

University of Rhode Island

DigitalCommons@URI

Computer Science and Statistics Faculty
Publications

Computer Science and Statistics

2019

Clustered Hierarchical Entropy-Scaling Search of Astronomical and Biological Data

Najib Ishaq

University of Rhode Island

George Student

University of Rhode Island

Noah M. Daniels

University of Rhode Island, noah_daniels@uri.edu

Follow this and additional works at: https://digitalcommons.uri.edu/cs_facpubs

Citation/Publisher Attribution

N. Ishaq, G. Student and N. M. Daniels, "Clustered Hierarchical Entropy-Scaling Search of Astronomical and Biological Data," *2019 IEEE International Conference on Big Data (Big Data)*, Los Angeles, CA, USA, 2019, pp. 780-789, doi: 10.1109/BigData47090.2019.9005688.

Available at: <http://dx.doi.org/10.1109/BigData47090.2019.9005688>

This Article is brought to you by the University of Rhode Island. It has been accepted for inclusion in Computer Science and Statistics Faculty Publications by an authorized administrator of DigitalCommons@URI. For more information, please contact digitalcommons-group@uri.edu. For permission to reuse copyrighted content, contact the author directly.

Clustered Hierarchical Entropy-Scaling Search of Astronomical and Biological Data

The University of Rhode Island Faculty have made this article openly available.
Please let us know how Open Access to this research benefits you.

This is a pre-publication author manuscript of the final, published article.

Terms of Use

This article is made available under the terms and conditions applicable towards Open Access Policy Articles, as set forth in our [Terms of Use](#).

Clustered Hierarchical Entropy-Scaling Search of Astronomical and Biological Data

Najib Ishaq
Computer Science and Statistics
University of Rhode Island
Kingston, RI 02881
najib_ishaq@zoho.com

George Student
Computer Science and Statistics
University of Rhode Island
Kingston, RI 02881

Noah M. Daniels
Computer Science and Statistics
University of Rhode Island
Kingston, RI 02881
noah_daniels@uri.edu

Abstract—Both astronomy and biology are experiencing explosive growth of data, resulting in a “big data” problem that stands in the way of a “big data” opportunity for discovery. One common question asked of such data is that of approximate search (ρ -nearest neighbors search). We present a hierarchical search algorithm for such data sets that takes advantage of particular geometric properties apparent in both astronomical and biological data sets, namely the metric entropy and fractal dimensionality of the data. We present CHES (Clustered Hierarchical Entropy-Scaling Search), a search tool with virtually no loss in specificity or sensitivity, demonstrating a $13.6\times$ speedup over linear search on the Sloan Digital Sky Survey’s APOGEE data set and a $68\times$ speedup on the GreenGenes 16S metagenomic data set, as well as asymptotically fewer distance comparisons on APOGEE when compared to the FALCONN locality-sensitive hashing library. CHES demonstrates an asymptotic complexity not directly dependent on data set size, and is in practice at least an order of magnitude faster than linear search by performing fewer distance comparisons. Unlike locality-sensitive hashing approaches, CHES can work with any user-defined distance function. CHES also allows for implicit data compression, which we demonstrate on the APOGEE data set. We also discuss an extension allowing for efficient k -nearest neighbors search.

1. Introduction

Astronomers have learned much of what they know about celestial objects by studying their electromagnetic spectra. By attaching an instrument known as a spectrograph to a telescope, astronomers are able to separate the incoming light by its wavelength, generating a fingerprint of the object being observed. We call this fingerprint a spectrum. Spectra can reveal information about characteristics of celestial objects including chemical composition, luminosity, and temperature.

The amount of astronomical data is massive and has been growing rapidly in recent years due to improved storage capacity and scientific instruments [1]. The rate of growth of this data is outpacing the computational improvements from Moore’s Law [2]. As a result, better algorithms

are required to keep up with this growth of data. This has prompted the emergence of astroinformatics, a new discipline at the intersection of astronomy and computer science [2].

The Sloan Digital Sky Survey (SDSS) [3], the first in a series of survey telescopes, is designed to catalog vast numbers of celestial objects. SDSS comprises several projects, including the Apache Point Observatory Galactic Evolution Experiment (APOGEE) [4], which by 2020 will have cataloged over 263,000 stars in the Milky Way galaxy. SDSS will be succeeded by the Large Synoptic Survey Telescope (LSST) [5] which is expected to catalog 37 billion stars and galaxies over ten years, producing 15 terabytes of data every night. Other large astronomical data sets are becoming available, such as the recent release of the National Optical Astronomy Observatory’s All-Sky Survey [6], which contains 2.9 billion unique objects.

Prior work by Yu, et al. demonstrated that a flat clustering approach could accelerate approximate search on biological and biochemical data sets [7]. As the growth of biological data has also outpaced Moore’s law [8], we also build on these results by demonstrating the applicability of CHES to a biological data set. Metagenomics is the study of entire microbial communities in environments such as seawater, soil, or the human gut. The compositional and functional analysis of the human gut microbiome is essential in the study of human health and disease [9], [10], [11]. One common approach to identifying what species or subspecies of bacteria are present is 16S ribosomal gene amplicon sequencing [12]. The ribosome, essential to all known life, comprises both protein and RNA elements; the 16th component of the small ribosomal subunit (16S) is highly conserved among bacteria, but has enough variation to make it a useful “fingerprint” to identify bacterial species [13]. The GreenGenes project [14] provides a multiple-sequence-alignment of over one million bacterial 16S sequences.

Recent approaches to dealing with the exponential growth of data have included locality-sensitive hashing [15], clever text indices such as FM Index [16], and a recent paper introducing “entropy-scaling search” for large biological data sets [7]. The key contribution of [7] is an asymptotic complexity for ρ -nearest neighbors search. This asymptotic complexity is given by:

$$O\left(\underbrace{k}_{\text{metric entropy}} + \underbrace{|B_D(q, r)|}_{\text{output size}} \underbrace{\left(\frac{r + 2r_c}{r}\right)^d}_{\text{scaling factor}}\right) \quad (1)$$

where k denotes the metric entropy of the data at cluster radius r_c (equivalently, the number of clusters), $B_D(q, r)$ denotes the number of data points in a ball of radius r around a query q (i.e. the output size of the search), and d denotes the fractal dimension in the range from query radius to cluster radius. Notably, n (the size of the database) does not appear in the complexity (though $B_D(q, r)$ is unlikely to be independent of the size and density of the data). Entropy-scaling search demonstrates particular effectiveness when data appear confined to a lower-dimensional manifold of a high-dimensional space, as hypothesized in [7] and further explored in [8].

The tasks of searching, analyzing, and storing the astronomical and biological data discussed above pose a true “big data” problem. Approximate search is useful for “guilt by association” to inexact matches that have been well studied, and is also a building block for the ubiquitous k -nearest neighbors (KNN) algorithm [17]. In this paper, we focus on the specific problem of approximate search (ρ -nearest neighbors search). Building on the concept of “entropy-scaling search” as applied to biological and biochemical data [7], we extend entropy-scaling ρ -nearest neighbors search to a hierarchical clustering approach, with similarities to binary search, called CHES (Clustered Hierarchical Entropy-Scaling Search). In doing so, we gain an expected logarithmic improvement in the coarse search, bringing the k term to $\log_2 k$. We demonstrate this approach on the Green-Genes biological data set and the Sloan Digital Sky Survey’s APOGEE astronomical data set. It should be noted CHES is applicable to any problem with a well-defined distance function. In the case of astronomical data, we compare the performance of CHES to FALCONN [18], a state-of-the-art locality sensitive hashing library. CHES also allows for lossless compression for storage and transmission of data.

