

Accurate Reconstruction of Microbial Strains Using Representative Reference Genomes

Zhemín Zhou, Nina Luhmann, Nabil-Fareed Alikhan, Christopher Quince, and Mark Achtman

Warwick Medical School, University of Warwick, Coventry, United Kingdom

Abstract. Exploring the genetic diversity of microbes within the environment through metagenomic sequencing first requires classifying these reads into taxonomic groups. Current methods compare these sequencing data with existing biased and limited reference databases. Several recent evaluation studies demonstrate that current methods either lack sufficient sensitivity for species-level assignments or suffer from false positives, overestimating the number of species in the metagenome. Both are especially problematic for the identification of low-abundance microbial species, e.g. detecting pathogens in ancient metagenomic samples. We present a new method, SPARSE, which improves taxonomic assignments of metagenomic reads. SPARSE balances existing biased reference databases by grouping reference genomes into similarity-based hierarchical clusters, implemented as an efficient incremental data structure. SPARSE assigns reads to these clusters using a probabilistic model, which specifically penalizes non-specific mappings of reads from unknown sources and hence reduces false-positive assignments. Our evaluation on simulated datasets from two recent evaluation studies demonstrated the improved precision of SPARSE in comparison to other methods for species-level classification. In a third simulation, our method successfully differentiated multiple co-existing *Escherichia coli* strains from the same sample. In real archaeological datasets, SPARSE identified ancient pathogens with $\leq 0.02\%$ abundance, consistent with published findings that required additional sequencing data. In these datasets, other methods either missed targeted pathogens or reported non-existent ones. SPARSE and all evaluation scripts are available at <https://github.com/zheminzhou/SPARSE>.

1 Introduction

Shotgun metagenomics generates DNA sequences directly from environmental samples, revealing unculturable organisms in the community as well as those that can be isolated. The resulting data represents a pool of all species within a sample, thus raising the problem of identifying individual microbial species and their relative abundance within these samples. Methods for such taxonomic assignment are either based on *de novo* assembly of the metagenomic reads, or take advantage of comparisons to existing *reference* genomes. Here we concentrate on the latter strategy, which relies on the diversity of genomes in ever-growing reference databases. This strategy has been instrumental in identifying many causative agents of ancient pandemics in reads obtained from archaeological samples by detecting genetic signatures of modern human pathogens [26].

Published methods for taxonomic assignment can be divided into two categories. *Taxonomic profilers* maintain a small set of curated genomic markers, which can be universal (e.g. used in MIDAS [16]) or clade-specific (e.g. used in MetaPhlan2 [24]). Metagenomic reads that align onto these genomic markers are used to extrapolate the taxonomic composition of the whole sample. These tools are usually computationally efficient with good precision. However, they also tend to show reduced resolution for species-level assignment [23], especially when a species has a low abundance in the sample and, hence, may have few reads mapping to a restricted set of markers.

Alternatively, *taxonomic bidders* compare metagenomic reads against reference genomes to achieve read-level taxonomic classification. The comparisons can be kmer-based (e.g. Kraken [25] and One Codex [15]) or alignment-based (MEGAN [6], MALT [5] and Sigma [1]). Binning methods based on kmers are usually fast, whilst alignment-based methods have greater sensitivity to distinguish the best match across similar database sequences. Benefiting from much larger databases in comparison to genomic markers used by profiling methods, binning methods usually detect more microbial species at very low abundance. However,

45 they also tend to accumulate inaccurate assignments (false positives) [23] due to the incompleteness of the
46 databases, resulting in reads from unrepresented taxa being erroneously attributed to multiple relatives.

47 While microbial species of low abundance are hard to identify by marker-based taxonomic profilers, the
48 estimations of taxonomic bidders can be hard to interpret due to their low precision. This problem especially
49 limits their application to the *in silico* screening of microbial content in sequenced archaeological materials [8].
50 Given that the ancient DNA fragments are expected to exist in low proportions in these samples, methods
51 need to identify weak endogenous signatures hidden within a complex background that is governed by modern
52 (environmental) contamination. Furthermore, reads from archaeological samples are fragmented and have
53 many nucleotide mis-incorporations due to postmortem DNA damage.

54 We identify two challenges that limit the performance of species-level assignments. First and foremost,
55 the reference database used for all taxonomic binnings are not comprehensive. The vast majority of microbial
56 genetic diversity reflect uncultured organisms, which have only rarely been sequenced and analyzed. Even
57 for the bacteria that have genomic sequences, their data are biased towards pathogens over environmental
58 species. This leads to the next challenge where, due to the lack of proper references, reads from unknown
59 sources can accidentally map onto distantly related references, mainly in two scenarios: 1) Foreign reads
60 originating from a mobile element can non-specifically map to an identical or similar mobile element in a
61 known reference. 2) Reads originated from Ultra-Conserved Elements (UCEs), which preserve their nucleotide
62 sequences between species, can also non-specifically map to the same UCE in an existing genome.

63 Addressing both of these challenges, we designed SPARSE (**S**train **P**rediction and **A**nalysis using **R**epre-
64 sentative **S**Equences). In SPARSE, we index all genomes in large reference databases such as RefSeq into
65 hierarchical clusters based on different sequence identity thresholds. A representative database that chooses
66 one sequence for each cluster is then compiled to facilitate a fast but sensitive analysis of metagenomic
67 samples with modest computational resources. Details are given in Section 2. Further, SPARSE implements
68 a probabilistic model for sampling reads from a metagenomic sample, which extends the model described in
69 Sigma [1] by weighting each read with its probability to stem from a genome not included in the reference
70 database, hence considered as an unknown source. Details are given in Section 3.

