM Symposium on Theory of Computing (STOC 2002)*, pages 741–750, 2002.
- [7] W.B. March, P. Ram, and A.G. Gray. Fast Euclidean minimum spanning tree: algorithm, analysis, and applications. In *Proceedings of the 16th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '10)*, pages 603–612, 2010.
- [8] P. Ram, D. Lee, W.B. March, and A.G. Gray. Linear-time algorithms for pairwise statistical problems. *Advances in Neural Information Processing Systems 22 (NIPS 2009)*, pages 1527–1535, 2009.
- [9] P. Ram, D. Lee, H. Ouyang, and A.G. Gray. Rank-approximate nearest neighbor search: Retaining meaning and speed in high dimensions. *Advances in Neural Information Processing Systems*, 22, 2009.