
Hierarchical Agglomerative Graph Clustering in Nearly-Linear Time

Laxman Dhulipala¹ David Eisenstat² Jakub Łacki² Vahab Mirrokni² Jessica Shi¹

Abstract

We study the widely used hierarchical agglomerative clustering (HAC) algorithm on edge-weighted graphs. We define an algorithmic framework for hierarchical agglomerative graph clustering that provides the first efficient $\tilde{O}(m)$ time exact algorithms for classic linkage measures, such as complete- and WPGMA-linkage, as well as other measures. Furthermore, for average-linkage, arguably the most popular variant of HAC, we provide an algorithm that runs in $\tilde{O}(n\sqrt{m})$ time. For this variant, this is the first exact algorithm that runs in subquadratic time, as long as $m = n^{2-\epsilon}$ for some constant $\epsilon > 0$. We complement this result with a simple ϵ -close approximation algorithm for average-linkage in our framework that runs in $\tilde{O}(m)$ time. As an application of our algorithms, we consider clustering points in a metric space by first using k -NN to generate a graph from the point set, and then running our algorithms on the resulting weighted graph. We validate the performance of our algorithms on publicly available datasets, and show that our approach can speed up clustering of point datasets by a factor of 20.7–76.5x.

1. Introduction

Clustering is a fundamental and widely used unsupervised learning technique with numerous applications in data mining, machine learning, and social network analysis. Hierarchical clustering is a popular approach to clustering which outputs a hierarchy of clusters, where the input data objects are singleton clusters at the bottom of the tree, with interior vertices corresponding to merging the two children clusters. In this paper, we consider the family of *hierarchical agglomerative clustering (HAC)* algorithms, which have attracted significant theoretical and practical attention since they were

¹MIT CSAIL ²Google Research. Correspondence to: Laxman Dhulipala <laxmandhulipala@gmail.com>.

first proposed nearly 50 years ago (King, 1967; Lance & Williams, 1967; Sneath & Sokal, 1973).

A HAC algorithm takes as input a set of n points and proceeds in $n - 1$ steps. In each step, it finds two most similar points and merges them together. Here, the notion of similarity is defined by a configurable *linkage measure*.

The specific choice of the linkage measure affects both the clustering quality and the computational complexity of the HAC algorithm. In the case of a general linkage measure, HAC can be implemented in $O(n^3)$ time, assuming that we are given the similarity between each two points as input. For the most commonly used linkage measure (single, average, Ward’s and complete linkage) this complexity can be improved to $O(n^2)$ by using the nearest-neighbor chain algorithm (Benzécri, 1982).

The $O(n^2)$ time algorithms for HAC are often referred to as optimal, given that they take the entire $n \times n$ similarity matrix as the input. However, from a practical point of view, this quadratic lower bound is very pessimistic, as in many applications only a small fraction of the $n \times n$ similarities are non-negligible. As an example, consider the problem of clustering search engine queries studied in (Beeferman & Berger, 2000). Each query is assigned a set of relevant URLs and the similarity between two queries is based on the overlap between their URL sets. Clearly, for the vast majority of query pairs, the similarity is zero.

Knowing that in practice only $o(n^2)$ pairs of input points have non-negligible similarity scores results in two natural questions. First, is it possible to design a subquadratic HAC algorithm in this case? And if so, does this algorithm lead to improved running times in practical applications? In this paper, we answer both these questions affirmatively.

1.1. Our Contributions

In this paper we study the HAC algorithm on edge-weighted graphs. Formally, we consider a graph $G(V, E)$ with n vertices and m edges, where each vertex represents one input point and edge weights describe the similarities between the endpoints.

We develop a general framework that encompasses both

average-linkage, and complete- and WPGMA-linkage, where common primitives used in HAC are modularized into a neighbor-heap data structure, which offers tradeoffs in the theoretical guarantees depending on the representation used. In particular, different applications of this framework result in the first subquadratic (and in many cases near-linear time) algorithms for several variants of HAC.

Complete-linkage and WPGMA-linkage. As a direct application of our framework, we obtain $\tilde{O}(m)$ -time exact algorithms for the complete-linkage and WPGMA-linkage measures. For these measures we assume that the similarity between pairs of vertices not connected by an edge is “undefined”. If we denote the undefined similarity by \perp , for each $x \in \mathbb{R}$ we have $\max(x, \perp) = x$ and $\mathcal{L}(x, \perp) = x$ for any linkage measure \mathcal{L} .

Average-linkage. As our main algorithmic result, we give an $\tilde{O}(n\sqrt{m})$ time-exact algorithm for HAC with the average-linkage (UPGMA) measure, as well as an $\tilde{O}(m)$ time approximate algorithm¹. Both algorithms assume that any two vertices not connected by an edge have a similarity equal to 0.

Average-linkage is arguably the most commonly used variant of HAC, due to its very good empirical performance (Shao et al., 2007; Moseley & Wang, 2017; Cohen-Addad et al., 2019). However, to the best of our knowledge, no subquadratic time algorithm for average-linkage HAC has been described prior to our work, even when $m = \Theta(n)$.

Our exact average-linkage HAC algorithm dynamically maintains a low-outdegree orientation of the graph, i.e., it assigns a direction to each edge, attempting to minimize the outdegrees of all vertices. At each step, the maximum outdegree is $O(\alpha)$, where α is the *arboricity* (see Section 1.3 for the definition) of the graph, which allows us to bound the amortized number of updates per cluster merge by $O(\alpha + \text{poly log } n)$.

The running time of the exact algorithm is derived from a more general bound. Consider a sequence of graphs $G_1 = G, G_2, \dots, G_{n-1}$ computed by the HAC algorithm. That is, G_{i+1} is obtained by contracting the largest weight edge in G_i and updating the edge weights using the chosen linkage measure. We show that the exact algorithm runs in $\tilde{O}(n \cdot \bar{\alpha})$ time, where $\bar{\alpha}$ is the maximum arboricity of the graphs G_1, \dots, G_{n-1} . The $\tilde{O}(n\sqrt{m})$ bound comes from the fact that arboricity of any m -edge graph is upper-bounded by $O(\sqrt{m})$. However, in many cases better bounds on graph arboricity are known. In particular, on planar graphs (and more generally graphs with an excluded minor) our algorithm runs in $\tilde{O}(n)$ time. Similarly, for graphs with treewidth at most t (Robertson & Seymour, 1986), the

¹Formally, for a constant ϵ we obtain an ϵ -close algorithm according to the definition of (Moseley et al., 2019).

overall running time is $\tilde{O}(nt)$.