71 We evaluate SPARSE on three simulated datasets published previously [14, 21, 23]. Comparing SPARSE
72 to several other taxonomic binning software in these simulations shows its improved precision and sensitivity
73 for assignments on the species-level or even strain-level. We further evaluate SPARSE on three ancient
74 metagenomic datasets, demonstrating the application of SPARSE for ancient pathogen screening. For all
75 three datasets, SPARSE is able to correctly identify small amounts of ancient pathogens in the metagenomic
76 samples that have subsequently been confirmed by additional sequencing in the respective studies.

77 2 Database indexing

78 2.1 Background

79 *Average nucleotide identity.* To catalog strain-level genomic variations within an evolutionary context, we
80 need to reconcile all the references in a database into comprehensive classifications. Since its first publication,
81 the average nucleotide identity (ANI) in the conserved regions of genomes has been widely used for such a
82 purpose [10]. In particular, 95 – 96% ANI roughly corresponds to a 70% DNA-DNA hybridization value,
83 which has been used for ~ 50 years as the definition for prokaryotic species.

84 Marakeby et al. [13] proposed a hierarchical clustering of individual genomes based on multiple levels of
85 ANIs. Extending from the 95% ANI species cut-off, it allows the classification of further taxonomic levels from
86 superkingdoms to clones. However, the standard ANI computation adopts BLASTn [2] to align conserved
87 regions between genomes, which is intractable to catalog large databases of reference genomes. We therefore
88 rely on an approximation of the ANI by MASH [18] to speed-up comparisons.

89 *ANI approximation.* MASH uses the MinHash dimensionality-reduction technique to reduce large genomes
90 into compressed sketches. A sketch is based on a hash function applied to a kmer representation of a genome,
91 and compression is achieved by only including the s smallest hash values of all kmers in the genome in the

92 sketch. Comparing the sketches of two genomes, MASH defines a distance measure under a simple Poisson
93 process of random site mutation that approximates ANI values as shown in [18].

94 *Parameter estimation.* Ondov et al. [18] already used MASH to group all genomes in RefSeq into ANI 95%
95 clusters. We adopted slightly different parameters and extended it to an incremental, hierarchical clustering
96 system. The accuracy of the MASH distance approximation is determined by both the kmer length k and
97 the sketch size s . Increasing k can reduce the random collisions in the comparison but also increase the
98 uncertainty of the approximation. We can determine k according to equation (2) in [18]:

$$k = \lceil \log_{|\Sigma|}(n(1-q)/q) \rceil,$$

99 where Σ is the set of all four possible nucleotides $\{A, C, G, T\}$, n is the total number of nucleotides and
100 q is the allowed probability of a random kmer to be found in a dataset. Given $n = 1$ terabase-pairs (Tbp;
101 current size of RefSeq) and $q = 0.05$, which allows a 5% chance for a random k-mer to be present in a 1 Tbp
102 database, we obtain a desired kmer size $k = 23$. Increasing the sketch size s will improve the accuracy of
103 the approximation, but will also increase the run time linearly. We chose $s = 4000$ such that for 99.9% of
104 comparisons that have a MASH distance of 0.05, the actual ANI values fall between 94.5 – 95.5%.

105 2.2 SPARSE reference database

106 We combine the hierarchical clustering of several ANI levels with the MASH distance computation to generate
107 a representation of the current RefSeq [17] database. The construction of the SPARSE reference database is
108 parallelized and incremental, thus the database can be easily updated with new genomes without a complete
109 reconstruction.

110 *Hierarchical clustering.* In order to cluster genomes in different levels, we defined 8 different ANI values
111 $L = [0.9, 0.95, 0.98, 0.99, 0.995, 0.998, 0.999, 0.9995]$, in which the genetic distances of two sequential levels
112 differ by ~ 2 fold. The first four ANI levels differentiate strains of different species, or major populations
113 within a species. The latter four levels give fine-grained resolutions for intra-species genetic diversities, which
114 can be used to construct clade-specific databases for specific bacteria.

115 The SPARSE database $D(S, L, K)$ is extended incrementally as shown in Algorithm 1, with S listing the
116 sketches of all genomes already in the database and K being a hash containing the cluster assignments at
117 each level $l \in L$ for each key $s \in S$. A new genome is integrated by finding another genome in the database
118 with the lowest distance using MASH, and clustering it with its nearest neighbour s_n depending on the ANI.

Algorithm 1 Incremental SPARSE database clustering

Input: SPARSE database $D(S, L, K)$, list of new genomes G

Output: Extended SPARSE database $D'(S, L, K)$

```
1: for each genome  $g \in G$  do
2:    $s_g = MashSketch(g)$ 
3:    $s_n = argmin_{s \in S} MashDistance(s_g, s)$ 
4:   for  $0 \leq i \leq |L| - 1$  do
5:     if  $L[i] \leq 1 - MashDistance(s_g, s_n)$  then
6:       Push  $K[s_n][i]$  to  $K[s_g]$ 
7:     else
8:       Push  $|S|$  to  $K[s_g]$ 
9:   Push  $s_g$  to  $S$ 
```

119 In the SPARSE implementation, we parallelized the database construction by inserting batches of genomes
120 at once and parallelizing sketch and distance computation, thereby scaling to the complexity of the problem.
121 After being added to the database, the cluster assignment for a genome is fixed and never redefined. Therefore,

122 the insertion order of genomes can influence the database structure. Here we utilize prior knowledge from the
123 community, so the SPARSE database is initialized first with all *gold standard* complete genomes in RefSeq,
124 followed by representative and curated genomes. With this strategy, the whole RefSeq database with 101,680
125 genomes (Aug. 2017) can be downloaded and assigned into ANI levels in ~ 23 hrs, using 20 processes on a
126 standalone server. Further insertion of 1,000 new genomes (~ 5 MB) into an already established database
127 takes ~ 15 mins.