Given the rise of “big data,” there has been an interest in sublinear-time algorithms [15], [19], [20]. With CHES, we demonstrate a search algorithm whose asymptotic complexity is *not a function of data set size* but rather geometric properties of the data. CHES’s advantages over naïve search will grow as data sets grow in size. We demonstrate that CHES is an order-of-magnitude (or more) faster than naïve search.

2. Methods

In this manuscript, we are primarily concerned with the problem of ρ -nearest neighbors search in a finite-dimensional (typically high-dimensional) vector space.

For our purposes, a *data point* is one datum or observation; it might represent the electromagnetic spectrum of a star, the entire genome of an organism, or any other entity on which a distance function is defined.

A *distance function* is any function $d : \mathbb{D} \times \mathbb{D} \rightarrow \mathbb{R}_+$, where \mathbb{D} is the set of data points. A distance function may also map to a subset of the nonnegative real numbers (for instance, cosine distance). A distance function that obeys the triangle inequality is also a *distance metric*. Sensible distance functions are chosen for the data: for spectral data, both Euclidean ($L2$) and cosine distance are used; for biological sequence data, Levenshtein edit distance or Hamming distance make sense [21].

We distinguish the *embedding dimension* from another concept, the discrete *fractal dimension* of the data at some particular length scale. We define *local fractal dimension* as:

$$\log_2 \left(\frac{|B_D(q, r_1)|}{|B_D(q, r_2)|} \right) \quad (2)$$

where $B_D(q, r)$ is the set of points contained in a ball on the dataset D of radius r centered on a point q ; here, fractal dimension is computed for a radius r_1 and a smaller radius $r_2 = \frac{r_1}{2}$.

For instance, the SDSS’s APOGEE data are embedded in \mathbb{R}_+^{8575} (they are nonnegative real-valued vectors of length 8,575) and they do not uniformly occupy that space. This is due to physical constraints (namely, the laws that govern stellar fusion and spectral emission lines); as famously illustrated by the Hertzsprung-Russell diagram [22] plotting luminosity against temperature. The notion of high-dimensional data constrained to a lower-dimensional manifold is known as the *manifold hypothesis* [23].

We define a *cluster* as a set of points with a *center*, a *radius*, and an approximation of the *local fractal dimension*. The *center* is the geometric median of the set of points in the cluster, and so is a real data point. The *radius* is the maximum distance from the points to the center. The *local fractal dimension* is estimated at the cluster radius and half the cluster radius. Each *cluster* (unless it is a *leaf cluster*) has two *child clusters* in much the same way as a node in a binary tree has two child nodes.

We define the *metric entropy* of a data set under a hierarchical clustering scheme as a refinement of [7], where metric entropy for a given cluster radius r_c was the number of clusters of that radius needed to cover all data. Here, we use a binary, divisive clustering approach, but with early-stopping criteria; clusters smaller than a certain threshold are not further split. Since we frame the asymptotic complexity of search in terms of the number of leaf clusters, the metric entropy is best thought of in terms of the number of leaf clusters, which is an emergent property for any given data set. Divisive clustering is chosen because we want a hierarchical “walk” of the data during search; determining an optimal clustering is not the goal.

2.1. Data Sets

2.1.1. Sloan Digital Sky Survey. The Sloan Digital Sky Survey [3] Apache Point Observatory Galaxy Evolution Experiment (APOGEE) [4] data set contains near-infrared

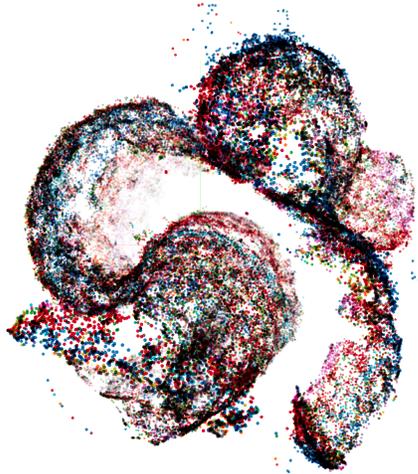


Figure 1. T-SNE visualization of the SDSS Apogee data set, using L2 norm as the distance function.

spectra of approximately 130,000 stellar objects in the Milky Way galaxy. Each spectral datum describes a star’s flux in Janskys ($1 \text{ Jy} = 10^{-26} \text{ W} \cdot \text{m}^{-2} \cdot \text{Hz}^{-1}$) at each of 8,575 wavelengths ranging from $\lambda = 1.51 \mu\text{m}$ to $1.70 \mu\text{m}$. Thus, each datum is a real-valued vector in \mathbb{R}_+^{8575} . These near-infrared spectra are used by astronomers to achieve a better understanding of how our galaxy has evolved. The APOGEE dataset used was downloaded in June 2017 from SDSS Data Release 12 [24].

Figure 1 shows a t-Stochastic Neighbor Embedding (t-SNE) [25] visualization of the APOGEE data, where colors represent leaf clusters of spectra (some colors are reused due to a limited palette). This visualization suggests support for the “manifold hypothesis.” Note that while the variation among spectra appears relatively continuous (there are few distinct, isolated clusters), there are large regions of unoccupied space in the reduced-dimension representation. In contrast, uniformly distributed data, when visualized via t-SNE, appear as a uniform cloud. This apparent constraint to a “sheetlike” manifold embedded in the high-dimensional vector space suggests that the APOGEE data is a reasonable candidate for CHES.

Figure 2 shows the mean local fractal dimension of the APOGEE data at each level of hierarchical clustering, under the L2 (Euclidean) distance metric. Each plot line represents a decile of fractal dimension; note that, other than the most extreme 10% of clusters, virtually all clusters have a local fractal dimension of less than 2. Figure 3 provides the same visualization under cosine distance and, once again, indicates that most clusters have a local fractal dimension of less than 2.

2.1.2. GreenGenes bacterial genomes. The GreenGenes project [14] provides a set of 1,027,383 bacterial 16S sequences. These sequences contain several duplicates so 805,434 unique sequence user used in clustering. Given a metagenomic sample, a bioinformatician might search

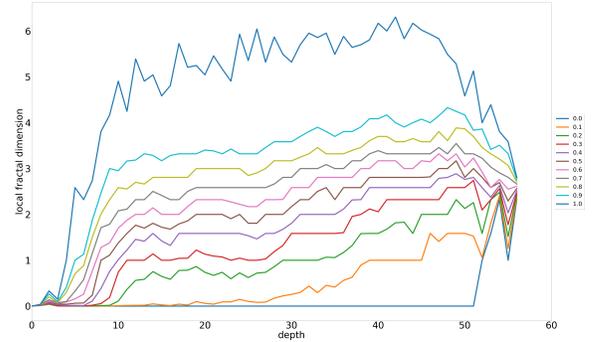


Figure 2. Mean fractal dimension of APOGEE clusters as a function of depth when clustered based on L2 norm. Each plotline represents a distinct decile of fractal dimension. Beyond a depth of 56, no clusters are further divided due to the minimum-cluster-cardinality stopping criteria.