Empirical evaluation. We provide efficient implementations of all the near-linear time algorithms that we give and study their empirical performance on large data sets.

Specifically we show an average speedup of 6.9x over a simple exact average-linkage algorithm on large real-world graphs. We have made our implementations publicly available on Github.²

Finally, we show that our efficient implementations can be used to obtain a significantly faster HAC algorithm, even if the input is a collection of points. Namely, we leverage an existing approximate nearest neighbor computation library (Guo et al., 2020) to compute a k-nearest neighbor graph on the input data points. We then cluster this graph using our efficient graph-based HAC implementation. As we show, the end-to-end time of finding nearest neighbors followed by running our HAC implementation is 20.7-76.5x faster than an efficient $O(n^2)$ implementation on point sets. At the same time, somewhat surprisingly, the quality of the clustering obtained by using our approximate method is on average the same as what is computed by the exact HAC, where for average-linkage, we achieve on average a 1.13% increase on the Adjusted Rand-Index score and a 1.06% increase on the Normalized Mutual Information score.

1.2. Related Work

The theoretical foundations of HAC algorithms have been well-studied (Dasgupta, 2016; Moseley & Wang, 2017), and have provided motivation for the use of certain variants of HAC in real-world settings (Roy & Pokutta, 2016; Charikar & Chatziafratis, 2017; Cohen-Addad et al., 2017; Charikar et al., 2019; Cohen-Addad et al., 2019). Moreover, the version of HAC that takes a graph as input has been studied before, especially in the context of graphs derived from point sets (Guha et al., 1999; Karypis et al., 1999; Franti et al., 2006), although without strong theoretical guarantees.

A rich body of work has focused on breaking the quadratic time barrier for HAC. In a recent major advancement, Abboud et al. (Abboud et al., 2019) showed that if the input points are in Euclidean space and Ward’s linkage method is used (Ward, 1963), only $\tilde{O}(n)$ similarity computations are needed. By using an efficient nearest-neighbor data structure, they obtained a subquadratic approximate HAC algorithm for Ward’s linkage measure.

Another special case is the single-linkage measure, again in the setting of Euclidean space and using approximate distances. By computing an approximate minimum spanning tree, one can obtain an approximate HAC solution

²<https://github.com/ParAlg/gbbs/tree/master/benchmarks/Clustering/SeqHAC>

that runs in $O(n \log n)$ total time (Smid, 2018). A related method is affinity clustering, which provides a parallel HAC algorithm inspired by Borůvka’s minimum spanning tree algorithm (Bateni et al., 2017).

A number of papers proposed subquadratic HAC algorithms by sacrificing theoretical guarantees on the quality of the result, see (Cochez & Mou, 2015) and the references therein. A prominent line of work leverages locality sensitive hashing to obtain improved running bounds. However, the improvements in the running time often come at a significant cost. For example, some of the algorithms (Cochez & Mou, 2015) do not come with a closed form expression giving the approximation ratio, or produce an incomplete dendrogram (by merging some datapoints together) (Gilpin et al., 2013).

Additionally, many different HAC implementations are available, both open source and proprietary (Wikipedia contributors, 2020). However, the fastest available implementations require $\Theta(n^2)$ time in the worst case (e.g. (Müller, 2013; The SciPy community, 2020; SAS Institute Inc., 2020)). Such implementations are typically capable of clustering up to tens of thousands of points in few minutes, and in these cases they are heavily optimized, for example by taking advantage of GPUs (Gorg Sissons, 2015). In comparison, by combining our efficient implementation with an efficient similarity search, we are able to cluster a dataset of almost $3 \cdot 10^5$ points in under two minutes.

Finally, we note that although subquadratic approximate algorithms for HAC are known in a few settings, to the best of our knowledge, the implementations of these algorithms are not publicly available.

1.3. Preliminaries

We denote a weighted graph by $G = (V, E, w)$, where $w : E \rightarrow \mathbb{R}$ assigns a weight to each edge. The number of vertices in a graph is $n = |V|$, and the number of edges is $m = |E|$. When reporting asymptotic bounds, we assume that $m = \Omega(n)$. Vertices are assumed to be indexed from 0 to $n - 1$. Unless otherwise mentioned, all graphs considered in this paper are undirected, and we assume that there are no self-edges or duplicate edges. We use $N(v)$ to denote the neighbors of vertex v and $d(v) = |N(v)|$ to denote its degree. We use $\text{Cut}(X, Y)$ to denote the set of edges between two sets of vertices X and Y .

The *arboricity* (α) of a graph is the minimum number of spanning forests needed to cover the graph. Note that α is upper bounded by $O(\sqrt{m})$ and lower bounded by $\Omega(m/n)$ (Chiba & Nishizeki, 1985). A *c-orientation* of a graph $G(V, E)$ directs each edge $e \in E$ such that the maximum out-degree of each vertex is at most c . It is well known that every arboricity α graph admits an α -orientation. The *dynamic graph-orientation* problem is to maintain an

orientation of a graph as it is modified over a series of edge insertions and deletions. The two quantities of interest are the *out-degree* of the maintained orientation and the *update time*, or the time taken to process each edge update. We say that an algorithm maintaining a c -orientation in Δ time per update is a (c, Δ) -*dynamic graph-orientation* algorithm.

2. Graph-Based HAC

Given a graph $G(V, E, w)$, the *hierarchical agglomerative clustering (HAC)* problem for a given linkage measure \mathcal{L} is to compute a dendrogram by repeatedly merging the two most similar clusters (the two clusters connected by the largest-weight edge) until only a single cluster remains. For simplicity, we assume that the graph contains a single connected component, although our algorithms and implementations handle graphs with multiple components. We treat clusters and vertices interchangeably in this paper. For example, we often refer to the degree $d(C)$ or neighbors $N(C)$ of a cluster C . We use $\mathcal{W}(C, D)$ to denote the weight of an edge between two clusters C and D . The size of a cluster $|C|$ is defined to be the number of initial (singleton) clusters that it contains.

Initially, there are n singleton clusters containing the vertices v_0, \dots, v_{n-1} . Clusters containing multiple vertices are generated over the course of the algorithm by merging existing clusters. The *merge* of two clusters X, Y results in a new cluster Z . The weights of edges incident to the cluster formed by the merge are given by a *linkage measure* (discussed below). A *dendrogram* is a vertex-weighted tree where the leaves are the initial clusters of the graph, the internal vertices correspond to clusters generated by merges, and the weight of a vertex created by merging two clusters X, Y is given by the weight (similarity) of the edge between the two vertices at the time they are merged.