128 *Representative database.* To avoid mapping metagenomic reads to redundant genomes within the database,
129 we construct a database of genome representatives for read assignment, similar to [9]. The representative
130 database consists of the first genome from each cluster defined by ANI 99% and is indexed using bowtie2-
131 build [11] with standard parameters. SPARSE indexes 20,850 bacterial representative genomes in ~ 4 hours
132 using 20 computer processes. Representative databases of other ANI levels or clade-specific databases can also
133 be built by altering the parameters. Furthermore, traditional read mapping tools such as bowtie2 [11] show
134 reduced sensitivity for divergent reads. This is not a problem for many bacterial species, especially bacterial
135 pathogens, because these organisms have been selectively sequenced. However, fewer reference genomes are
136 available for environmental bacteria and eukarya. In order to map reads from such sources to their distantly
137 related references, SPARSE also provides an option to use MALT [5], which is slower than bowtie2 and needs
138 extensive computing memory, but can efficiently align reads onto references with $<90\%$ similarity.

139 3 Metagenomic read sampling

140 Given read mappings to the representative databases as input, we adapt a probabilistic model reconstructing
141 the process of sampling reads from a metagenomic sample to assign reads onto reference genomes. We extend
142 the model implemented in Sigma [1] by also considering that reads aligned to a genome in the reference
143 database could still be originating from an unknown source, thus avoiding to overestimate the number of
144 genomes present in the sample. We introduce a weighting for each read reflecting the probability to be
145 sampled from an unknown genome, and show in Section 4 how this improves the precision of taxonomic
146 assignments.

147 Let E denote the set of both known and potentially unknown genomes in a metagenomic sample, and
148 the set of reference genomes included in the SPARSE database is a subset $G \in E$. Let $Pr(r_i|E)$ be the
149 probability of sampling a random read r_i from any possible source, we have

$$Pr(r_i | E) = Pr(r_i, G | E)Pr(r_i | G).$$

150 We denote $w_i = Pr(r_i, G|E)$ as the *sampling probability*, indicating the probability that r_i is sampled from
151 any known reference genome in G . On the other hand, $Pr(r_i | G)$ is the probability of generating r_i given G
152 and can be further separated as

$$Pr(r_i | G) = \sum_{g_j \in G} Pr(r_i | g_j)Pr(g_j | G),$$

153 where $Pr(g_j|G)$ is the probability that a genome $g_j \in G$ was chosen to generate the read, and $Pr(r_i|g_j)$ is
154 the probability of obtaining read r_i from g_j . As in Sigma, given a uniform mismatch probability $\sigma = 0.05$,
155 $Pr(r_i|g_j)$ can be directly calculated from the alignment of r_i to genome g_i with x mismatches, and can be
156 stored in a matrix Q , such that

$$Q_{i,j} = Pr(r_i | g_j) = \sigma^x(1 - \sigma)^{l-x},$$

157 where l is the length of read r_i . We next describe how the sampling probability w_i is inferred, by giving a
158 weight to each read that indicates the probability of being sampled from a known reference genome. Reads
159 with a low weight do not influence the optimization process used to infer the optimal $Pr(g_j|G)$ for a complete
160 metagenomic read dataset.

161 3.1 SPARSE sampling probability

162 We model two scenarios that can lead to non-specific mappings of foreign reads.

163 1) Since there is no systematic way of masking all mobile elements in a reference sequence, we evaluate the
164 probability of a read being drawn from the core genome. We assume that highly conserved regions are part
165 of the core genome, which has been vertically inherited, whereas variable regions likely represent horizontal
166 gene transfers (HGTs). We denote this *HGT probability* as m_i .

167 2) We evaluate the probability of a read originating from an Ultra-Conserved Element (UCE), by com-
168 paring the read depths of the aligned genome fragments with other regions in the genome. UCEs are so
169 highly conserved that additional reads from divergent genomes are likely to map on to them, which results
170 in a higher read depth than other regions. We denote this *UCE probability* as n_i . Combining both cases as
171 a joint probability, we infer a weight w_i for each read as

$$w_i = m_i n_i.$$

172 *HGT probability.* Given any cluster t in ANI level k that consists of u references, a read r_i can be assigned
173 to either the core genome g_c or accessory genome g_a of this cluster. Given the number of references $v \subseteq u$
174 the read aligns to, we can formulate the probability of the read originating from the core genome as

$$Pr_t(g_c|r_i) = \frac{Pr_t(r_i|g_c)Pr(g_c)}{Pr(r_i)} = \frac{Pr_t(r_i|g_c)Pr(g_c)}{Pr_t(r_i|g_c)Pr(g_c) + Pr_t(r_i|g_a)(1 - Pr(g_c))}, \quad (1)$$

$$Pr_t(r_i|g_c) = p_c^v(1-p_c)^{u-v}, \quad Pr_t(r_i|g_a) = p_a^v(1-p_a)^{u-v}$$

175 where $Pr(g_c)$ is the prior probability of any read originating from a core genomic region, and p_c and p_a
176 are the respective probabilities for core genomic fragments or accessory genomic fragments. Default prior
177 probabilities in SPARSE are given in Table 1. Furthermore, a read can align to multiple clusters in the same
178 ANI level k , so we average the probabilities of all such clusters for each read weighted by Q inferred from
179 the read alignment:

$$Pr_k(g_c|r_i) = \frac{\sum_t \max_{g_j \in t} Q_{i,j} Pr_t(g_c|r_i)}{\sum_t \max_{g_j \in t} Q_{i,j}}.$$

