skDER: microbial genome dereplication approaches for comparative and metagenomic applications

Rauf Salamzade[1,2], Lindsay R. Kalan[1,3]

[1]Department of Medical Microbiology and Immunology, School of Medicine and Public Health, University of Wisconsin-Madison, Madison, WI, USA
[2]Microbiology Doctoral Training Program, University of Wisconsin-Madison, Madison, WI, USA
[3]M.G. DeGroote Institute for Infectious Disease Research, David Braley Centre for Antibiotic Discovery, Department of Biochemistry and Biomedical Sciences, McMaster University, Hamilton, Ontario, Canada

Abstract

skDER (https://github.com/raufs/skDER) combines recent advances to efficiently estimate average nucleotide identity (ANI) between thousands of microbial genomes by skani[1] with two low-memory methods for genomic dereplication. The first method implements a dynamic algorithm to determine a concise set of representative genomes. This approach is well-suited for selecting reference genomes to align metagenomic reads onto for tracking strain presence across related microbiome samples. This is because fewer representative genomes should alleviate the concern that reads belonging to the same strain get falsely partitioned across closely related genomes. The other method, which uses a greedy approach, is better suited for use in comparative genomics, where users might be overwhelmed with the high number of genomes available for certain taxa and aim to reduce redundancy and, therefore, computational requirements for downstream analytics. This method selects a larger number of representative genomes to comprehensively sample the pangenome space for the taxon of interest. To further aid usage for comparative genomics studies, skDER also features an option to automatically download genomes classified as a particular species or genus in the Genome Taxonomy Database[2–4] and we provide precomputed representative genomes for 17 commonly studied bacterial taxa[5].

Statement of Need

The boundaries of bacterial species are commonly defined using ANI-based thresholds [6–8]. For delineation of more resolute clusters of strains or lineages within species and the selection of individual genomes to serve as representatives, the use of phylogenetic [9], alignment [10,11], and k-mer based [12] approaches have been used as alternatives to ANI-based methods. While these methods have demonstrated scalability to species featuring thousands of genomes, they typically involve multi-step workflows, such as preliminary construction of core genome sequence alignments or phylogenies. In contrast, an ANI-cutoff based approach is more direct to use and simply takes genomic assemblies as input.

For broad comparative genomic investigations of a species or genus, highly sequenced lineages can skew evolutionary statistics [13], especially if such lineages are associated with a single environment, such as hospitals [14]. Thus, for the purposes of removing such biases or simply reducing redundancy to save on computational costs for downstream analyses, genome dereplication can be extremely useful. Methods that have used ANI to perform dereplication of an input set of genomes - to select representatives - have largely been designed for metagenomic applications. A common strategy of these methods [2,15,16] has been to first perform a preliminary clustering using a more efficient though less precise ANI calculator to bin related genomes together. Afterwards, a secondary, more stringent, clustering is performed using a high precision but low efficiency approach, such as FastANI [7]. While this heuristic works well when applied for selecting representative genomes from a set of diverse metagenomic genome assemblies (MAGs), it is not designed to handle high-redundancy datasets with thousands of genomes belonging to the same species or genus. Hence, we developed skDER, which implements methods for high-resolution selection of representative genomes for both comparative genomic and metagenomic investigations based on ANI estimates computed using skani[1].

Methods

A dynamic dereplication method to select distinct reference genomes for metagenomic investigations

In the first dereplication approach, referred to as the “dynamic” algorithm, an empty set, redundant_genomes, is first initialized. Next, pairwise genome-to-genome similarity information from skani is assessed and a genome specific score is used to determine the more preferable genome in the pair to use as a reference. The score for a genome is computed as the product of its assembly N50 and the number of genomes that are regarded as highly similar to it based on user-adjustable ANI and aligned fraction (AF) cutoffs. Assuming the genomes are similar to each other at the required ANI and AF cutoffs and the AF for one genome is not substantially larger than the other, the genome with the lower score, less ideal to serve as a representative, is added to the set of redundant_genomes. This dereplication approach also relies on a third filter, the difference in the AF, which is non-symmetric, between a pair of genomes. If this difference exceeds a certain threshold, the ANI between the genomes meets the ANI cutoff, and the AF for at least one genome exceeds the AF cutoff, then the smaller genome, with the higher AF, will by default be added to the redundant_genomes. The reason being because the smaller genome is largely contained within the larger genome and should not serve as a representative. Once all pairs of genomes are assessed, genomes which are absent in the redundant_genomes set are reported as representatives. It is important to note that this approach approximates selecting a single representative genome for each coarse transitive cluster of related genomes. Therefore, some non-representative genomes might exhibit greater distances to the nearest representative genome than the requested thresholds if they are indirectly connected to the representative genome.