Linkage Measures. A linkage measure specifies how to reweight edges incident to a cluster created by a merge. Many different linkage measures that have previously been studied are applicable in the graph-based setting. In *single-linkage*, the weight between two clusters (X, Y) is $\max_{(x,y) \in \text{Cut}(X,Y)} w(x, y)$, or the maximum-similarity edge between two vertices in X and Y . In *complete-linkage* the weight between two clusters (X, Y) is $\min_{(x,y) \in \text{Cut}(X,Y)} w(x, y)$, or the minimum-similarity edge between two vertices in X and Y . In the popular *unweighted average-linkage (UPGMA-linkage)* measure (often called the average-linkage measure) the similarity between two clusters (X, Y) is $\sum_{(x,y) \in \text{Cut}(X,Y)} w(x, y) / (|X| \cdot |Y|)$ or the total weight of inter-cluster edges between X and Y , normalized by the number of possible inter-cluster edges. The *weighted average-linkage (WPGMA-linkage)* measure is similar, but is defined in terms of the current weights. If a cluster Z is created by merging clusters X, Y , the similar-

ity of the edge between Z and a neighboring cluster U is $\frac{\mathcal{W}(X,U)+\mathcal{W}(Y,U)}{2}$ if both the (X,U) and (Y,U) edges exist, and otherwise just the weight of the existing edge.

We define the *best-neighbor* of a cluster X to be the cluster $\operatorname{argmax}_{Y \in N(X)} \mathcal{W}(X,Y)$. We call the edge connecting X and Y the *best-edge* of X . We say that a linkage measure is *reducible* (Benzécri, 1982), if for any three clusters X, Y, Z where X and Y are mutual best-neighbors, it holds that $\mathcal{W}(X \cup Y, Z) \geq \max(\mathcal{W}(X, Z), \mathcal{W}(Y, Z))$.

Our framework yields exact HAC algorithms for any reducible linkage measure that also satisfies the following property.

Definition 2.1. *A linkage measure is called triangle-based if it satisfies the following property. Consider any step of the algorithm which merges clusters B and C into a cluster $B \cup C$. Let A be a cluster distinct from B and C . Then, if edge (A, C) does not exist, $\mathcal{W}(A, B) = \mathcal{W}(A, B \cup C)$.*

In other words, the weight of an edge in an triangle-based linkage measure changing implies that the affected cluster (a cluster not participating in the merge) was part of a triangle with both of the clusters being merged.

Lemma 2.2. *Single-, complete-, and WPGMA-linkage are all triangle-based linkages.*

Unfortunately, while unweighted average-linkage is a reducible linkage measure, *it is not a triangle-based linkage*. We design a special algorithm for unweighted average-linkage in Section 4.

3. Algorithmic Framework

In this section we design an algorithmic framework for solving graph-based HAC for triangle-based linkage measures. We give two different algorithms, based on the nearest-neighbor chain and heap-based algorithms respectively from the classic literature on HAC.

Overview. There are two key substeps found in both the classic nearest-neighbor chain and heap-based algorithms: (i) *merging* two clusters to obtain a new cluster and (ii) finding the *best-edge* (edge to the most similar neighbor) out of a given cluster. The classic algorithms use simple ideas to implement both steps, for example, by using a linear-time merge for step (i), or by checking the similarity between all pairs of points for step (ii) in the case of the chain-based algorithm.

Unfortunately, directly applying these simple ideas to graphs yields algorithms with quadratic time-complexity. For example, implementing (i) using a linear-time merge algorithm will, on a n -vertex star graph ($m = O(n)$), take $\Theta(n^2)$ time for any sequence of merges. A similar example shows that exhaustively searching a neighbor-list for step (ii) in a

Algorithm 1 MERGE(A, B, \mathcal{L})

Input: Active clusters A and B , triangle-based linkage \mathcal{L} .

Output: Cluster ID of the remaining active cluster.

- 1: $(A, B) = (\operatorname{argmin}(d(A), d(B)), \operatorname{argmax}(d(A), d(B)))$.
 - 2: Remove B from $\text{HEAP}(A)$ (and vice versa).
 - 3: $\text{HEAP}(B) = \text{HEAP}(A) \cup \text{HEAP}(B)$ (using \mathcal{L} to merge the weights of $C \in \text{HEAP}(A) \cap \text{HEAP}(B)$).
 - 4: Mark cluster A as inactive.
 - 5: For each $C \in \text{HEAP}(A)$, update the cluster ID from A to B in $\text{HEAP}(C)$, using \mathcal{L} to merge if $B \in \text{HEAP}(C)$.
 - 6: Return B .
-

chain-based algorithm may take up to $\Theta(n^2)$ time.

We address these problems by noting that we can reuse the data structures for clusters that are being merged, and by using heap data structures that support efficient updates. For example, when merging two clusters, we can reuse data structures associated with both merged clusters since they are logically deleted after the merge. Furthermore, for efficiency we represent neighbors of a cluster using data structures that can merge two instances of size s, l where $s \leq l$ in $O(s \log(l/s + 1))$ time. Our analysis shows that we can perform any sequence of merges, while performing best-edge queries on the intermediate graphs in $\tilde{O}(m)$ time.

3.1. Algorithms

Data Structures and Common Primitives. Initially all $v \in V$ are *active* singleton clusters. Each cluster A maintains a *neighbor-heap* data structure $\text{HEAP}(A)$, which is abstractly a max-heap that stores the neighbors of cluster A . The heap elements are key-value pairs containing the neighbor’s cluster id and the weight (similarity) to the neighbor. The neighbor-heap supports the BESTEDGE operation, which returns the best (highest-priority) edge in H . In addition, the data structure supports the UNION operation, which given two neighbor-heaps H_1, H_2 , and a triangle-based linkage \mathcal{L} , creates $H_1 \cup H_2$, where the weights of the pairs in $\text{KEYS}(H_1) \cap \text{KEYS}(H_2)$ are merged using \mathcal{L} .

We consider two different representations of a neighbor-heap. The first is a deterministic representation using augmented balanced trees where the augmented values are the heap priorities (Blelloch et al., 2016). We also consider a representation of neighbor-heaps using mergeable heaps (e.g., Fibonacci heaps) combined with hash tables, which enables a faster implementation of the chain-based algorithm (discussed further in the appendix). If $s = \min(|H_1|, |H_2|)$, $l = \max(|H_1|, |H_2|)$, the cost of UNION is $O(s \log(l/s + 1))$ using augmented trees, and $O(s)$ amortized using mergeable heaps. The cost for BESTEDGE in both implementations is $O(\log n)$.