180 Finally, we consider three different ANI levels for the core genome analysis (by default 90%, 95% and 98%),
181 assigning a lower value for m_i if the read does not map to the core genome at any of these ANI levels:

$$m_i = 1 - \prod_k (1 - Pr_k(g_c|r_i)). \quad (2)$$

182 Default values for the prior probabilities were inferred from a published study of core genes across multiple
183 bacterial species [3]. We account for 1% of random deletions of core genes, which gives $p_c = 0.99$. We also
184 observed that <10% of all genes are core genes in bacterial species represented by many genomes. This results
185 in $\sum Pr(g_c) < 0.1$ over all three ANI levels. We arbitrarily assigned a higher $Pr(g_c)$ for levels with lower
186 ANI, because a sequence fragment is less likely to be part of a mobile element if it is coincidentally present in
187 more divergent genomes. Finally, ~40% of the genes in a random genome are core genes. This gives $m_i \approx 0.6$
188 when $v = 1$ and $u = 1$, which can be used to find empirical values of p_a via equations 1 and 2.

189 *UCE probability.* In order to compare the read coverage of each fragment in a reference genome g_j with other
190 fragments of the same genome, we split its sequence into k consecutive fragments $f_{j,k}$ using two uniform
191 arbitrary lengths, 487 bps and 2000 bps. Here 487 is used because it is a prime, such that the ends of
192 two fragments overlap only once per Mbp. Then the read depth in each fragment, d_k , follows a Poisson
193 distribution with parameter λ as the average number of reads per region and probability mass function

Table 1: Default prior probabilities for three ANI levels, values inferred from [3].

ANI	$Pr(g_c)$	p_c	p_a
90%	0.05	0.99	0.1
95%	0.02	0.99	0.2
98%	0.01	0.99	0.5

194 $f(k, \lambda)$. Because of the complexity of the read alignments, we relax the probability of read depth in each
 195 fragment such that a wide range of read depths retain high probabilities:

$$Pr(r_i | f_{j,k}) = \begin{cases} \frac{f(d_k, \lambda/\sqrt{2})}{f(\lambda/\sqrt{2}, \lambda/\sqrt{2})} & \text{for } d_k < \lambda/\sqrt{2}, \\ 1 & \text{for } \lambda/\sqrt{2} \leq d_k \leq \sqrt{2}\lambda, \\ \frac{f(d_k, \sqrt{2}\lambda)}{f(\sqrt{2}\lambda, \sqrt{2}\lambda)} & \text{for } \sqrt{2}\lambda < d_k, \end{cases}$$

196 Since a read can again align to multiple genomes g_j , we compute the UCE probability of a read as a weighted
 197 average of all its alignments. If a read aligns multiple times to the same genome g_j with equal alignment
 198 score, we choose one fragment randomly. The UCE probability is then defined as

$$n_i = \frac{\sum_j (Q_{i,j} Pr(r_i | f_{j,k}))}{\sum_j Q_{i,j}}.$$

199 Thus a lower value of n_i is the result from a deviation of the general coverage at the read position in
 200 comparison to the average coverage in the genome, indicating that the read is likely mapping to an ultra-
 201 conserved region in the genome.

202 3.2 Optimization problem

203 Knowing the weight w_i for all reads r_i in a whole metagenomic read set R , the task is then simplified to
 204 finding optimal $Pr(g_j | G)$ values that maximize the probability of the whole read set:

$$\max Pr(R|E) = \max \prod_{r_i \in R} Pr(r_i | E) = \max \prod_{r_i \in R} (w_i \sum_{g_j \in G} Q_{i,j} Pr(g_j | G)).$$

205 The optimization problem can be solved by a non-linear programming (NLP) method. In SPARSE, we rely
 206 on a modified version of the function provided in Sigma [1].

207 After optimizing $Pr(g_j | G)$, we finally assign a read to a potential reference by checking the following
 208 ratio of the computed probabilities:

$$P(r_i, g_j) = \frac{Pr(r_i, g_j | G)}{Pr(r_i, G)} = \frac{Q_{i,j} * Pr(g_j | G)}{\sum_{g_j \in G} Q_{i,j} * Pr(g_j | G)}. \quad (3)$$

209 We may assign a read to multiple references, as long as $\frac{P(r_i, g_j)}{\max_g P(r_i, g)} \geq 0.1$. This allows a better abundance
 210 estimation for multiple strains from the same species, in which case a read cannot be assigned unambiguously
 211 to a single reference.

212 Further, let $r_i \in B \subset R$ be all reads assigned to g_j . For a read r_i of length l with x mismatches in
 213 the alignment to its assigned reference, we have a nucleotide similarity of $s_{i,j} = \frac{l-x}{l}$. The weighted average
 214 similarity $\bar{s}_{B,j}$ can be calculated as

$$\bar{s}_{B,j} = \frac{\sum_{r_i \in B} s_{i,j} w_i P(r_i, g_j)}{\sum_{r_i \in B} w_i P(r_i, g_j)}.$$

215 Potentially, reads assigned to a single reference could still originate from several co-existing genomes,
 216 with varying degrees of diversity, in the metagenome. We can identify reads from more divergent sources
 217 by comparing $s_{i,j}$ to their average similarity. If all reads assigned to a single reference originate from the
 218 same genome in the metagenome, we assume that the similarity of most reads complies with the average
 219 similarity over all reads. However, reads originating from very conserved regions show higher similarity than
 220 the average and provide a sampling bias. On the other hand, reads originating from different more divergent
 221 genomes, will show lower similarity which can be used to avoid overestimating the abundance of each cluster.
 222 Therefore we compute the expected average nucleotide identity s' for r_i as

$$s'_{i,j} = \min(s_{i,j}, \bar{s}_{B,j}).$$

223 This similarity reflects the ANI between each read and the assigned reference and, as described in the next
 224 section, can be used to compute the abundance of each cluster in the metagenomic sample.