A greedy dereplication method to sufficiently sample the pangenome space for a taxon and simplify comparative genomic studies

The greedy algorithm based approach begins by initializing a set called representative_genomes after which pairwise genome-to-genome similarity information from skani is assessed. For each genome, a list is constructed of alternate genomes deemed similar to it using user-adjustable ANI and AF cutoffs. Similar to the dynamic algorithm based approach, a score indicating the value of a genome to serve as a representative is computed. Genomes are sorted in descending order by this score. The first genome, with the greatest score, is automatically selected as a representative and appended to the representative_genomes set. The direct list of alternate genomes regarded as redundant to this genome are noted and kept track of in another set, nonrepresentative_genomes, disqualifying them from serving as representatives. As the full list of genomes is traversed linearly, they are appended to the set of representative_genomes if they have not previously been added to the set of nonrepresentative_genomes. Similar to the first representative genome, genomes subsequently selected as representatives also have their respective set of genomes deemed as redundant to them appended to the nonrepresentative_genomes set.

Comparison of skDER to other ANI-based approaches for representative genome selection across *Enterococcus*

We compared the performance of skDER (v1.0.4) for representative genome selection across the diverse genus of *Enterococcus*[17] to two other ANI-based dereplication programs, dRep (v 3.2.2) and galah (v0.3.1), which are primarily designed for selecting reference genomes for mapping metagenomic readsets. Both the dynamic and greedy based approaches implemented in skDER as well as dRep and galah were independently applied to a set of 5,291 genomes classified as *Enterococcus* or *Enterococcus*-like (e.g. Enterococcus_A, Enterococcus_B, etc.) in GTDB release 207[3]. FastANI[7] was requested for secondary clustering in dRep. For all four methods, dereplication was performed using an ANI-threshold of 99% with otherwise default parameters. dRep and galah selected 463 and 472 representative genomes in approximately 266 and 546 minutes using 30 threads, respectively. The dynamic algorithm of skDER selected 437 genomes as representatives in only 22 minutes, also using 30 threads. All non-representative genomes exhibited $\geq$ 98.48 ANI and $\geq$ 70% AF to their nearest representative genome, with 91.8% exhibiting $\geq$ 99% ANI to a representative. The greedy algorithm ran for approximately the same time, 25 minutes, and selected 787 representative genomes. All four methods selected at least one genome representative for each individual *Enterococcus* species in GTDB. The faster runtime of skDER relative to dRep and galah is largely due to skani's improved efficiency in computing pairwise ANI relative to FastANI.

One of the key differences between the four methods was the number of representative genomes selected for the two most highly sequenced species in the genus, *E. faecium* and *E. faecalis*. Therefore, we further investigated the coverage of the *E. faecalis* pangenome by representative genomes selected using skDER in comparison to those selected by dRep and galah. Panaroo[18] was used to group genes into ortholog groups for all genomes regarded as

belonging to *E. faecalis* by GTDB (**Figure 1**). Importantly, Prokka [19] was used for gene calling and predicted coding sequences on the edge of scaffolds were disregarded to avoid overinflating the pangenome with partial CDS features due to assembly fragmentation [20]. Of the total 1,902 *E. faecalis* genomes considered, dRep and galah had selected a similar number of representatives, 101 and 99, respectively. These representative genomes featured approximately 60.5% to 61.8% of the total number of distinct genes in the *E. faecalis* pangenome. The dynamic dereplication approach within skDER achieved a similar level of pangenome coverage, finding 10,051 (56.3%) of 17,829 distinct genes, but through selection of only 63 representative genomes. This concise selection approach thus makes it well-suited for metagenomic applications through avoiding over-selection of multiple reference genomes for the same strain. In contrast, the greedy approach in skDER with default parameters led to the greatest number of distinct representative genomes being selected, 198, however achieved the best recovery of the full *E. faecalis* pangenome, featuring 81.5% of the total distinct genes. This showcases the greedy dereplication approach's benefit for simplifying comparative genomics analyses through reducing the number of genomes needed to adequately sample the taxa-wide pangenome.