Algorithm 2 GraphHAC-Chain($G = (V, E, w), \mathcal{L}$)

Input: Edge weighted graph, G , triangle-based linkage \mathcal{L} .

Output: Dendrogram D for \mathcal{L} -HAC.

```

1: for each cluster  $v \in V$  do
2:   if  $v$  is active then
3:     Initialize stack  $S$ , initially containing only  $v$ .
4:     while  $S$  is not empty do
5:       Let  $t$  be TOP( $S$ ).
6:       Let  $(t, b, \mathcal{W}(t, b)) = \text{BESTEDGE}(t)$ .
7:       if  $b$  is already on  $S$  then
8:         POP( $S$ ).
9:         MERGE( $t$ , TOP( $S$ )). [Algorithm 1]
10:        POP( $S$ ).
11:       else
12:         Push  $b$  onto  $S$ .
    
```

Merging in both of our algorithms is performed using Algorithm 1. Given two clusters A and B , suppose without loss of generality that A has fewer neighbors. The algorithm merges A into B , where the weights to neighbors in $N(A) \cap N(B)$ are computed using the linkage measure \mathcal{L} (Lines 1–3). A is then marked as inactive (Line 4). Lastly, the algorithm maps over each neighboring cluster $C \in N(A)$, deletes the entry for A in $\text{HEAP}(C)$, and inserts this entry with the new cluster ID B , merging using \mathcal{L} if B already exists in $\text{HEAP}(C)$. Critically, this algorithm ensures that after it runs, all clusters which previously had edges to A now point to B , and that the weights of all edges to B are correctly updated using \mathcal{L} .

Next, we provide pseudocode for the two HAC algorithms in our framework. Note that both algorithms output a dendrogram D , but for simplicity we do not show the pseudocode of this step (the dendrogram can easily be maintained as part of Algorithm 1).

Algorithm 2 shows the pseudocode for our chain-based algorithm. The structure of our algorithm is similar to the classic nearest-neighbor chain algorithm (Murtagh, 1983), using a stack to maintain a path of best-neighbors and merging two vertices that are connected by a reciprocal best-edge.

Algorithm 3 shows the pseudocode for our heap-based algorithm. Our algorithm uses a *lazy* approach for handling edges in the global heap H which point to inactive clusters (Lines 5–7). Although we could potentially eagerly update the heap H when deactivating vertices in Algorithm 1 without an asymptotic increase in the running time, the lazy version is simpler to describe.

3.2. Analysis

We start with the following lemma, which bounds the cost of the MERGE operations performed in Algorithm 1 for any sequence of $n - 1$ merges.

Algorithm 3 GraphHAC-Heap($G = (V, E, w), \mathcal{L}$)

Input: Edge weighted graph, G , triangle-based linkage \mathcal{L} .

Output: Dendrogram D for \mathcal{L} -HAC.

```

1: Let  $H$  be a max-heap storing the highest-weight edge
   incident to each active cluster in the graph.
2: while  $|H| > 1$  do
3:   Let  $e = (u, v, \mathcal{W}(u, v))$  be the max edge in  $H$ .
4:   Delete  $e$  from  $H$ .
5:   if  $v$  is inactive then
6:     Let  $e' = (u, v', \mathcal{W}(u, v')) = \text{BESTEDGE}(u)$ .
7:     Insert  $e'$  into  $H$ .
8:   else
9:      $x = \text{MERGE}(u, v)$ . [Algorithm 1]
10:    Let  $e' = (x, y, \mathcal{W}(x, y)) = \text{BESTEDGE}(x)$ .
11:    Insert  $e'$  into  $H$ .
    
```

Lemma 3.1. *Let m_1, \dots, m_{n-1} be the sequence of merge operations performed by an algorithm where $m_i = (u_i, v_i)$. Let $\text{Cost}(m_i) = \min(d(u_i), d(v_i))$, where $d(u_i)$ and $d(v_i)$ are the degrees of the clusters when they are merged. Then, the total cost $\sum_{i=1}^{n-1} \text{Cost}(m_i) = O(m \log n)$.*

Theorem 3.2. *There are deterministic implementations of the chain-based and heap-based algorithms that run in $O(m \log^2 n)$ time for any triangle-based linkage \mathcal{L} .*

Proof sketch. We sketch the proof for the chain-based algorithm. The complete proof is given in the appendix.

Clearly, there are at most $n - 1$ MERGE operations and the total number of merged elements is bounded by $O(m \log n)$ due to Lemma 3.1. Using an augmented tree, the amortized cost of merging a single element is bounded by $O(\log n)$, which results in the total time of $O(m \log^2 n)$.

From the properties of the nearest-neighbor chain algorithm at most $2n$ elements are ever added to the stack, which implies that the BESTEDGE operation is called $O(n)$ times. Hence all BESTEDGE operations take $O(n \log n)$ time. \square

We show that in the case of the chain-based algorithm, we can obtain an asymptotically faster algorithm by using hash tables and Fibonacci heaps to represent the neighbor-heaps.

Theorem 3.3. *There is a randomized implementation of the chain-based algorithm that runs in $O(m \log n)$ time in expectation for any triangle-based linkage \mathcal{L} .*

4. Average Linkage

The key challenge for HAC using the average-linkage (UPGMA-linkage) measure is that merging two clusters into a new cluster affects the weights of *all edges* incident to the new cluster, and thus our framework from Section 3 is not directly applicable. We show that by carefully modifying

Algorithm 4 FLIPEDGE(A, B)

Input: Edge oriented from active clusters A to B .

Output: Edge oriented from B to A .

- 1: $w = \mathcal{W}(A, B)$. [true weight of the edge]
 - 2: Update the edge (A, B, w) in $\text{HEAP}(A)$.
-

Algorithm 5 UPDATEORIENTATION(A, B, \mathcal{EO})

Input: Active clusters A and B , dynamic orientation structure \mathcal{EO} .