225 3.3 ANI cluster abundances

226 The equation $m_i n_i P(r_i, g_j)$ describes the probability, for each read $r_i \in R$, to be drawn from a region in
 227 reference g_j that is part of the core genome (m_i) and has even read depth in comparison to the whole
 228 chromosome (n_i). In summary for all reads assigned to g_j , $\sum_i m_i n_i * P(r_i, g_j)$ gives the frequency of reads
 229 originating from the core genome of g_j . However, the desired read abundance for a reference g_j needs to
 230 also include reads from the accessory genome. Such reads have been previously suppressed when computing
 231 m_i . If we assume that all species have the same proportion of core genome, the relative abundances of their
 232 core genomes will be equal to the relative abundance of their whole genomes. However, since this is not the
 233 case [3], we need to normalize each m_i computed previously. Given $P(r_i, g_j)$ from Equation 3, for any ANI
 234 90% cluster t , we normalize m_i for a read r_i as

$$m'_i = \frac{\sum_{g_j \in t} \sum_{\substack{r_k \in R, \\ s'_{k,j} \geq 0.9}} P(r_k, g_j)}{\sum_{g_j \in t} \sum_{\substack{r_k \in R, \\ s'_{k,j} \geq 0.9}} m_k P(r_k, g_j)} * m_i.$$

235 Finally, we assign reads into clusters of all ANI levels according to the references contained in the cluster. For
 236 each cluster, we only assign reads if its similarity complies with the ANI level l of the cluster, i. e. $s'_{i,j} \geq l$.

237 Thus the abundance of a cluster t_l is computed as the sum of all read abundances assigned to all
 238 genomes in the cluster weighted by their probability to originate from an unknown genome. Therefore
 239 clusters containing only reads with small n_i and m_i probabilities will receive a low abundance value even if
 240 many reads are assigned to it.

$$a_{t_l} = \sum_{g_j \in t_l} \sum_{\substack{r_i \in R \\ s'_{i,j} \geq l}} m'_i n_i P(r_i, g_j)$$

241 3.4 Taxonomic labels for ANI clusters

242 We finally assign standard taxonomic designations to all clusters at all ANI levels, in order to interpret their
 243 biological meaning. Here we rely on a majority vote of all genomes in a cluster. However, the taxonomic
 244 levels are restricted to certain ANI levels. For example, species are distinguished at the ANI 95% level, and
 245 a species designation is therefore inappropriate for an ANI 90% cluster. Similarly, the taxonomic label for
 246 an ANI 95% cluster should not include any subspecies designations.

247 4 Evaluation

248 4.1 Representative Database

249 We ran SPARSE to index the RefSeq database that consists of 101,680 complete or draft genomes into
250 28,732 clusters at ANI 99% level, which were further grouped into 18,205 clusters at 95% ANI level, as
251 shown in Fig. 1. Grouping all the genomes according to their species, the resulting representative database
252 is much more evenly distributed, with a Pielou's evenness [19] of $J' = 0.9$, comparing to $J' = 0.51$ for the
253 whole RefSeq database. Over-representation of pathogenic organisms in the RefSeq database are largely due
254 to repeated sequencing of nearly identical genomes rather than sequencing of intra-species genetic diversities.
255 In particular, nearly half of the genomes in RefSeq are from the top 10 most sequenced bacterial species,
256 which are all human pathogens. All these genomes were grouped into 615 clusters at ANI 99% level, which
257 gives a 65-fold reduction of the data indexed for these species.

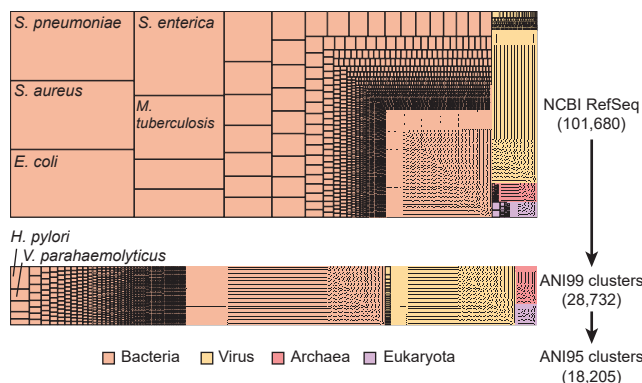


Fig. 1: Hierarchical clustering of 101,680 genomes in NCBI RefSeq database (Aug. 2017) into 18,205 ANI 95% clusters using SPARSE. Each rectangle represents such a cluster at ANI 95% level, with its area relative to the total number of genomes (top) or clusters at ANI 99% (bottom).

258 4.2 Simulated Data

259 We ran SPARSE on three recent simulated datasets (Sczyrba et al. [23], McIntyre et al. [14] and Quince
260 et al. [21]). For a fair comparison, the analyses for all datasets were based on a database built from NCBI
261 RefSeq and taxonomy databases dated 22th June, 2015, which is the deadline for the comparison in [23] and
262 also pre-dates the other two comparisons. We evaluated the performance of SPARSE as described in the
263 respective papers for the read-level taxonomic binners, adopting their results for the compared methods. We
264 also included Sigma using the same database as SPARSE in the comparison. We calculated sensitivity and
265 precision based on the number of true-positives (TP; correctly assigned reads), false-positives (FP; incorrectly
266 assigned reads), and false-negatives (FN; unassigned reads).

