

phylobar: an R package for multiresolution compositional barplots in omics studies

Megan Kuo¹, Kim-Anh Lê Cao^{2,*}, Saritha Kodikara^{2,*}, Jiadong Mao^{2,*}, Kris Sankaran^{1,*}

¹Department of Statistics, University of Wisconsin-Madison, Madison, WI 53703, United States

²Melbourne Integrative Genomics, School of Mathematics and Statistics, University of Melbourne, Melbourne, VIC 3052, Australia

*Corresponding authors. Kris Sankaran, Department of Statistics, University of Wisconsin-Madison, 1205 University Ave Rm 5669, Madison, WI 53703, United States. E-mail: ksankaran@wisc.edu; Jiadong Mao, School of Mathematics and Statistics, University of Melbourne, Grattan Street, Parkville, VIC 3010, Australia. Email: jiadong.mao@unimelb.edu.au; Saritha Kodikara, School of Mathematics and Statistics, University of Melbourne, Grattan Street, Parkville, VIC 3010, Australia. Email: saritha.kodikara@unimelb.edu.au.

Associate Editor: Macha Nikolski

Abstract

Summary: Stacked barplots, though widely used in microbiome studies, can obscure important patterns in microbiome data. They omit rare taxa and can mask shifts that emerge at finer taxonomic levels. To address this issue, we introduce phylobar, an R package that interactively links stacked barplots with overview phylogenetic or taxonomic hierarchies. The interface allows users to collapse or expand subtrees, paint color palettes interactively, and search for specific taxa. This allows comparison across taxonomic resolutions that are hidden in static overviews. phylobar works with any omics data with hierarchical organization, including cell type hierarchies, as we demonstrate in a case study of immune cell composition in COVID-19 patients.

Availability and implementation: phylobar is available as an R package on GitHub. It uses the htmlwidgets library to link interactive D3 visualizations with R. The interactive plots can be embedded within R Markdown or Quarto notebooks, and views can be exported as vector graphics files. The package is open source and documented at <https://mkdiro-O.github.io/phylobar>.

1 Introduction

Stacked barplots are widely used to visualize microbiome community structure. These plots divide each bar—representing a single sample—into colored segments. The colors denote taxonomic categories; the heights reflect group frequencies. This provides a useful exploratory overview for understanding the main taxonomic groups that are present across the dataset. The view also supports comparisons in community structure across samples, especially when the samples are ordered in a meaningful way—e.g. sorting by time or sample similarities. These plots have driven biological discovery in the microbiome literature. For example, they have helped group samples into community state types (DiGiulio *et al.* 2015), highlight heterogeneity among dysbiotic samples (Symul *et al.* 2023), and clarify response to environmental interventions (Raimondi *et al.* 2021). Several R packages are now available to generate these plots (Reeder *et al.* 2021, Lu *et al.* 2023, Yang *et al.* 2025). They can be made from community composition matrices or directly from microbiome-specific data structures, like those exported by phyloseq (McMurdie and Holmes 2013) or TreeSummarizedExperiment (Huang *et al.* 2020) objects.

Stacked barplots use color to graphically encode taxonomic categories. This imposes a basic constraint—when the color palette grows, distinctions become more and more challenging to make. Researchers typically solve this in one of two ways: visualize only high-level taxonomic categories, or show low-level categories for abundant groups and aggregate rare taxa into a bin for “Other” (Reeder *et al.* 2021, Xia and Sun 2024). Yet some widely used packages demonstrate this problem in their own documentation. The qiime2 viewer uses 200 colors in a single bar plot (<https://go.wisc.edu/357516>), and an example in MicrobiomeStat uses 43 (<https://go.wisc.edu/o7uv18>). These palettes make the plots extremely difficult to use.

Recently, the StackbarExtend package (Cuisinieri and Santos 2024) addressed this problem by assigning similar colors to related taxonomic categories. In this way, color indicates taxonomic assignment while ensuring broader categories remain distinguishable. We also use taxonomy to guide the stacked bar plot. However, instead of color gradients, phylobar uses interactivity. Specifically, phylobar exploits hierarchical structure and the focus-plus-context principle (Shneiderman 1996, Sankaran and Holmes 2018, Fukuyama *et al.* 2023), allowing users to collapse and expand the hierarchy to different resolutions and

Received: 5 November 2025. Revised: 26 February 2026. Accepted: 17 March 2026

© The Author(s) 2026. Published by Oxford University Press.

This is an Open Access article distributed under the terms of the Creative Commons Attribution License (<https://creativecommons.org/licenses/by/4.0/>), which permits unrestricted reuse, distribution, and reproduction in any medium, provided the original work is properly cited.

paint the barplot by hovering the mouse over different subtrees. Hence, users can interactively test different color palettes and identify taxonomic groups with interesting variation before finalizing the visualization. Further, this approach allows direct comparison of multiple levels of taxonomic resolution, since large and small subtrees can be painted simultaneously. A detailed comparison between phylobar and both static (Cuisiniere and Santos 2024) and interactive (Ondov *et al.* 2011, Reeder *et al.* 2021) alternatives is given in [supplementary data](#) at *Bioinformatics* online. phylobar applies to compositional visualization in other hierarchically structured omics data. This includes pseudobulk analysis of single-cell data, since cell types can be organized into nested categories across resolution levels. Example 2 below demonstrates phylobar's utility in this single-cell setting.