- 1: $(A, B) = (\text{argmin}(d(A), d(B)), \text{argmax}(d(A), d(B)))$.
 - 2: For each $C \in N(A)$, delete (C, A) and insert (C, B) into the orientation data structure. Edge flips are handled using FLIPEDGE [Algorithm 4]
 - 3: **for** each edge (B, C) oriented out of B **do**
 - 4: $w = \mathcal{W}(B, C)$. [true weight of the edge]
 - 5: Update the edge (B, C, w) in $\text{HEAP}(C)$.
-

our framework, we can obtain an efficient exact algorithm for average-linkage that runs in sub-quadratic time.

Overview. Two simple ideas to support average-linkage in our framework are a *fully eager* approach, which updates all of the edges incident to a merged cluster, and a *fully lazy* approach, which updates none of the edges incident to a merged cluster and forces a $\text{BESTEDGE}(v)$ computation to spend $d(v)$ time. Unfortunately, simple examples show that both approaches can be forced to spend $\Theta(n^2)$ time for an $m = O(n)$ edge graph (e.g., a star on n vertices).

Our approach is to enable the HAC algorithm to perform BESTEDGE queries while only updating a *subset* of the edges incident to a newly-merged cluster. We achieve this by using a dynamic graph-orientation data structure, which maintains an dynamic $O(\alpha)$ -outdegree orientation where α is the arboricity of the current graph. Our observation is that using this data structure lets us maintain information about the current clustered graph using a *bounded amount of laziness*. Specifically, we maintain the invariant that each vertex stores the up-to-date weight of all its *incoming* edges.

We show that we can perform the i -th merge with an extra cost of $O(\alpha_i)$ (in addition to the merge-cost in the framework) where α_i is the arboricity of the current graph at the i -th merge. Similarly, we can perform a BESTEDGE operation in $O(\alpha_i)$ time. Using the fact that $\forall i, \alpha_i \leq \sqrt{m}$, we can obtain a sub-quadratic bound as long as the number of BESTEDGE operations is $o(n^2/\sqrt{m})$. Unfortunately, the heap-based algorithm could perform $O(m \log n)$ BESTEDGE computations, but we are guaranteed that the chain-based algorithm will only perform $O(n)$ BESTEDGE computations, and thus gives an algorithm with $\tilde{O}(n\sqrt{m})$ time-complexity.

Algorithm. The differences between our exact average-

link algorithm and Algorithm 2 are (1) that we maintain a dynamic graph-orientation structure \mathcal{EO} and (2) that we run extra procedures before performing the BESTEDGE and MERGE algorithms used by Algorithm 2. We also make a minor change in how the weights of edges in each cluster’s neighbor-heaps are stored.

Before a MERGE(A, B). Before merging two active clusters using MERGE (Algorithm 1) we call UPDATEORIENTATION (Algorithm 5). This algorithm updates \mathcal{EO} by deleting all edges going to the smaller deactivated cluster (A), and relabeling and inserting these edges to refer to the remaining active cluster (B). Note that the orientation data structure could cause a number of edges to have their orientation *flipped*. We inject Algorithm 4, which is called each time an edge is flipped and which updates the weight of the edge to its correct weight at the head of the new direction of the edge. The last step in Algorithm 5 is to update the weights for each of the edges oriented out of the active cluster B in the heap of this directed neighbor of B .

Before a BESTEDGE(A). Before extracting the best-edge from $\text{HEAP}(A)$, the algorithm updates each of the edges currently oriented out of A in \mathcal{EO} . Specifically, for such a directed edge (A, B) , it computes the true weight of this edge and updates this value in $\text{HEAP}(u)$. Performing this update is necessary, since B could have updated its size since the last time the (A, B) edge was updated in $\text{HEAP}(u)$.

Weight Representation in Neighbor-Heaps. If we store the weights of edges in each neighbor-heap, then when a cluster’s size increases through a merge, we must update all of the edges incident to the new cluster since the weights of all edges change. Instead, we store only part of the edge weights. Specifically, for an edge to B incident to cluster A we store $\frac{1}{|B|} \sum_{(a,b) \in \text{Cut}(A,B)} w(a, b)$ in $\text{HEAP}(A)$, and implicitly multiply this quantity by $1/|A|$, explicitly multiplying by this quantity when extracting an actual weight.

We obtain our exact average-linkage algorithm using a specific dynamic graph-orientation data structure. Specifically, we use the recent data structure of (Henzinger et al., 2020) which maintains a $(O(\alpha_i), O(\log^2 n))$ -dynamic graph-orientation data structure \mathcal{EO} . The out-degree of the orientation maintained is *adaptive*, i.e., the out-degree of vertices in \mathcal{EO} after the i -th update is $O(\alpha(G_i))$ where G_i is the graph at the time of the i -th update.

Theorem 4.1. *The exact average-linkage algorithm is correct and runs in $\tilde{O}(n\sqrt{m})$ time for arbitrary graphs.*

Proof. We provide a proof-sketch here, and provide a detailed proof in the appendix. Note that the dynamic graph-orientation data structure is only updated and used during the MERGE and BESTEDGE operations. Consider the i -th such operation, and let G_i be the graph induced by the current clustering at the time of this operation. It

is easy to check that the cost for a MERGE operation is $O(\alpha(G_i) + \text{Cost}(m_i) \log^2 n)$ and the cost for a BESTEDGE operation is $O(\log n + \alpha(G_i))$. Therefore, the total number of updates to \mathcal{EO} is $\sum_{i=1}^{n-1} \text{Cost}(m_i) = O(m \log n)$, and the overall time-complexity of updating \mathcal{EO} is $O(m \log^3 n)$. Since $\forall i, \alpha(G_i) \leq \sqrt{m}$ and there are $O(n)$ merge and best-edge operations, the time for updating the neighbor-heaps is $O(n\sqrt{m} + n \log n)$. Thus the total time-complexity of the algorithm is $\tilde{O}(n\sqrt{m} + m) = \tilde{O}(n\sqrt{m})$. \square

Our approach yields near-linear time exact algorithms for graphs whose minors all have bounded arboricity. Specifically, we show that for minor-closed graphs, such as planar graphs, bounded genus graphs, and bounded treewidth graphs we obtain algorithms with $\tilde{O}(n)$ time-complexity.

Corollary 4.2. *The exact average-linkage algorithm runs in $\tilde{O}(n \cdot t)$ time on any graph from a minor-closed graph family whose elements all have arboricity at most t .*

4.1. Approximation algorithm

Next, we show that average-linkage can be approximated in nearly-linear time. An ϵ -close HAC algorithm is an algorithm which only merges edges that have similarity at least $(1 - \epsilon) \cdot \mathcal{W}_{\max}$ where \mathcal{W}_{\max} is the largest weight currently in the graph (Moseley et al., 2019). An ϵ -close algorithm is constrained to merge an edge that is “close” in similarity to the merge the exact algorithm would perform. At the same time, the definition gives the algorithm flexibility in which edge it chooses to merge, which it can exploit to save work.