267 All simulated reads in the McIntyre et al. [14] study were generated from published complete genomes.
268 This dataset is suitable for comparing the completeness of the databases, as well as the sensitivity of the
269 read mapping approaches in different tools. Both SPARSE and Sigma were run on 18 samples that have
270 read-level taxonomic labels. SPARSE binned all the samples in 10 hours with 20 processes. The precision
271 and sensitivity of both tools in addition to six binning tools from [14] are summarized in Figure 2A. As
272 expected, all tools reached a high precision of $> 97\%$, but differed in their sensitivity. Benefiting from the
273 representative database, SPARSE and Sigma assigned the highest numbers of reads into correct species. The
274 difference between the two methods is due to their different strategies in the modeling, where Sigma assigned
275 all reads to their possible references, whereas SPARSE filtered out unreliable mappings. An independent run

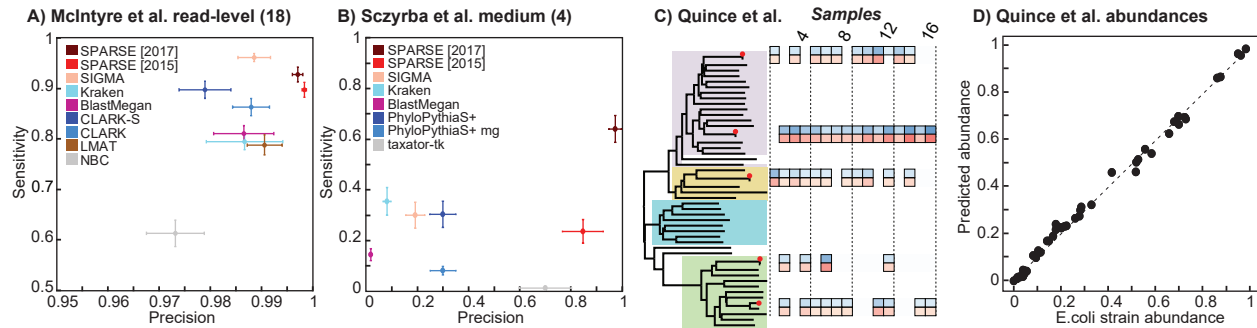


Fig. 2: Performances of SPARSE in simulated published datasets. The performance of all the tools in A and B, except for SPARSE and Sigma, are obtained from the respective publications [23, 14]. SPARSE was run in parallel using two different databases. [2015] uses database built from RefSeq at 2015, whereas [2017] uses up-to-date database. A) All the simulated reads in McIntyre et al. [14] were derived from published genomes. B) The Sczyrba et al. [23] used unpublished genomes for read simulations. C+D) Strain-level identification using the mocked *E. coli* datasets as published in [21]. C) Left: The distance-based species tree for *E. coli* for 45 ANI 99% representative genomes plus the five genomes used in [21] for mocked reads. The four largest ANI 98% clusters in *E. coli* are highlighted with colors. Right: Each column shows one of the 16 mocked samples. The true relative abundances of *E. coli* strains in samples (blue) and the relative abundances of predicted strains (red) in samples are shown as colored squares. D) Comparison of true *E. coli* strain abundances versus SPARSE predictions. The dashed line indicates the linear regression of the two values, with $R^2 = 0.9948$ and $p < 2.2e-16$.

276 of SPARSE using the latest RefSeq database (Aug. 2017) assigned slightly more reads into species, but does
 277 not improve precision. This database consists of 20,850 representative genomes, which is > 2 fold the number
 278 of representatives (9,707) in RefSeq 2015.

279 The datasets in Sczyrba et al. [23] are much more challenging, because all the reads were generated
 280 from sequencing of environmental isolates, many of which do not have closely related references in the 2015
 281 database. Furthermore, many reads do not have a known microbial species label, because they are not similar
 282 to any species in SILVA [20], which was used as the gold standard in this study. We ran both Sigma and
 283 SPARSE on the medium complexity datasets, and compared the results with the other methods (see Fig. 2f
 284 in [23]) for the recovery of microbial species (Fig. 2B). Using 80 processes, SPARSE ran through all four
 285 datasets in ~ 40 hours. All the taxonomic binners published in [23] obtained an average precision of $< 30\%$
 286 at species level, except for taxator-tk [4] with a precision of 70% along with the lowest sensitivity ($\sim 1.25\%$).
 287 The performance of Sigma is comparable to other binning tools, whereas SPARSE obtained an exceptionally
 288 high precision of $\sim 85\%$ while still maintaining a sensitivity of $\sim 23\%$. Many incorrect taxonomic bins
 289 predicted in Sigma were suppressed in SPARSE, because they have low sampling probability w_i to any of
 290 the existing references. Again, SPARSE was also run independently against the database built Aug. 2017.
 291 We recovered 63% of the species in the CAMI median datasets, with an average precision of 97%.

292 Both benchmarks evaluate the performances of taxonomic binnings on or above species level, but give
 293 no resolution in intra-species diversity. DESMAN [21] allows reference-free recovery of strain-level variations
 294 based on uneven read depths of different strains across multiple samples. It has been compared with two other
 295 strain-level binning methods using mock *E. coli* samples [21]. Applying SPARSE to the same 20 genome
 296 mocks, we recovered 50/51 *E. coli* strains in all 16 samples without any additional strains (false positives),
 297 as shown in Fig. 2C. The only strain that was not recovered by SPARSE is 2011C-3493 in the 12th sample
 298 (Sample733 in [21]), which accounts for only $\sim 0.03\%$ of all *E. coli* reads in the sample. We also obtained
 299 an almost exact correspondence between the relative abundances of the strains and the predictions (Fig. 2
 300 D). A linear regression of real abundances and the predictions gives an $R^2 = 0.9948$ and $p < 2.2e-16$.