Limitations and future directions

Notably, bacterial genomes often feature mobile genetic elements (MGEs), such as plasmids and temperate phages, that might be appropriate to disregard during dereplication for certain research objectives. Therefore, in the future, we plan to incorporate an option to use established tools for prediction of such MGEs[21–23] and mask them prior to dereplication. In addition, contamination is currently not considered, which might be particularly problematic for metagenomic assembled genomes (MAGs). We do not plan to account for contamination directly, but do recommend users assess genome contamination using tools such as CheckM[24] or charcoal [25] and remove genomes or individual contigs deemed as contaminated from consideration prior to running skDER.

Acknowledgments

This work was supported by grants from the National Institutes of Health awarded to L.R.K (NIAID U19AI142720 and NIGMS R35GM137828). The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health. We thank Dr. C. Titus Brown, Dr. N. Tessa Pierce-Ward, and Dr. Karthik Anantharaman for helpful discussions on the development of skDER.

References

1. Shaw J, Yu YW. Fast and robust metagenomic sequence comparison through sparse chaining with skani. bioRxiv. 2023. p. 2023.01.18.524587. doi:10.1101/2023.01.18.524587

2. Parks DH, Chuvochina M, Chaumeil P-A, Rinke C, Mussig AJ, Hugenholtz P. A complete domain-to-species taxonomy for Bacteria and Archaea. Nat Biotechnol. 2020;38: 1079–

1086.

3. Parks DH, Chuvochina M, Rinke C, Mussig AJ, Chaumeil P-A, Hugenholtz P. GTDB: an ongoing census of bacterial and archaeal diversity through a phylogenetically consistent, rank normalized and complete genome-based taxonomy. Nucleic Acids Res. 2021. doi:10.1093/nar/gkab776

4. Blin K. ncbi-genome-download: Scripts to download genomes from the NCBI FTP servers. Github; Available: https://github.com/kblin/ncbi-genome-download

5. Salamzade R, Kalan L. skDER Representative Genomes for 17 Bacterial Taxa. 2023. doi:10.5281/zenodo.8267523

6. Goris J, Konstantinidis KT, Klappenbach JA, Coenye T, Vandamme P, Tiedje JM. DNA-DNA hybridization values and their relationship to whole-genome sequence similarities. Int J Syst Evol Microbiol. 2007;57: 81–91.

7. Jain C, Rodriguez-R LM, Phillippy AM, Konstantinidis KT, Aluru S. High throughput ANI analysis of 90K prokaryotic genomes reveals clear species boundaries. Nat Commun. 2018;9: 1–8.

8. Olm MR, Crits-Christoph A, Diamond S, Lavy A, Matheus Carnevali PB, Banfield JF. Consistent Metagenome-Derived Metrics Verify and Delineate Bacterial Species Boundaries. mSystems. 2020;5. doi:10.1128/mSystems.00731-19

9. Menardo F, Loiseau C, Brites D, Coscolla M, Gygli SM, Rutaihwa LK, et al. Treemmer: a tool to reduce large phylogenetic datasets with minimal loss of diversity. BMC Bioinformatics. 2018;19: 164.

10. Cheng L, Connor TR, Sirén J, Aanensen DM, Corander J. Hierarchical and spatially explicit clustering of DNA sequences with BAPS software. Mol Biol Evol. 2013;30: 1224–1228.

11. Tonkin-Hill G, Lees JA, Bentley SD, Frost SDW, Corander J. Fast hierarchical Bayesian analysis of population structure. Nucleic Acids Res. 2019;47: 5539–5549.

12. Lees JA, Harris SR, Tonkin-Hill G, Gladstone RA, Lo SW, Weiser JN, et al. Fast and flexible bacterial genomic epidemiology with PopPUNK. Genome Res. 2019;29: 304–316.

13. Salamzade R, Tran P, Martin C, Manson AL, Gilmore MS, Earl AM, et al. zol & fai: large-scale targeted detection and evolutionary investigation of gene clusters. bioRxiv. 2023. doi:10.1101/2023.06.07.544063

14. David S, Reuter S, Harris SR, Glasner C, Feltwell T, Argimon S, et al. Epidemic of carbapenem-resistant Klebsiella pneumoniae in Europe is driven by nosocomial spread. Nat Microbiol. 2019;4: 1919–1929.