The idea of our algorithm is to maintain an extra counter for each cluster which stores the size the cluster had at the last time that the algorithm updated *all* of the incident edges of the cluster. Call this variable the *staleness*, $\mathcal{S}(A)$, of a given cluster A . Our algorithm maintains the following invariant:

Invariant 1. *For any active cluster A , $|A| < (1 + \epsilon)\mathcal{S}(A)$.*

We maintain this invariant by checking after merging two clusters if the size of the remaining active cluster A is still smaller than $(1 + \epsilon)\mathcal{S}(A)$. If the invariant is violated, then the algorithm performs a *rebuild* which updates the similarities of all edges incident to A to their true weights. Note that these updates are performed on the neighbor-heaps for both endpoints of the updated edges.

Since Invariant 1 can be violated at most $O(\log_{1+\epsilon} n)$ times, any given edge will be processed during a rebuild at most $O(\log_{1+\epsilon} n)$ times over the course of the algorithm, for a total cost of $O(m \log_{1+\epsilon} n)$.

We note that we only apply this approximation idea with the *heap-based* algorithm from our framework. Although it may be possible to combine this notion of approximation with the chain-based algorithm, the analysis needs to handle

the fact that the chain-based algorithm merges local minima instead of global minima, and so for simplicity we only consider the heap-based approach.

Theoretical Guarantees. To argue that our approximation algorithm yields a ϵ -close HAC algorithm for average-linkage, we show that the similarity of any edge stored within an active cluster’s neighbors cannot be much larger than its true edge similarity. Let the stored similarity of an edge (u, v) in the neighborhood be denoted $W_{\mathcal{S}}(u, v)$ and the true similarity of this edge be $W_{\mathcal{T}}(u, v)$.

Lemma 4.3. *Let $e = (U, V)$ be an edge in the neighborhood of an active cluster U . Then, $(1 + \epsilon)^{-2} W_{\mathcal{S}}(U, V) \leq W_{\mathcal{T}}(U, V)$.*

This lemma implies that if we merge this edge, then in the worst case, we will get a $(1 + \epsilon)^{-2}$ -close algorithm since the highest similarity edge could have a true similarity of $W_{\mathcal{S}}(e)$, but the edge merged by the algorithm could have a true similarity of $W_{\mathcal{T}}(e) = (1 + \epsilon)^{-2} W_{\mathcal{S}}(e)$. By setting the value of ϵ that we use internally appropriately, we obtain the following result.

Theorem 4.4. *There is an ϵ -close HAC algorithm for the average-linkage measure that runs in $O(m \log^2 n)$ time.*

5. Empirical Evaluation

We implemented our framework algorithms in C++ using the Graph Based Benchmark Suite (GBBS) (Dhulipala et al., 2018; 2020), and using the augmented maps from PAM (Blelloch et al., 2016) to represent neighbor-heaps. We provide more details about our implementations in the appendix. We run our experiments on a 72-core machine with 4×2.4 GHz Intel 18-core E7-8867 v4 Xeon processors, and 1TB of memory. Both GBBS and PAM are designed for parallel algorithms, but to enable a fair comparison with other sequential algorithms we disabled parallel execution.

5.1. Experiments

Approximate vs. Simple-Exact Algorithm. We start by evaluating the performance of our near-linear time approximation algorithm for average-linkage vs. a simple implementation of an exact average-linkage algorithm which updates the weights of all edges incident to a newly merged cluster. We ran this experiment on a collection of large real-world graphs from the SNAP datasets. Since these graphs are originally unweighted, we set the similarity of an edge (u, v) to $\frac{1}{\log(d(u)+d(v))}$. We provide more details about our graph inputs in the appendix. Figure 1 shows the result of the experiment. Our approximate average-linkage algorithm (using $\epsilon = 0.1$) achieves an average speedup of 6.9x over the exact average-linkage algorithm. We note that the dendrograms in the cases where the simple-exact

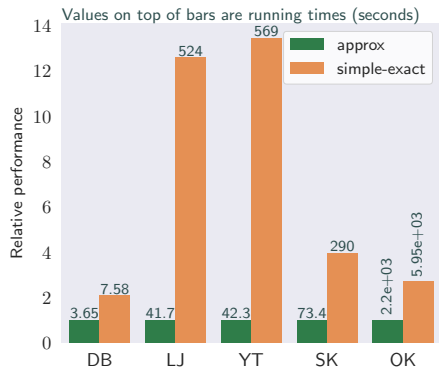


Figure 1. Relative performance of the approximate and simple-exact average-linkage algorithms on real-world graphs from the SNAP datasets, normalized to the fastest time per graph. The value on top of each bar is the running time in seconds.

algorithm performs reasonably well are shallower than in cases where the algorithm performs poorly (e.g., the DB dendrogram is 99x shallower than that of YT, although the number of vertices in DB is only 2.6x smaller). A simple upper-bound for the time-complexity of the simple-exact algorithm is $O(nD)$ where D is the depth of the dendrogram. For other linkage measures, such as single-, complete-, and WPGMA-linkage, we achieve up to 730x speedup over the simple-exact algorithm that spends $O(d(u) + d(v))$ time to merge two clusters, and note that the dendrograms observed for these measures have very high depth.

Comparison with Metric Clustering. Next, we study the quality and scalability of our graph-based HAC algorithms compared to metric HAC algorithms. Given an input pointset, P , we first apply an approximate k -NN algorithm to P to build an approximate k -NN graph. We use the state-of-the-art ScaNN k -NN library (Guo et al., 2020) for this graph-building step. We note that ScaNN internally uses multithreading, which we did not disable. We then symmetrize the k -NN graph and run our graph-based HAC implementation on it. We compare our results with those of the widely-used Scikit-learn (*sklearn*) package.

Quality. In the first set of experiments, we evaluate our algorithms and the four HAC variants supported by sklearn on the *iris*, *wine*, *digits*, and *cancer*, and *faces* classification datasets. We note that the heap-based and chain-based algorithms yielded the same dendrograms. To measure quality, we use the Adjusted Rand-Index (ARI) and Normalized Mutual Information (NMI) scores. The level of the tree with the highest score is used for evaluation.