Table 2: Real archaeological datasets.

Accession (Details, # reads)	Pathogen (% reads)	SPARSE	Sigma	Metagenomic binning tools ^a			
				Kraken	One Codex	MetaPhlan	MIDAS
ERR650978 [7] (1794AD Hungarian, 1.7M)	MT (0.02%)	+	+	- (CD,ML)	+	-	-
ERR1094783 [12] (5300-yr-old Iceman; 15M)	<i>H. pylori</i> (0.01%)	+	-	+	+	+	-
ERR1018927 [22] (Bronze Age human; 1.6M)	<i>Y. pestis</i> (0.01%)	+	+	-	+	-	-

^a +/− for the identification of the pathogen. Abbreviations for suspicious predictions in bracket (CD: *Corynebacterium diphtheriae*; MA: *M. avium*; ML: *M. leprae*; MT: *M. tuberculosis*; SA: *Staphylococcus aureus*; SE: *S. enterica*; VC: *V. cholerae*; VP: *V. parahaemolyticus*; YE: *Y. enterocolitica*; YP: *Y. pseudotuberculosis*).

301 4.3 Ancient Metagenomes

302 We further evaluated SPARSE and five additional metagenomic tools on three real sets of ancient DNA
303 reads (*Mycobacterium tuberculosis* from [7], *Yersinia pestis* from [22] and *Helicobacter pylori* from [12]) and
304 summarised their results in Table 2. For all samples, the presence of the targeted pathogen, although in
305 very low frequencies ($\leq 0.02\%$), has been confirmed by additional sequencing in the respective publications.
306 MIDAS [16] failed in all three samples and MetaPhlan2 [24] managed to identify *H. pylori* but failed in the
307 other two samples. The results for these two marker-based approaches are consistent with the simulations
308 discussed earlier. Kraken [25] and One Codex [15] are both based on kmer-based taxonomic assignment, but
309 yielded different results. Kraken only identified *H. pylori*, whereas One Codex got positive results in all three
310 samples. However both methods incurred a high number of false positives. For example, Kraken reported
311 *Salmonella enterica* and *Vibrio cholerae* in the Iceman sample, whereas One Codex predicted two *Yersinia*.
312 All these predictions are inconsistent with results from other tools and analyses presented in the publications.
313 Sigma identified two of three pathogens but inaccurately predicted *V. parahaemolyticus*, which is normally
314 associated with seafood, for the human remains from the Bronze Age. SPARSE successfully identified all three
315 targeted species without any additional suspicious pathogen, which highlights its application to archaeological
316 samples.

317 5 Conclusion

318 The genetic signatures of specific microbes in metagenomic data, such as human pathogens, are often buried
319 behind the majority of reads from genetically diverse environmental organisms. This is exemplified in the
320 metagenomic sequencing of archaeological samples. Current taxonomic assignment methods compare the
321 metagenomic data with databases that do not fully capture the diversity of microbial genomes. Among these
322 tools, the marker-based taxonomic profilers fail to identify species at low abundances whereas whole genome
323 based taxonomic bidders give inaccurate predictions due to non-specific read mappings on ultra-conserved
324 or horizontally transferred elements.

325 SPARSE indexes existing reference genomes into a comprehensive database with automatic hierarchical
326 clusterings of related organisms. This database is used as a reference for mapping of metagenomic reads.
327 SPARSE penalizes unreliable mappings of reads from unknown sources, and integrates all remaining into
328 a probabilistic model, in which reads were assigned to either an existing reference or unknown sources. In
329 both simulations and real archaeological data, SPARSE outperforms all existing methods, especially in the
330 precisions of species-level assignment. Furthermore, SPARSE managed to identify multiple strains of the
331 same species even when they co-exist in the same sample.