In addition to initial exploratory analysis, phylobar also enables post-modeling exploration. For example, when a model identifies differentially abundant taxa, the same interactive barplot can be used to visualize the distribution and variability of those taxa across samples and groups, providing a clear, sample-level view of model-derived findings. This approach preserves the benefits of the original stacked bar plot, allowing overviews of the entire dataset and comparison across taxonomic categories. But viewers can now interactively browse taxonomic groups (including rare ones), no longer needing to accept a static set of coarse or pre-defined categories.

1.1 Overview of package functionality

The main visualization function in the phylobar package is called `phylobar`. This function has two required inputs, (i) a data matrix of abundances with samples in rows and taxa in columns, and (ii) a phylo structure defining the taxonomic hierarchy. The leaf nodes of the phylo must match the column (taxa) names of the sample abundance matrix. A leaf node is the lowest level of the tree hierarchy and represents the finest level of taxonomic resolution. Sample abundances need not sum to one, and users can provide non-compositional normalizations to phylobar, resulting in stacked bar plots where the heights across bars are not necessarily equal. phylobar requires a phylo object representing the tree. Considering that hierarchical information often exists in less structured formats, like taxonomic tables, the package includes a function (`taxonomy_to_tree`) that converts these formats into the required phylo structure. Note that the resulting tree does not have to be a binary tree; i.e. each internal node is not restricted to having only two child nodes.

phylobar includes two panels (Fig. 1a). The left panel shows a taxonomic hierarchy drawn as a tree. The right panel shows a stacked barplot where each rectangle within a sample's bar corresponds to one leaf from the tree. Samples can be ordered by hierarchical clustering or by their input data matrix rows. The latter option enables users to reorder samples based on predefined characteristics, such as chronological time, experimental group, or other relevant study factors.

Hovering near a node paints the associated subtree with the user-provided color palette and highlights the associated leaves in the stacked bar plot. Pressing the Control key fixes the currently painted subtrees and adds a new color for painting. Colors can be added until the initial palette is exhausted. By

default, the initial palette is six colors. Once the palette is exhausted, Control cycles through previously used colors. Alternatively, the number keys select specific colors in the palette. While painting, a color legend is interactively updated to display each highlighted subtree's root label. This tree painting strategy makes it possible to compare the abundance of taxonomic categories, some of which may be rare or at different levels of taxonomic resolution, across the full collection of samples in the data.

Large trees benefit from trimming taxonomic subgroups of limited interest. To this end, users can click on a node to collapse all of its descendants. The collapsed node appears with a thicker border to indicate the hidden subtree. The entire tree layout is redrawn to make use of the newly freed space. When a subtree is collapsed, the associated leaf nodes in the stacked bar plot are merged into a single rectangle corresponding to the entire subtree. If a color had been present only within the descendants of a collapsed subtree, that color is removed from the palette. Clicking the collapsed subtree expands it to its previous state.

Hovering over a bar in the stacked barplot magnifies its sample name. When bars are tightly packed, labels can optionally be hidden until the sample is hovered over. Similarly, when a dataset includes many samples, plotting all samples becomes impractical. In this case, the package can select representatives from hierarchical clustering centroids or Bray-Curtis clustering medoids derived from the dataset, and the stacked bar plot will show one representative from each cluster.

Altogether, this interactivity supports efficient comparison of taxonomic categories across subtrees. The visualization's appearance is customizable, and users can modify color palettes, font sizes, and component dimensions. Labels can be placed directly on subtrees rather than in a separate legend. If the analyst finds a useful view, then pressing the Escape key freezes the view for export into a Scalable Vector Graphics (SVG) file. Pressing escape again resumes the interaction. While simple, phylobar supports exploratory analysis of microbiome community structure at different levels of resolution, preserving distinctions even at strain-level variation.

2 Materials and methods

phylobar is an R package whose visualizations are built using the D3 JavaScript library (Bostock *et al.* 2011). D3 transforms JavaScript data arrays into graphical layout coordinates and provides methods for Document Object Model (DOM) manipulation. phylobar uses `d3.hierarchy` and `d3.stack` for tree and stacked bar plot visualizations, respectively. The resulting hierarchical layout includes methods for enumerating children, descendants, and leaf nodes, which support subtree painting and collapsing. A Delaunay triangulation around the tree nodes triggers mouseover events when the cursor moves between node neighborhoods, allowing smoother interaction than direct node manipulation. In response to mousemoves and clicks, D3 manipulates the DOM to change subtree and barplot colors and resolution. The DOM is modified in place without re-rendering the entire view. For example, hovering a subtree updates fill attributes in DOM elements without restructuring them. This