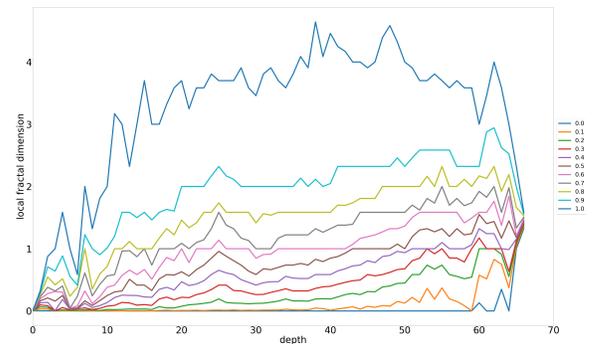


Figure 3. Mean fractal dimension of APOGEE clusters as a function of depth when clustered based on cosine distance. Each plotline represents a distinct decile of fractal dimension. Beyond a depth of 66, no clusters are further divided due to the minimum-cluster-cardinality stopping criteria.

those million-plus sequences for close matches to the 16S sequences identified in a sample of interest. A typical distance function here would be edit (Levenshtein) distance or Hamming distance, which roughly models the insertions, deletions, and substitutions that happen during the course of evolution [21]. These sequences are already aligned and, thus, form sequences of $\{A, C, G, T, -\}$ (where $-$ represents a gap corresponding to an insertion or deletion during the course of evolution). Thanks to the alignment, every sequence in the set is 7,682 characters long, so Hamming distance is an appropriate metric.

During the course of evolution, organisms adapt to their environment by random mutation and the retention of (or selection for) those mutations that are beneficial. In [7], it was hypothesized that this process could result in the low fractal dimension and low metric entropy for which entropy-scaling search is suited. Figure 4 shows a t-SNE visualization of the GreenGenes dataset under Hamming distance, where colors indicate the genus of a bacterium (some colors are reused due to a limited palette). Clearly, these data appear fundamentally different from the APOGEE data. There are distinct clusters of species separate from one another, which tend to be grouped by genus, and the continuity seen in the

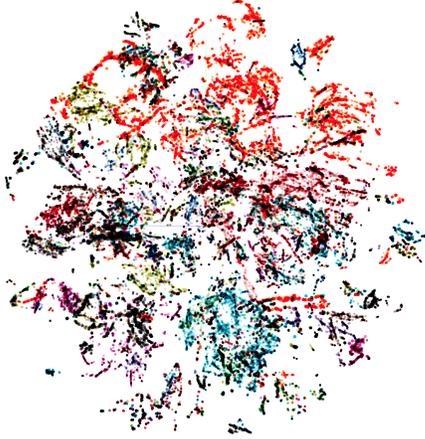


Figure 4. T-SNE visualization of the GreenGenes 16S RNA data set, using Hamming distance.

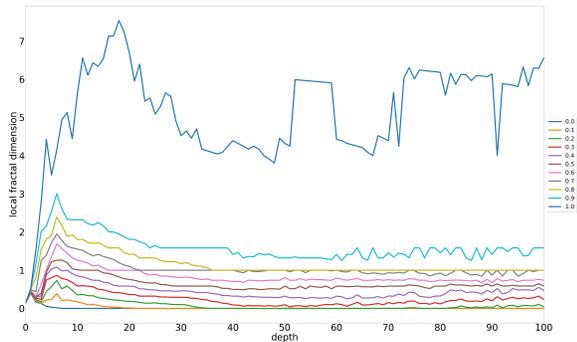


Figure 5. Mean fractal dimension of GreenGenes clusters as a function of depth when clustered based on Hamming distance. Each plotline represents a distinct decile of fractal dimension.

APOGEE data is not present. Like the APOGEE data, entire regions of the embedding space are unoccupied, suggesting that CHESS is well-suited to the GreenGenes dataset.

Figure 5 shows the local fractal dimension of the GreenGenes data at each level of hierarchical clustering under Hamming distance. Each plot line represents a decile of fractal dimension; once again, other than the most extreme 10% of clusters, virtually all clusters have a local fractal dimension of less than 2.

For this study, we used the October 2012 release of Greengenes, downloaded in April 2019.

2.2. Hierarchical clustering and search

2.2.1. Clustering. Hierarchical clustering lends itself to a binary search on clusters, followed by one or more “fine” (linear) searches inside clusters. As such, the k term in the asymptotic complexity from flat-clustered search becomes $\log_2 k$. However, the fine search term is more difficult to quantify, as cluster radius is no longer an input parameter, but instead an emergent property. Unlike classical binary search, there are cases (such as the “ball” around a

query crossing into multiple child clusters) where multiple branches of the binary tree must be searched. However, under an assumption of low fractal dimension and low metric entropy, where much of the high-dimensional space is unoccupied by data, traversing multiple branches of the tree should be rare. This is demonstrated in Section 3.

Hierarchical clustering follows the procedure outlined in Algorithm 1.

Algorithm 1 Cluster(data)

Require: $maxdepth > 0$

Require: $minsize > 0$

$d \leftarrow 0$

$clusters = \{\}$

$n = |data|$

while $d \leq maxdepth$ **do**

$numseeds \leftarrow \sqrt{n}$

$seeds \leftarrow numseeds$ chosen randomly from $data$

$l, r \leftarrow \{l, r \mid l, r \in seeds \wedge l, r = \arg \max d(x, y) \mid x, y \in seeds\}$

$clusters[l] \leftarrow \{x \mid x \in data \wedge d(l, x) \leq d(r, x)\}$

$clusters[r] \leftarrow \{x \mid x \in data \wedge d(r, x) < d(l, x)\}$

if $|clusters[l]| > minsize$ **then**

 Cluster($clusters[l]$)

end if

if $|clusters[r]| > minsize$ **then**

 Cluster($clusters[r]$)

end if

end while

Hierarchical clustering is essentially a recursive procedure that relies on some distance function $d(x, y)$. In order to cluster n data points, \sqrt{n} are chosen at random as possible cluster centers (seeds). All pairwise distances are computed among these \sqrt{n} seeds, and the two furthest seeds are chosen as the “left” and “right” cluster centers. All n data points are then partitioned into “left” and “right” “child” clusters by whichever seed they are closer to. This procedure recurses until a user-specified maximum depth is reached. Any cluster containing fewer than a user-specified minimum count (typically 10) is not recursively clustered further.

2.2.2. Search. Rather than an exact search (which may not return any hits), we conduct an approximate (ρ -nearest neighbor) search. Given a hierarchical clustering, the search procedure begins much like classical binary search. During this search, it is possible that both children of a cluster contain valid hits. Thus, a search radius of $r + r_{c_i}$ (the query radius plus the cluster radius for a cluster c_i) is used to determine whether a subtree might contain valid hits. Search proceeds recursively down the binary tree of clusters, exploring any child of a cluster if that child’s cluster center is within $r + r_c$ of the query, where r_c is the radius of that child. Once leaves are reached, the union of all leaves identified as containing candidate hits are searched exhaustively. The procedure is outlined in Algorithm 2.