15. Olm MR, Brown CT, Brooks B, Banfield JF. dRep: a tool for fast and accurate genomic comparisons that enables improved genome recovery from metagenomes through de-replication. ISME J. 2017;11: 2864–2868.

16. Woodcroft BJ. galah: More scalable dereplication for metagenome assembled genomes. Github; Available: https://github.com/wwood/galah

17. Lebreton F, Manson AL, Saavedra JT, Straub TJ, Earl AM, Gilmore MS. Tracing the Enterococci from Paleozoic Origins to the Hospital. Cell. 2017;169: 849–861.e13.

18. Tonkin-Hill G, MacAlasdair N, Ruis C, Weimann A, Horesh G, Lees JA, et al. Producing polished prokaryotic pangenomes with the Panaroo pipeline. Genome Biol. 2020;21: 180.

19. Seemann T. Prokka: rapid prokaryotic genome annotation. Bioinformatics. 2014;30: 2068–2069.

20. Klassen JL, Currie CR. Gene fragmentation in bacterial draft genomes: extent, consequences and mitigation. BMC Genomics. 2012;13: 14.

21. Robertson J, Nash JHE. MOB-suite: software tools for clustering, reconstruction and typing of plasmids from draft assemblies. Microb Genom. 2018;4. doi:10.1099/mgen.0.000206

22. Kieft K, Zhou Z, Anantharaman K. VIBRANT: automated recovery, annotation and curation of microbial viruses, and evaluation of viral community function from genomic sequences. Microbiome. 2020;8: 90.

23. Camargo AP, Roux S, Schulz F, Babinski M, Xu Y, Hu B, et al. Identification of mobile genetic elements with geNomad. Nat Biotechnol. 2023. doi:10.1038/s41587-023-01953-y

24. Chklovski A, Parks DH, Woodcroft BJ, Tyson GW. CheckM2: a rapid, scalable and accurate tool for assessing microbial genome quality using machine learning. Nat Methods. 2023;20: 1203–1212.

25. charcoal: Remove contaminated contigs from genomes using k-mers and taxonomies. Github; Available: https://github.com/dib-lab/charcoal

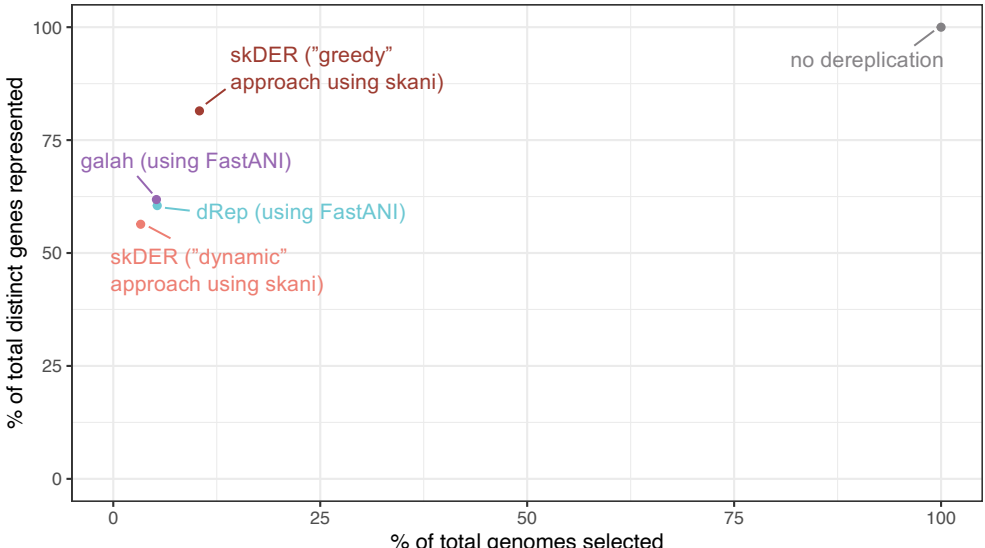

Figure 1: The relationship between the percentage of representative *E. faecalis* genomes selected by different dereplication methods to the percentage of distinct genes in the species' pangenome sampled.