We show the full quality scores in the appendix. Overall, our graph-based algorithms produce results that are comparable with, and sometimes superior to the results of the metric-

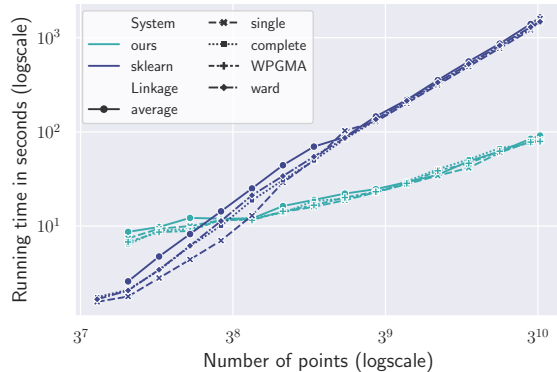


Figure 2. End-to-end running times of the sklearn and our graph-based algorithms on varying-size slices of the Fashion dataset.

based algorithms in sklearn. One exception is our complete-linkage algorithm, which is almost always worse than the sklearn algorithm, which is because the k -NN graph is missing large-distance edges which prevent cluster formation in the metric setting. We note that running our complete-linkage algorithm on the complete graph (with all distance edges) results in clustering results that match the quality of the sklearn algorithm. Our simple-exact and approximate average-linkage algorithm (using $\epsilon = 0.1$) achieve essentially the same quality results as the exact sklearn algorithm (our algorithms achieve on average 1.8% better ARI and 0.5% better NMI). Furthermore, the approximate and simple-exact algorithms yield identical quality results for all but the digits dataset, where the simple-exact algorithm is less than 1% better for both quality measures.

Scalability. In the second set of experiments, we study whether our approach can yield end-to-end speedups over the sklearn algorithms on large pointsets. We use the Fashion-MNIST (764-dimensions), Last.fm (65 dimensions), and NYTimes (256 dimensions) datasets in these experiments. We run both the sklearn and our algorithms on slices of these datasets to understand how the running time scales as the number of points to cluster increases.

Our results for the Fashion-MNIST dataset (shown in Figure 2) show that after about 10000 points, the end-to-end time of using the graph-based approach is always faster than using the $O(n^2)$ time metric-based algorithm. For the full Fashion-MNIST dataset, which contains 60,000 points, our approach yields an overall speedup of 20.7x.

In Figure 3 we show the results of the same experiment using the Last.fm and NYTimes datasets, but using only our approximate average-linkage algorithm to reduce clutter (this is the slowest algorithm out of all of our linkage-measures). We terminated algorithms that ran for longer than 1 hour, and were therefore unable to finish running

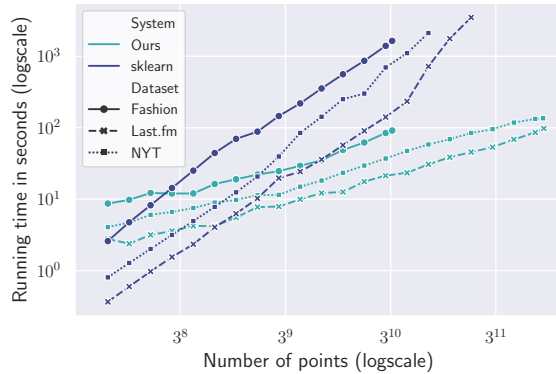


Figure 3. End-to-end running times of sklearn’s average-linkage and graph-based approximate average-linkage on varying-size slices of the Last.fm and NYTimes datasets.

the metric-based algorithm on the full Last.fm (292,385 points) and NYTimes (290,000 points) datasets. Our algorithms achieve up to 36.2x speedup on the NYTimes dataset and 76.5x speedup on the Last.fm dataset over the available datapoints for sklearn. Extrapolating from the trends of the sklearn implementations, a rough estimate suggests speedups of between 200x–500x for these datasets.

We note that we also ran the same experiments with the C++-based HAC implementations provided in SciPy ([The SciPy community, 2020](#)) and Fastcluster ([Müller, 2013](#)). We obtained running times that were within 10% of the running times of sklearn for all of our datasets and linkage functions, and thus only report the running times for sklearn in Figures 2 and 3.

Limits of our Approach. We observed that once we have generated a graph input, our algorithm’s performance scales almost linearly with the number of edges in the graph. Currently, the main bottleneck in our experiments for pointsets is the graph-building step which generates the k -NN graph using ScaNN and writes the k -NN graph to disk. Supplying the generated k -NN graph to our HAC algorithms without first writing it to disk will further accelerate our algorithms.

Ignoring the cost of the writing to disk, both the memory usage and running time of the graph-clustering step is lower than that of ScaNN. Specifically, the memory usage of our algorithms (excluding ScaNN) is approximately $56 \cdot m + O(n)$ bytes (where the constant on the n term is small). Therefore, our implementations can solve graphs with up to several billion vertices and 2–3 billion edges on a machine with 256GB of memory (obtainable from Google Cloud for a few dollars per hour). In terms of running time, we observed that our graph-based algorithms scale linearly with the number of edges in practice, and we could thus solve such a graph in between 12–24 hours.

6. Conclusion

In this paper we designed efficient HAC algorithms which run in near-linear time with respect to the number of input similarity pairs. We conducted a preliminary empirical evaluation, which shows that our algorithms achieve significant speedups while maintaining competitive clustering quality.

For future work, it would be very interesting to understand the parallel complexity of graph-based HAC, and to design efficient exact and approximate algorithms for these problems in a parallel or dynamic setting. From an experimental perspective, a significant challenge is to design HAC implementations that can be run on graphs with tens of billions of edges in a reasonable amount of time. Combining the ideas in this paper with an efficient dynamic graph processing system such as Aspen ([Dhulipala et al., 2019](#)) may be a first step towards such a result. Finally, an interesting open question is to design a near-linear time exact HAC algorithm for the unweighted average-linkage measure.

Acknowledgements. Thanks to D. Ellis Hershkowitz for helpful comments about this paper. We would also like to thank the anonymous reviewers for their helpful feedback and suggestions.