Algorithm 2 Search($q,r,clusters$)

Require: $r \geq 0$
Require: $clusters \neq \emptyset$
 $results \leftarrow \{\}$
if $clusters.left$ **then**
 if $d(q, clusters.l.center) \leq r + clusters.l.radius$
 then
 Search($q,r,clusters.l$)
 end if
end if
if $clusters.right$ **then**
 if $d(q, clusters.rt.center) \leq r + clusters.rt.radius$
 then
 Search($q,r,clusters.rt$)
 end if
end if
if $\neg clusters.l \wedge \neg clusters.rt$ **then**
 for $p \in clusters.points$ **do**
 if $d(q, p) \leq r$ **then**
 $results \leftarrow r$
 end if
 end for
end if
return $results$

2.3. Data compression

Each leaf cluster contains a cluster center (a real data point, rather than a centroid). Therefore, each member of that cluster can be represented as the set of differences from that center. The optimal representation of these differences depends on the nature of the data as well as the distance function. In the case of string data (such as biological sequence data) under a Hamming distance metric, the representation is naturally a set of edits to transform the cluster center into any data point in the cluster. *The number of these edits is bounded by the cluster radius.* In the case of astronomical spectra, the data are real-valued vectors in (\mathbb{R}_+^{8575}) and so a compact encoding appears elusive. This is because the upper bound difference of one cluster radius from the center could be distributed in any manner among the 8,575 bins.

However, the underlying APOGEE data is not of arbitrary precision. Due to the physical limitations inherent in measurement, these values are quantized [26] and this makes an integer encoding of differences possible. Consider, by way of analogy, a meter rule that is marked in increments of $1mm$. Any length measured by the meter rule is measured as some integer multiple of $1mm$ even though the actual length is some real number. Any difference of less than $1mm$ between two lengths cannot be measured and is ignored. Lengths measured by such a meter rule can, thus, be thought to be quantized by $1mm$ sized quanta. The telescope and spectrometer used for APOGEE have an analogous limitation. APOGEE reports that it can detect stars of magnitude $H = 12.2$ [4] (where $H = -2.5 \log_{10}(I)$ and I is the intensity of the star). This leads to quanta of

size $10^{-12.2/2.5}$.

For lossless compression of clustered data, we calculate the difference between each point in a cluster and the cluster center. This difference is then quantized to give an integer encoding. These encodings can then be compressed using any suitable compression library. We used the zipfile library in Python for our implementation.

2.4. Benchmarks

2.4.1. Sloan Digital Sky Survey APOGEE. Fifty randomly-selected data points were held out from the data prior to hierarchical clustering. Each of those points was treated as a query, and the data set was searched to return all spectra within a specific radius of the query. For L2 (Euclidean) distance, these radii were 2,000 and 4,000. We measured the wall-clock time for each query, as well as the number of distance comparisons required to perform the search. We report the mean and standard deviation for distance comparisons, search time (in seconds), the fraction of the dataset that was searched, and the mean speedup factor compared to a naïve, linear search implementation. This linear search must make n comparisons given a database of size n .

All benchmarks were conducted on an Intel Xeon E5-2690 v4 2.60GHz (single threaded), with 512GB RAM and CentOS 7 Linux with kernel 3.10.0-862.9.1.el7.x86_64. FALCONN is not versioned, but commit `a4c0288edb2575b0306c5f6b4ab1f559b1e2a` from GitHub was used.

FALCONN is written in C++ (and heavily optimized for CPUs) while CHESS is written in Python. As such, wall-clock time comparisons are not truly fair; a GPU-optimized or multithreaded FALCONN would likely provide improved performance, while CHESS rewritten in C++ would likely boast improved all-round performance. To remedy this inconsistency, we report the *number of distance comparisons made* by the two methods while conducting the same search on the same dataset. We used Valgrind’s `callgrind` tool [27] to produce a performance profile of FALCONN’s underlying C++ library, and we used internal source-code instrumentation to profile CHESS.

2.4.2. GreenGenes. Fifty randomly-selected data points were held out from the data prior to hierarchical clustering. Each of those points was treated as a query, and the data set was searched to return all 16S sequences within a specific radius of the query. We used Hamming distances corresponding to 99.9% and 99% sequence identity (note that at 95% sequence identity a query returns, on average, nearly 75% of the data set, while at 90% sequence identity a query would return almost the entire dataset, so we expect no improvement over naïve search). We measured the wall-clock time for each query, as well as the number of distance comparisons required to perform the search. We report the mean and standard deviation for comparisons, search time (in seconds), the fraction of the dataset that was searched,

and the mean speedup factor compared to a naïve, linear search implementation.

The GreenGenes benchmarks for CHESS were conducted on a Xeon E5-2620 v4 2.10GHz with 64GB RAM, hosted on Ubuntu Linux with kernel 4.15.0-54-generic, Python 3.6.8.

FALCONN only supports L2 and cosine distances, not Hamming distance, so it is not directly comparable on the GreenGenes data set. Thus, we did not benchmark FALCONN on GreenGenes.

3. Results

3.1. Asymptotic complexity

To analyze the computational complexity of hierarchical clustering, consider the top level of the hierarchy. There are n data points, from which \sqrt{n} seeds are chosen. All pairwise distances between these seeds are computed, i.e. n comparisons are made. The furthest pair of points are chosen as the centers of child-clusters. Every data point except the two chosen centers is then compared to both centers; i.e., an additional $2n$ comparisons. In general, at depth d of the hierarchy, we have $\frac{n}{2^d}$ points per cluster (assuming a perfectly balanced clustering, which will turn out to not be required), $\sqrt{\frac{n}{2^d}}$ seeds, $\frac{n}{2^d}$ pairwise seed comparisons, and $2 \times \frac{n}{2^d}$ additional comparisons when each other data point is compared to the two cluster centers, for a total of $3 \times \frac{n}{2^d}$ distance comparisons. At each level, there are 2^d clusters, so the total number of distance comparisons per level is $3n$. Thus, the asymptotic complexity in terms of distance comparisons for a cluster tree of depth d is $O(dn)$.

Since hierarchical entropy-scaling search begins as a search over a binary tree, it would appear that the coarse search (the tree traversal) would take $O(\log_2 k)$ time, if there are k leaf clusters. However, there is a complication: since this is approximate rather than exact search, a nonzero query radius, r , is used. The “ball” of possible points around the query may overlap multiple branches of the tree. While it is clear that some of the search space is pruned during the tree traversal, there is no guarantee that the fine search will be over only one leaf cluster.

Consider a clustering using L2 norm. Suppose that the data comprise a one-dimensional manifold in a two-dimensional embedding. Then, every level of the binary tree represents some line dividing the parent cluster into two (not necessarily equally-sized) clusters. If a parent cluster is not identified as a target for search, neither of its children will be. However, it is possible that *both* children of a given parent cluster could be considered. Once cluster radius has decreased to be smaller than the search radius, there is no advantage to further tree traversal, and the entire subtree must be searched. However, above the level of the tree where search radius approximates cluster radius, at most both children of an internal cluster node need to be searched; even in this instance, only half of their descendents would need to be searched, because of the two-dimensional nature of the data

set. Thus, the coarse search (in one dimension) requires only a constant factor of 2; $O(2 \log_2 k) = O(\log_2 k)$. However, as dimensionality increases the argument requires using the geometry of an n -sphere.