332 References

- 333 1. Ahn, T.H., Chai, J., Pan, C.: Sigma: Strain-level inference of genomes from metagenomic analysis for biosurveil-
334 lance. *Bioinformatics* 31(2), 170–177 (2015)
- 335 2. Altschul, S.F., Gish, W., Miller, W., Myers, E.W., Lipman, D.J.: Basic local alignment search tool. *Journal of*
336 *Molecular Biology* 215(3), 403–410 (1990)
- 337 3. Ding, W., Baumdicker, F., Neher, R.A.: panX: pan-genome analysis and exploration. *bioRxiv* 10.1101/072082
338 (2016)
- 339 4. Dröge, J., Gregor, I., McHardy, A.C.: Taxator-tk: precise taxonomic assignment of metagenomes by fast approx-
340 imation of evolutionary neighborhoods. *Bioinformatics* 31(6), 817–824 (2014)
- 341 5. Herbig, A., Maixner, F., Bos, K.I., Zink, A., Krause, J., Huson, D.H.: Malt: Fast alignment and analysis of
342 metagenomic dna sequence data applied to the tyrolean iceman. *bioRxiv* 10.1101/050559 (2016)
- 343 6. Huson, D.H., Beier, S., Flade, I., Górská, A., El-Hadidi, M., Mitra, S., Ruscheweyh, H.J., Tappu, R.: MEGAN
344 community edition-interactive exploration and analysis of large-scale microbiome sequencing data. *PLoS Com-*
345 *putational Biology* 12(6), e1004957 (2016)
- 346 7. Kay, G.L., Sergeant, M.J., Zhou, Z., Chan, J.Z.M., Millard, A., Quick, J., Szikossy, I., Pap, I., Spigelman,
347 M., Loman, N.J., Achtman, M., Donoghue, H.D., Pallen, M.J.: Eighteenth-century genomes show that mixed
348 infections were common at time of peak tuberculosis in Europe. *Nature Communications* 6, 6717 (2015)
- 349 8. Key, F.M., Posth, C., Krause, J., Herbig, A., Bos, K.I.: Mining metagenomic data sets for ancient DNA: recom-
350 mended protocols for authentication. *Trends in Genetics* 33(8), 508–520 (2017)
- 351 9. Kim, D., Song, L., Breitwieser, F.P., Salzberg, S.L.: Centrifuge: rapid and sensitive classification of metagenomic
352 sequences. *Genome research* 26(12), 1721–1729 (2016)
- 353 10. Konstantinidis, K.T., Tiedje, J.M.: Genomic insights that advance the species definition for prokaryotes. *Pro-*
354 *ceedings of the National Academy of Sciences* 102(7), 2567–72 (2005)
- 355 11. Langmead, B., Salzberg, S.L.: Fast gapped-read alignment with Bowtie 2. *Nature Methods* 9(4), 357–9 (2012)
- 356 12. Maixner, F., Krause-Kyora, B., Turaev, D., Herbig, A., Hoopmann, M.R., Hallows, J.L., Kusebauch, U., Vigl,
357 E.E., Malfertheiner, P., Megraud, F., et al.: The 5300-year-old *Helicobacter pylori* genome of the Iceman. *Science*
358 351(6269), 162–165 (2016)
- 359 13. Marakeby, H., Badr, E., Torkey, H., Song, Y., Leman, S., Monteil, C.L., Heath, L.S., Vinatzer, B.A.: A system to
360 automatically classify and name any individual genome-sequenced organism independently of current biological
361 classification and nomenclature. *PLoS ONE* 9(2) (2014)
- 362 14. McIntyre, A.B.R., Ounit, R., Afshinnekoo, E., Prill, R.J., Hénaff, E., Alexander, N., Minot, S.S., Danko, D., Foon,
363 J., Ahsanuddin, S., et al.: Comprehensive benchmarking and ensemble approaches for metagenomic classifiers.
364 *Genome Biology* 18(1), 182 (2017)
- 365 15. Minot, S.S., Krumm, N., Greenfield, N.B.: One Codex: A Sensitive and Accurate Data Platform for Genomic
366 Microbial Identification. *bioRxiv* 10.1101/027607 (2015)
- 367 16. Nayfach, S., Rodriguez-Mueller, B., Garud, N., Pollard, K.S.: An integrated metagenomics pipeline for strain
368 profiling reveals novel patterns of bacterial transmission and biogeography. *Genome Research* 26(11), 1612–1625
369 (2016)
- 370 17. O’Leary, N.A., Wright, M.W., Brister, J.R., Ciufu, S., Haddad, D., McVeigh, R., Rajput, B., Robbertse, B.,
371 Smith-White, B., Ako-Adjei, D., et al.: Reference sequence (RefSeq) database at NCBI: current status, taxonomic
372 expansion, and functional annotation. *Nucleic Acids Research* 44(D1), D733–D745 (2015)
- 373 18. Ondov, B.D., Treangen, T.J., Melsted, P., Mallonee, A.B., Bergman, N.H., Koren, S., Phillippy, A.M.: Mash: fast
374 genome and metagenome distance estimation using MinHash. *Genome Biology* 17(1), 132 (2016)
- 375 19. Pielou, E.C.: *Ecological diversity*. Wiley New York (1975)
- 376 20. Quast, C., Pruesse, E., Yilmaz, P., Gerken, J., Schweer, T., Yarza, P., Peplies, J., Glöckner, F.O.: The SILVA
377 ribosomal RNA gene database project: improved data processing and web-based tools. *Nucleic Acids Research*
378 41(D1), D590–D596 (2012)
- 379 21. Quince, C., Delmont, T.O., Raguideau, S., Alneberg, J., Darling, A.E., Collins, G., Eren, A.M.: DESMAN: a new
380 tool for de novo extraction of strains from metagenomes. *Genome Biology* 18(1), 181 (2017)
- 381 22. Rasmussen, S., Allentoft, M.E., Nielsen, K., Orlando, L., Sikora, M., Sjögren, K.G., Pedersen, A.G., Schubert,
382 M., Van Dam, A., Kapel, C.M.O., et al.: Early divergent strains of *Yersinia pestis* in Eurasia 5,000 years ago.
383 *Cell* 163(3), 571–582 (2015)
- 384 23. Sczyrba, A., Hofmann, P., Belmann, P., Koslicki, D., Janssen, S., Dröge, J., Gregor, I., Majda, S., Fiedler, J.,
385 Dahms, E., Bremges, A., Fritz, A., Garrido-Oter, R., Jørgensen, T.S., et al.: Critical Assessment of Metagenome
386 Interpretation-a benchmark of metagenomics software. *Nature Methods* (2017)

- 387 24. Truong, D.T., Franzosa, E.A., Tickle, T.L., Scholz, M., Weingart, G., Pasolli, E., Tett, A., Huttenhower, C.,
388 Segata, N.: Metaphlan2 for enhanced metagenomic taxonomic profiling. *Nature methods* 12(10), 902–903 (2015)
- 389 25. Wood, D.E., Salzberg, S.L.: Kraken: ultrafast metagenomic sequence classification using exact alignments.
390 *Genome Biology* 15(3), R46 (2014)
- 391 26. Zhou, Z., Lundstrøm, I., Tran-Dien, A., Duchêne, S., Alikhan, N.F., Sergeant, M.J., Langridge, G., Fotakis, A.K.,
392 Nair, S., Stenøien, H.K., et al.: Millennia of genomic stability within the invasive Para C Lineage of *Salmonella*
393 *enterica*. *bioRxiv* 10.1101/105759 (2017)