References

- Abboud, A., Cohen-Addad, V., and Houdrouge, H. Subquadratic high-dimensional hierarchical clustering. In *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 11580–11590, 2019.
- Batani, M., Behnezhad, S., Derakhshan, M., Hajiaghayi, M., Kiveris, R., Lattanzi, S., and Mirrokni, V. Affinity clustering: Hierarchical clustering at scale. In *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 6864–6874, 2017.
- Beeferman, D. and Berger, A. L. Agglomerative clustering of a search engine query log. In *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 407–416, 2000.
- Benzécri, J.-P. Construction d’une classification ascendante hiérarchique par la recherche en chaîne des voisins réciproques. *Cahiers de l’analyse des données*, 7(2):209–218, 1982.
- Blelloch, G. E., Ferizovic, D., and Sun, Y. Just join for parallel ordered sets. In *ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, pp. 253–264, 2016.
- Charikar, M. and Chatziafratis, V. Approximate hierarchical clustering via sparsest cut and spreading metrics. In *ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 841–854, USA, 2017.

- Charikar, M., Chatziafratis, V., and Niazadeh, R. *Hierarchical Clustering better than Average-Linkage*, pp. 2291–2304. 2019.
- Chiba, N. and Nishizeki, T. Arboricity and subgraph listing algorithms. *SIAM Journal on Computing*, 14(1):210–223, 1985.
- Cochez, M. and Mou, H. Twister tries: Approximate hierarchical agglomerative clustering for average distance in linear time. In *ACM SIGMOD International Conference on Management of Data*, pp. 505–517, 2015.
- Cohen-Addad, V., Kanade, V., and Mallmann-Trenn, F. Hierarchical clustering beyond the worst-case. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 30, pp. 6201–6209. Curran Associates, Inc., 2017.
- Cohen-Addad, V., Kanade, V., Mallmann-Trenn, F., and Mathieu, C. Hierarchical clustering: Objective functions and algorithms. *J. ACM*, 66(4), 2019.
- Dasgupta, S. A cost function for similarity-based hierarchical clustering. In *ACM Symposium on Theory of Computing (STOC)*, pp. 118–127, New York, NY, USA, 2016. Association for Computing Machinery.
- Dhulipala, L., Btleloch, G. E., and Shun, J. Theoretically efficient parallel graph algorithms can be fast and scalable. In *ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, pp. 293–304, 2018.
- Dhulipala, L., Btleloch, G. E., and Shun, J. Low-latency graph streaming using compressed purely-functional trees. In *ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*, pp. 918–934, 2019.
- Dhulipala, L., Shi, J., Tseng, T., Btleloch, G. E., and Shun, J. The graph based benchmark suite (GBBS). In *International Workshop on Graph Data Management Experiences and Systems (GRADES) and Network Data Analytics (NDA)*, pp. 11:1–11:8, 2020.
- Franti, P., Virtajoki, O., and Hautamaki, V. Fast agglomerative clustering using a k-nearest neighbor graph. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28(11):1875–1881, 2006.
- Gilpin, S., Qian, B., and Davidson, I. Efficient hierarchical clustering of large high dimensional datasets. In *ACM International Conference on Information and Knowledge Management*, pp. 1371–1380, 2013.
- Gorg Sissons. Gpu-accelerated R in the cloud with teraproc cluster-as-a-service, 2015. URL <https://developer.nvidia.com/blog/gpu-accelerated-r-cloud-teraproc-cluster-service/>. [Online; accessed 2-February-2021].
- Guha, S., Rastogi, R., and Shim, K. Rock: A robust clustering algorithm for categorical attributes. In *Proceedings of the International Conference on Data Engineering*, pp. 512, USA, 1999.
- Guo, R., Sun, P., Lindgren, E., Geng, Q., Simcha, D., Chern, F., and Kumar, S. Accelerating large-scale inference with anisotropic vector quantization. In *International Conference on Machine Learning (ICML)*, pp. 3887–3896, 2020.
- Henzinger, M., Neumann, S., and Wiese, A. Explicit and implicit dynamic coloring of graphs with bounded arboricity. *CoRR*, abs/2002.10142, 2020.
- Karypis, G., Han, E.-H. S., and Kumar, V. Chameleon: Hierarchical clustering using dynamic modeling. *Computer*, 32(8):68–75, August 1999.
- King, B. Step-wise clustering procedures. *Journal of the American Statistical Association*, 69:86–101, 1967.
- Lance, G. N. and Williams, W. T. A general theory of classificatory sorting strategies 1. Hierarchical systems. *Computer Journal*, 9(4):373–380, February 1967.
- Moseley, B. and Wang, J. R. Approximation bounds for hierarchical clustering: Average linkage, bisecting k-means, and local search. In *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 3094–3103, 2017.
- Moseley, B., Lu, K., Lattanzi, S., and Lavastida, T. A framework for parallelizing hierarchical clustering methods. In *ECML PKDD 2019*, 2019.
- Murtagh, F. A survey of recent advances in hierarchical clustering algorithms. *The computer journal*, 26(4):354–359, 1983.
- Müllner, D. fastcluster: Fast hierarchical, agglomerative clustering routines for R and python. *Journal of Statistical Software, Articles*, 53(9):1–18, 2013.
- Robertson, N. and Seymour, P. D. Graph minors II. Algorithmic aspects of tree-width. *J. Algorithms*, 7(3):309–322, 1986.
- Roy, A. and Pokutta, S. Hierarchical clustering via spreading metrics. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 29, pp. 2316–2324. Curran Associates, Inc., 2016.
- SAS Institute Inc. The CLUSTER procedure, 2020. URL <https://support.sas.com/rnd/app/stat/procedures/ClusterAnalysis.html>. [Online; accessed 2-February-2021].

Shao, J., Tanner, S. W., Thompson, N., and Cheatham, T. E. Clustering molecular dynamics trajectories: 1. characterizing the performance of different clustering algorithms. *Journal of Chemical Theory and Computation*, 3(6):2312–2334, 2007.

Smid, M. H. The well-separated pair decomposition and its applications., 2018.

Sneath, P. H. and Sokal, R. R. *Numerical Taxonomy: The Principles and Practice of Numerical Classification*. W.H. Freeman, San Francisco, 1973. ISBN 0 7167 0697 0.

The SciPy community. `scipy.cluster.hierarchy.linkage`, 2020. URL <https://docs.scipy.org/doc/scipy/reference/generated/scipy.cluster.hierarchy.linkage.html>. [Online; accessed 2-February-2021].

Ward, J. H. Hierarchical grouping to optimize an objective function. *Journal of the American Statistical Association*, 58(301):236–244, 1963. ISSN 01621459.

Wikipedia contributors. Hierarchical clustering — Wikipedia, the free encyclopedia, 2020. URL https://en.wikipedia.org/w/index.php?title=Hierarchical_clustering. [Online; accessed 2-February-2021].