Recall that our clustering approach chooses (or estimates) the two furthest points in a cluster to determine the cluster centers at the next level of the hierarchy. First, consider data in one dimension. The very first cut into two clusters results in a boundary on which a query with nonzero radius might hit results from both clusters. However, any further division of those clusters separates the data into non-adjacent regions, so continued binary clustering and then search will result in a pruning of the search space. Next, consider data in two dimensions. Without loss of generality, consider those data uniformly occupying the interior of a circle around some origin. A first cut will separate the data into two approximate semicircles. But, the furthest points on each semicircle will likely be 180 degrees apart, along the diameter of the cut. Thus, a second cut will separate each semicircle into a quarter-circle, which contains a right triangle. Thus, the furthest points in each quarter-circle could still be at the two intersections of the triangle with the circumference of the circle. It is only after this point that, under a uniform distribution of data, the third cut will separate each quarter-circle into an eighth-circle, leading the fourth cut to be orthogonal to a radius, separating “outer” points from “inner” points and guaranteeing that no query with radius smaller than the cluster radius would require searching inside every cluster.

Taken to three dimensions, orthogonal cuts are only guaranteed once the wedges of the sphere are less than 60 degrees in both ϕ and θ , which requires six cuts. Generalizing to the n -sphere, described in polar coordinates as containing angles $\phi_1 \cdots \phi_n$, the first $3 \times (d - 1)$ cuts are not guaranteed to be orthogonal to a radius. If our data were uniformly distributed in a high-dimensional space, this would bode poorly for CHESS.

Of course, with the APOGEE data set, we have not two dimensions but 8,575. By the curse of dimensionality, we might need to explore many thousands of “neighboring” tree branches. How can this explosion be reconciled with the empirically faster search as tree depth increases? There is no such explosion because the number of neighboring branches that may need to be searched grows not with the *embedding* dimension (i.e. the dimensionality of each data point, 8,575 for APOGEE and 7,682 for GreenGenes) but with the *local fractal dimension* on the length scale of the cluster and search radii. As in [7], we use Equation 2 to estimate the local fractal dimension between two radii at a given point. As illustrated in Figures 2, 3, and 5, this fractal dimension is low, typically less than 2, except for the most extreme decile of clusters. Thus, *empirically* for both the APOGEE and GreenGenes data sets, the multiplier is bounded by 4, which means that once clustering exceeds a depth of 3, we can claim an effective asymptotic complexity of:

TABLE 1. CHESS PERFORMANCE ON APOGEE vs. NAÏVE, L2 NORM

Depth	Comparisons $\times 10^4$		Search Time (s)		Fraction Searched		Speedup Factor
	μ	σ	μ	σ	μ	σ	
r = 2,000			Output Size: $\mu = 482 \sigma = 1,004$				
10	7.79	2.86	5.03	1.82	0.63	0.23	3.42
20	3.30	1.88	2.38	1.22	0.27	0.15	9.20
30	1.44	1.20	1.57	0.94	0.12	0.10	12.85
40	0.73	0.69	1.40	0.88	0.06	0.06	13.52
50	0.68	0.63	1.38	0.87	0.05	0.05	13.55
r = 4,000			Output Size: $\mu = 4,563 \sigma = 5,803$				
10	7.87	2.86	5.04	1.81	0.64	0.23	3.37
20	3.55	1.99	2.57	1.30	0.29	0.16	9.11
30	1.97	1.47	1.96	1.04	0.16	0.12	9.77
40	1.54	1.22	1.84	0.99	0.12	0.10	9.86
50	1.54	1.21	1.84	0.98	0.12	0.10	9.85

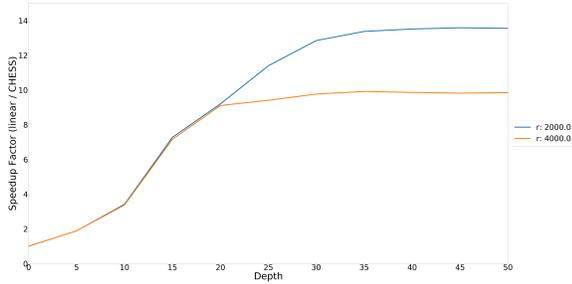


Figure 6. Speedup Factor of CHESS on APOGEE vs. Naïve Search Under L2 Norm as a Function of Search Depth

$$O\left(\underbrace{\log_2 k}_{\text{metric entropy}} + \underbrace{|B_D(q, r)|}_{\text{output size}} \underbrace{\left(\frac{r + 2\hat{r}_c}{r}\right)^d}_{\text{scaling factor}}\right) \quad (3)$$

where \hat{r}_c is the *mean* cluster radius of leaf clusters.

Similar to the original entropy-scaling search [7], this approach is ill-suited to data that are not constrained to low-dimensional manifolds in their higher-dimensional embedding. On the other hand, when data exhibit such constraints, entropy-scaling search is appropriate.

3.2. Benchmark Results

On the SDSS’s APOGEE data set, using L2 norm, there is a clear reduction in search time as the maximum tree depth increases. Table 1 shows the number of comparisons, search time in seconds, what fraction of the data needed to be searched, and the speedup factor vs. naïve linear search for L2 norm. Figure 6 illustrates the speedup as a function of depth.

On the APOGEE data set using cosine distance, there is no clear reduction in search time until CHESS begins to significantly prune the search space, which requires a much

TABLE 2. CHESS PERFORMANCE ON APOGEE vs. NAÏVE, COSINE DISTANCE

Depth	Comparisons $\times 10^4$		Search Time (s)		Fraction Searched		Speedup Factor
	μ	σ	μ	σ	μ	σ	
r = 0.0005			Output Size: $\mu = 3 \sigma = 8$				
10	12.24	0.02	7.75	0.09	0.99	0.00	1.00
20	10.82	1.08	7.58	0.73	0.88	0.09	1.04
30	4.19	2.11	4.93	1.97	0.34	0.17	2.06
40	0.88	0.72	3.85	1.92	0.07	0.06	2.76
50	0.62	0.49	3.83	1.93	0.05	0.04	2.78
r = 0.001			Output Size: $\mu = 73 \sigma = 163$				
10	12.24	0.02	7.78	0.09	0.99	0.00	0.99
20	10.85	1.08	7.63	0.76	0.88	0.09	1.03
30	4.40	2.21	5.19	2.14	0.36	0.18	2.03
40	1.07	0.90	4.30	2.34	0.09	0.07	2.64
50	0.78	0.66	4.34	2.43	0.06	0.05	2.64

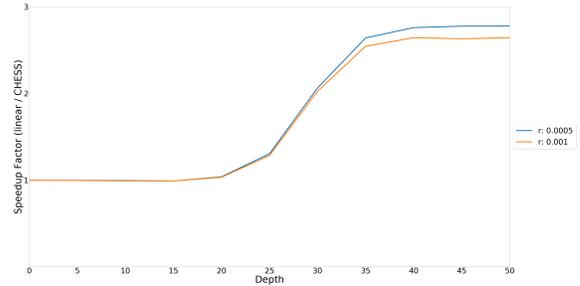


Figure 7. Speedup Factor of CHESS on APOGEE vs. Naïve Search Under Cosine Distance as a Function of Search Depth

higher depth than the L2 norm. Table 2 shows benchmark results under cosine distance at depths of 10 through 50 with increments of 10. At a depth of 30, CHESS begins to show moderate speedup, leveling off by depth 45. The greater depth requirement appears to be a result of the fact that cosine distance cannot make *orthogonal* cuts to the space the way L2 norm or Hamming distance can. While on a circle, with angle θ , a query with small radius can hit at most two clusters. On a 3-sphere, with angles ϕ and θ , a query could hit $2^{d-1} = 4$ clusters. Thus, it takes greater cluster depth to begin to see significant pruning of the search space.

On the GreenGenes data set, using Hamming distance, there is a clear reduction in search time as long as the search radius is small. Table 3 shows the number of comparisons, search time in seconds, the fraction of leaf clusters that must be searched, and the speedup factor vs. naïve linear search. At 99.9% and 99% sequence identity, speedup factors of 68.02 and 18.39 are observed. Figure 8 illustrates that, for small search radii, CHESS search performance improves with depth. These small radii are biologically meaningful, as every sequence represents the same functional and structural component of the ribosome and, thus, almost the entire data set exists within a diameter of 90% sequence identity.

FALCONN’s benchmark results on APOGEE appear in

TABLE 3. CHESS PERFORMANCE ON GREENGENES VS. NAÏVE, HAMMING DISTANCE

Depth	Comparisons $\times 10^5$		Search Time (s)		Fraction Searched	Speedup Factor
	μ	σ	μ	σ	μ	μ
99.9% Seq. Identity			Output Size: $\mu = 245$ $\sigma = 486$			
10	46.84	18.58	26.86	10.67	0.582	3.02
20	3.39	4.24	2.88	2.90	0.042	57.54
30	0.32	0.34	1.40	1.02	0.004	66.86
40	0.22	0.19	1.37	0.98	0.003	67.51
50	0.21	0.17	1.36	0.99	0.003	68.02
99% Seq. Identity			Output Size: $\mu = 1,924$ $\sigma = 2,364$			
10	59.09	14.46	33.86	8.21	0.734	1.45
20	12.59	9.05	9.82	6.45	0.156	12.57
30	2.79	2.82	5.77	4.22	0.035	18.20
40	2.66	2.69	5.58	3.98	0.033	18.34
50	2.66	2.69	5.57	3.98	0.033	18.39

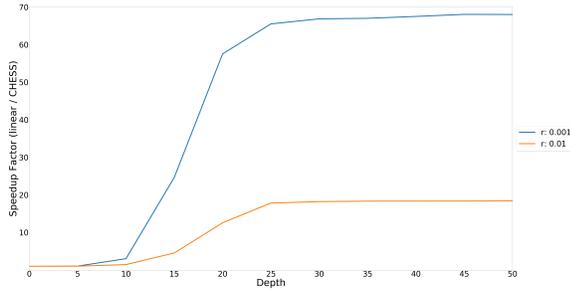


Figure 8. Speedup Factor of CHESS on GreenGenes vs. Naïve Search Under Hamming Distance as a Function of Search Depth.

TABLE 4. FALCONN SEARCH PERFORMANCE ON APOGEE

Distance	Radius	Comparisons $\times 10^5$	Search Time (s)	
			μ	σ
L2	2000	1.32	0.563	0.008
	4000	1.32	0.564	0.008
Cosine	0.0005	1.32	0.520	0.014
	0.001	1.32	0.517	0.011

Table 4. While FALCONN is uniformly faster in terms of wall-clock time than CHESS, we note that the number of comparisons required by FALCONN is significantly greater than that required by CHESS. In particular, at a radius of 2,000 and a depth of 50, CHESS under L2 Norm requires, on average, only $6.8 \cdot 10^3$ distance comparisons; and at a radius of 0.0005 and a depth of 50, CHESS under Cosine Distance requires, on average, only $6.4 \cdot 10^3$ distance comparisons. FALCONN requires $1.32 \cdot 10^5$ comparisons in all cases, showing that CHESS provides an order of magnitude algorithmic improvement over FALCONN. In all cases, naïve search, CHESS, and FALCONN agreed on all search results, so we can consider all searches to be exact.

3.3. Accuracy

Similar to the argument presented in [28], the triangle inequality guarantees perfect accuracy as long as the distance function used is a metric. To search for all points within radius r of a query, at each node we only need to look inside child-clusters whose centers are within radius $r + r_i$ of a query, where r_i is the cluster radius of a child of that node. False positives as compared against naïve search are not possible with CHESS. Under Hamming and Euclidean distances, which both obey the triangle inequality, the false negative rate is zero. Under cosine distance, which does not obey the triangle inequality, the false negative rate averages 4.4×10^{-4} and is only nonzero at clustering depths greater than 30.

3.4. Compression

Our implementation of data compression as described in Section 2.3 results in shrinking the APOGEE data set from 4 GB of storage (already represented as a NumPy memmap, a compact binary file) to 3.4 GB. With a larger data set and thus greater redundancy, we would expect greater compression. It may be worthwhile future work to investigate representing differences at each level of the binary tree, particularly for string or genomic data, where all differences are discrete.

3.5. k-Nearest Neighbors Search Extension

Machine learning and data science practitioners consider k-nearest neighbors (KNN) a nearly ubiquitous classification tool [29], [30], [31], [32]. While in this paper we have focused on the related ρ -nearest neighbors problem, it is worth discussing an extension to KNN. Under clustered search, KNN can be implemented as follows:

Perform CHESS (Algorithm 2) with $\rho = \hat{r}_c$.

If $|B_D(q, \rho)| > k$ and more than one leaf cluster was searched:

Repeat CHESS, halving ρ until $|B_D(q, \rho)| \leq k$.
Double ρ once.

If $|B_D(q, \rho)| < k$ and $\rho \leq r_{c_0}$:

Repeat CHESS, doubling ρ until $|B_D(q, \rho)| \geq k$.

Return k-nearest neighbors from $B_D(q, \rho)$.

where \hat{r}_c is the median cluster radius of leaf clusters, and r_{c_0} is the radius of the root cluster.

The loops inside both conditionals are bound by $\log(r_{c_0})$ calls to the search function. The return statement has an asymptotic complexity of $O(|B_D(q, \rho)| \cdot \log k)$ if a k-element min-heap is used to filter the results. Therefore the complexity of KNN search is:

$$O(S \cdot \log(r') + |B_D(q, \rho)| \cdot \log k) \quad (4)$$

where S is the asymptotic complexity of CHESS given in Equation 3 and the radius r' is r_{c_0} when r_{c_0} includes values greater than 1, and otherwise $\frac{1}{r_{c_0}}$.

4. Conclusion and Discussion

We have presented CHES (Clustered Hierarchical Entropy-Scaling Search), an algorithm and software implementation for ρ -nearest neighbors approximate search of large data sets, which improves the coarse search time of entropy-scaling search [7] from k to $\log_2 k$, where k is the number of clusters. Like [7], CHES's asymptotic complexity is not defined in terms of n (the size of the data set being searched). However, the $B_D(q, r)$ term (which is equivalent to the output size) has an *implicit* dependence on the size of the data set; as more data are collected, particularly if they are similar to existing data, they are expected to increase the density of each cluster. This can be addressed by further deepening the hierarchical clustering, but this also creates an implicit dependence on n (as n grows, so k should grow, but only as $\log(n)$).

CHES is most effective when data exhibit low metric entropy and fractal dimension. CHES requires an order of magnitude fewer comparisons than the FALCONN locality-sensitive hashing library, and provides data compression. While the wall-clock benchmarks slightly favor FALCONN, it is important to note that CHES's Python implementation imposes some overhead, despite using the `sklearn` library for distance computations. In contrast, FALCONN relies on the highly-optimized `Eigen` linear-algebra library, which uses AVX intrinsics to achieve vectorization on the CPU. In addition, CHES's I/O routines are in Python, which also adds overhead. A native implementation of CHES in C++ or Rust will be worthwhile future engineering work. Unlike locality-sensitive hashing approaches, CHES is extensible to any user-defined edit distance or similarity function. Examples include Jaccard similarity for sets, Wasserstein (Earth-Mover) distance for distributions, a variety of inter-graph distances for graphs, and others. As long as the distance function obeys the triangle inequality, search is guaranteed to be exact; non-metric distance functions such as cosine distance or BLAST E-values [33] can lead to false negatives, though not false positives. This is further discussed in section 3.3.

CHES did not begin to provide benefits on APOGEE with cosine distance until significantly higher cluster depths compared to L2 norm. This is at odds with the flat entropy-scaling search results from [7], though those were from a protein structure data set. Further investigation into cosine distance on APOGEE is warranted, but so is the investigation of other distance functions, such as Wasserstein or Earth Mover distance.

We have also discussed the theoretical properties of implementing k -nearest neighbor search using CHES (as well as flat entropy-scaling search). The implementation is a topic for future work.

In our comparison to FALCONN [18], it is worth noting that CHES can be thought of as a form of locality-sensitive hashing [34], [35], though one not based on the choice of random separating hyperplanes. Similar objects hash, or cluster, into like clusters, but (other than the possibility of near neighbors appearing in neighboring clusters) it lacks

the probabilistic nature of classical locality-sensitive hashing approaches. It is also reasonable to think of CHES as a form of database index; if one wished to dispense with the data compression aspect, the cluster hierarchy itself could be thought of as an index. Indeed, we do not foresee any serious theoretical obstacles to incorporating such an index into a relational or key-value database such as PostgreSQL or MongoDB.

CHES can easily be extended to use GPU-accelerated distance calculations. We provide some distance functions implemented in tensorflow. We did not use them for benchmarks because the overhead of sending data to the GPU is far greater than the computational cost of the distance functions used in this manuscript. For more costly distance functions, for example maximal common subgraph, GPU implementations may be worthwhile.

CHES can easily be adapted to allow for live updates to the data. As new data points become available, we can perform a binary search with zero search radius to identify the leaf cluster to which the new point would best belong. If a point is added that is far outside a leaf cluster (for example, twice the cluster radius or more), a new cluster is created.

The original entropy-scaling search manuscript [7] led to a thoughtful discussion of the fundamental limits of search [36]. In this study, we have suggested a further lowering of these limits to be logarithmic rather than linear in the number of leaf clusters.

The source code for CHES is available under an MIT license at <https://github.com/nishaq503/CHES>.

Acknowledgments

The authors would like to thank Bonnie Berger, Michael Baym, and Y. William Yu for suggesting the extension of entropy-scaling search to a hierarchical paradigm, and Tom Howard and Matthew Daily for helpful discussions.

References

- [1] Y. Zhang and Y. Zhao, "Astronomy in the big data era," *Data Science Journal*, vol. 14, 2015.
- [2] M. Brescia, S. Cavuoti, G. S. Djorgovski, C. Donalek, G. Longo, and M. Paolillo, "Extracting knowledge from massive astronomical data sets," in *Astrostatistics and Data Mining*. Springer, 2012, pp. 31–45.
- [3] M. R. Blanton, M. A. Bershady, B. Abolfathi, F. D. Albareti, C. Allende Prieto, A. Almeida, J. Alonso-García, F. Anders, S. F. Anderson, B. Andrews, and et al., "Sloan Digital Sky Survey IV: Mapping the Milky Way, Nearby Galaxies, and the Distant Universe," *The Astronomical Journal*, vol. 154, p. 28, Jul. 2017.
- [4] S. R. Majewski, R. P. Schiavon, P. M. Frinchaboy, C. Allende Prieto, R. Barkhouser, D. Bizyaev, B. Blank, S. Brunner, A. Burton, R. Carrera, S. D. Chojnowski, K. Cunha, C. Epstein, G. Fitzgerald, A. E. García Pérez, F. R. Hearty, C. Henderson, J. A. Holtzman, J. A. Johnson, C. R. Lam, J. E. Lawler, P. Maseman, S. Mészáros, M. Nelson, D. C. Nguyen, D. L. Nidever, M. Pinsonneault, M. Shetrone, S. Smee, V. V. Smith, T. Stolberg, M. F. Skrutskie, E. Walker, J. C. Wilson, G. Zasowski, F. Anders, S. Basu, S. Beland, M. R. Blanton, J. Bovy, J. R. Brownstein, J. Carlberg, W. Chaplin, C. Chiappini, D. J. Eisenstein, Y. Elsworth, D. Feuillet, S. W. Fleming, J. Galbraith-Frew, R. A.

- García, D. A. García-Hernández, B. A. Gillespie, L. Girardi, J. E. Gunn, S. Hasselquist, M. R. Hayden, S. Hekker, I. Ivans, K. Kinemuchi, M. Klaene, S. Mahadevan, S. Mathur, B. Mosser, D. Muna, J. A. Munn, R. C. Nichol, R. W. O’Connell, J. K. Parejko, A. C. Robin, H. Rocha-Pinto, M. Schultheis, A. M. Serenelli, N. Shane, V. Silva Aguirre, J. S. Sobeck, B. Thompson, N. W. Troup, D. H. Weinberg, and O. Zamora, “The Apache Point Observatory Galactic Evolution Experiment (APOGEE),” *The Astronomical Journal*, vol. 154, p. 94, Sep. 2017.
- [5] Ž. Ivezić, S. M. Kahn, J. A. Tyson, B. Abel, E. Acosta, R. Allsman, D. Alonso, Y. AlSayyad, S. F. Anderson, J. Andrew *et al.*, “Lsst: from science drivers to reference design and anticipated data products,” *The Astrophysical Journal*, vol. 873, no. 2, p. 111, 2019.
- [6] D. L. Nidever, A. Dey, K. Olsen, S. Ridgway, R. Nikutta, S. Juneau, M. Fitzpatrick, A. Scott, and F. Valdes, “First data release of the all-sky noao source catalog,” *The Astronomical Journal*, vol. 156, no. 3, p. 131, 2018.
- [7] Y. W. Yu, N. Daniels, D. C. Danko, and B. Berger, “Entropy-scaling search of massive biological data,” *Cell Systems*, vol. 1, no. 2, pp. 130–140, 2015.
- [8] B. Berger, N. M. Daniels, and Y. W. Yu, “Computational biology in the 21st century: Scaling with compressive algorithms,” *Communications of the ACM*, vol. 59, no. 8, p. 72, 2016.
- [9] S. R. Gill, M. Pop, R. T. DeBoy, P. B. Eckburg, P. J. Turnbaugh, B. S. Samuel, J. I. Gordon, D. A. Relman, C. M. Fraser-Liggett, and K. E. Nelson, “Metagenomic analysis of the human distal gut microbiome,” *science*, vol. 312, no. 5778, pp. 1355–1359, 2006.
- [10] M. Arumugam, J. Raes, E. Pelletier, D. Le Paslier, T. Yamada, D. R. Mende, G. R. Fernandes, J. Tap, T. Bruls, J.-M. Batto *et al.*, “Enterotypes of the human gut microbiome,” *nature*, vol. 473, no. 7346, p. 174, 2011.
- [11] T. Yatsunenko, F. E. Rey, M. J. Manary, I. Trehan, M. G. Dominguez-Bello, M. Contreras, M. Magris, G. Hidalgo, R. N. Baldassano, A. P. Anokhin *et al.*, “Human gut microbiome viewed across age and geography,” *nature*, vol. 486, no. 7402, p. 222, 2012.
- [12] J. Jovel, J. Patterson, W. Wang, N. Hotte, S. O’Keefe, T. Mitchell, T. Perry, D. Kao, A. L. Mason, K. L. Madsen *et al.*, “Characterization of the gut microbiome using 16s or shotgun metagenomics,” *Frontiers in microbiology*, vol. 7, p. 459, 2016.
- [13] M. G. Langille, J. Zaneveld, J. G. Caporaso, D. McDonald, D. Knights, J. A. Reyes, J. C. Clemente, D. E. Burkpile, R. L. V. Thurber, R. Knight, and C. Huttenhower, “Predictive functional profiling of microbial communities using 16S rRNA marker gene sequences,” *Nature Biotechnology*, vol. 31, no. 9, pp. 814–821, 2013.
- [14] T. Z. DeSantis, P. Hugenholtz, N. Larsen, M. Rojas, E. L. Brodie, K. Keller, T. Huber, D. Dalevi, P. Hu, and G. L. Andersen, “GreenGenes, a chimera-checked 16s rRNA gene database and workbench compatible with arb,” *Appl. Environ. Microbiol.*, vol. 72, no. 7, pp. 5069–5072, 2006.
- [15] P. Indyk, “Sublinear time algorithms for metric space problems,” *Symposium on Theory of Computing*, 1999.
- [16] J. T. Simpson and R. Durbin, “Efficient construction of an assembly string graph using the fm-index,” *Bioinformatics*, vol. 26, no. 12, pp. i367–i373, 2010.
- [17] T. M. Cover, P. Hart *et al.*, “Nearest neighbor pattern classification,” *IEEE transactions on information theory*, vol. 13, no. 1, pp. 21–27, 1967.
- [18] I. Razenshteyn, L. Schmidt, A. Andoni, P. Indyk, and T. Laarhoven, “Falconn: Similarity search over high-dimensional data,” 2015.
- [19] S. Bhattacharya, M. Henzinger, D. Nanongkai, and C. Tsourakakis, “Space-and time-efficient algorithm for maintaining dense subgraphs on one-pass dynamic streams,” in *Proceedings of the forty-seventh annual ACM symposium on Theory of computing*. ACM, 2015, pp. 173–182.
- [20] D. M. Kane, J. Nelson, and D. P. Woodruff, “On the exact space complexity of sketching and streaming small norms,” in *Proceedings of the twenty-first annual ACM-SIAM symposium on Discrete Algorithms*. Society for Industrial and Applied Mathematics, 2010, pp. 1161–1178.
- [21] S. B. Needleman and C. D. Wunsch, “A general method applicable to the search for similarities in the amino acid sequence of two proteins,” *Journal of molecular biology*, vol. 48, no. 3, pp. 443–453, 1970.
- [22] H. Rosenberg, “Über den zusammenhang von helligkeit und spektral-typus in den plejaden,” *Astronomische Nachrichten*, vol. 186, p. 71, 1910.
- [23] C. Fefferman, S. Mitter, and H. Narayanan, “Testing the manifold hypothesis,” *Journal of the American Mathematical Society*, vol. 29, no. 4, pp. 983–1049, 2016.
- [24] S. Alam, F. D. Albareti, C. A. Prieto, F. Anders, S. F. Anderson, T. Anderton, B. H. Andrews, E. Armengaud, É. Aubourg, S. Bailey *et al.*, “The eleventh and twelfth data releases of the sloan digital sky survey: final data from sdss-iii,” *The Astrophysical Journal Supplement Series*, vol. 219, no. 1, p. 12, 2015.
- [25] L. v. d. Maaten and G. Hinton, “Visualizing data using t-sne,” *Journal of machine learning research*, vol. 9, no. Nov, pp. 2579–2605, 2008.
- [26] M. Goldberg, P. Boucher, and S. Shlien, “Image compression using adaptive vector quantization,” *IEEE Transactions on Communications*, vol. 34, no. 2, pp. 180–187, 1986.
- [27] N. Nethercote and J. Seward, “Valgrind: a framework for heavyweight dynamic binary instrumentation,” in *ACM Sigplan notices*, vol. 42, no. 6. ACM, 2007, pp. 89–100.
- [28] Y. W. Yu, D. Yorukoglu, J. Peng, and B. Berger, “Quality score compression improves genotyping accuracy,” *Nature Biotechnology*, vol. 33, no. 3, pp. 240–243, 2015.
- [29] M. Muja and D. G. Lowe, “Fast approximate nearest neighbors with automatic algorithm configuration,” *VISAPP (1)*, vol. 2, no. 331–340, p. 2, 2009.
- [30] M. E. Houle and J. Sakuma, “Fast approximate similarity search in extremely high-dimensional data sets,” in *21st International Conference on Data Engineering (ICDE’05)*. IEEE, 2005, pp. 619–630.
- [31] K. Hajebi, Y. Abbasi-Yadkori, H. Shahbazi, and H. Zhang, “Fast approximate nearest-neighbor search with k-nearest neighbor graph,” in *Twenty-Second International Joint Conference on Artificial Intelligence*, 2011.
- [32] M. Muja and D. G. Lowe, “Scalable nearest neighbor algorithms for high dimensional data,” *IEEE transactions on pattern analysis and machine intelligence*, vol. 36, no. 11, pp. 2227–2240, 2014.
- [33] S. F. Altschul, W. Gish, W. Miller, E. W. Myers, and D. J. Lipman, “Basic local alignment search tool,” *Journal of Molecular Biology*, vol. 215, no. 3, pp. 403–410, 1990.
- [34] A. Gionis, P. Indyk, R. Motwani *et al.*, “Similarity search in high dimensions via hashing,” in *Vldb*, vol. 99, no. 6, 1999, pp. 518–529.
- [35] S. Har-Peled, P. Indyk, and R. Motwani, “Approximate nearest neighbor: Towards removing the curse of dimensionality,” *Theory of computing*, vol. 8, no. 1, pp. 321–350, 2012.
- [36] S. Kannan and D. Tse, “Fundamental limits of search,” *Cell systems*, vol. 1, no. 2, pp. 102–103, 2015.