Figure 1 (a) An overview of phylobar's display and interactivity. Clicking or hovering over tree nodes and barplot compositions helps focus comparisons across taxa. A search box supports lookup of individual taxa. (b) High-level overviews confirm Bacteroidia and Firmicutes shifts under the HFHS intervention from [Susin et al. \(2020\)](#). (c) Collapsing unrelated subtrees focuses the view on these two groups. Searching for family Lactobacillaceae confirms its HFHS association, and hovering over bars reveals the elevated abundance of OTU_6 in the HFHS group. (d) An immune cell type hierarchy coupled with cell type compositions from a study of immune factors in COVID-19 severity. (e) Hovering over nodes for T cell subtypes replaces the purple (upper level) stacked bar portions from panel (d) with colors associated with the CD8 T lymphocyte and $\gamma\delta$ subtypes.

ensures smooth response to user inputs and is a key advantage of browser-based visualization. A standalone JavaScript package, phylobar-js, wraps the core data manipulation and

visualization functions, and is available on NPM (<https://www.npmjs.com/package/phylobar>). This package exports reusable components for building interactive hierarchical visualizations,

including functions for interactively collapsing and expanding trees, searching for string matches across hierarchy labels, and detecting mouse hover events across node neighborhoods.

The `htmlwidgets` R package (Vaidyanathan *et al.* 2020) bridges phylobar's JavaScript implementation and R interface. This package converts R matrices and lists into JavaScript arrays and objects. `htmlwidgets` also creates an SVG canvas for D3 visualizations. This canvas can be embedded within Quarto (Evans 2025) or Rmarkdown (Gandrud 2018) notebooks with multiple visualizations per notebook. Views can also be exported to scalable vector graphics files, allowing customization (e.g. font families and edge stroke types) in vector graphics software like Illustrator or Inkscape.

3 Results

3.1 Example 1: high-fat high-sugar diet and the gut microbiome

The gut microbiome is known to be reshaped by high-fat, high-sugar (HFHS) diets. Susin *et al.* (2020) were the first to analyze the HFHS mouse model to characterize diet-driven microbiome effects using samples collected by Lê Cao at the University of Queensland Diamantina Institute. They randomized 47 C57BL/6 female mice into HFHS and control diet groups and gathered fecal samples at days 0, 1, 4, and 7. We use phylobar to re-analyze the preprocessed 16S rRNA sequencing data from Kodikara and Cao (2024). These data were generated using Illumina MiSeq and taxonomic representation was determined using QIIME 1.9.0. To concentrate on long-term effects of the diet intervention, we filter to samples from day 7. In our stacked barplot, we place all control and HFHS samples on the left and right, respectively. Within the two treatment groups, we apply complete link hierarchical clustering on the Euclidean distance between community composition vectors and sort samples according to the leaf order on the resulting trees. We first check whether we observe an effect on the Bacteroides versus Firmicutes ratio, which has previously been found to be associated with HFHS diet. Figure 1b shows the stacked barplots when hovering over class Bacteroidia and Phylum Firmicutes. Consistent with earlier studies, Bacteroidia appears more abundant in the control samples, while Firmicutes is elevated in the HFHS condition. Next we asked whether we could decompose this observation into finer effects within narrower taxonomic groups. Figure 1c shows the visualization after highlighting a particular strain, OTU_6, and searching for family Lactobacillaceae, both of which are members of phylum Firmicutes. Subtrees unrelated to Bacteroidia and Firmicutes have been collapsed. Consistent with the previous view, abundances of both groups increase in the HFHS condition, but they change in qualitatively different ways. OTU_6 is absent in all but a handful of control samples, but it is one of the most abundant taxa in the HFHS communities. In contrast, family Lactobacillaceae is present in most control samples, but at lower levels than in the HFHS condition. In this example, all three forms of interaction—browsing, sorting, and searching—proved important. Interactivity makes it straightforward to compare samples across multiple levels of taxonomic resolution and helps avoid the aggregation required in static alternatives. This encourages multiresolution exploration and helps generate precise, biologically relevant hypotheses.

3.2 Example 2: immune cell type composition in mild to severe COVID-19 patients

In addition to microbiome data, we can use phylobar to visualize other types of compositions. We applied phylobar to a single-cell RNA-seq dataset from Su *et al.* (2020) to examine how immune cell type composition changes across mild, moderate, and severe COVID-19. The input was the proportions of immune cell types from 120 COVID-19 patients' blood samples. Patients were first grouped by their disease severity and then ordered by hierarchical clustering, so that patients with similar cell type compositions within each severity group were more closely placed. A cell type tree representing the hierarchical structure of immune cells was manually defined.

Based on the phylobar visualization we observed that T cells and monocytes made up the majority of cellular populations in patient samples. We also observed an overall increase of monocytes and decrease of NK and T cells (Fig. 1d), especially CD8 T cells, in moderate and severe COVID-19 patients. These changes were also highlighted in Su *et al.* (2020). Another feature revealed by phylobar that was not characterized in Su *et al.* (2020) was the presence of $\gamma\delta$ (gd in the cell type tree) T cells in milder COVID-19 patients exclusively (Fig. 1e). This finding echoes some more recent works suggesting that $\gamma\delta$ T cells may play an important role in the immune response to the pathogenic SARS-CoV-2 virus (von Massow *et al.* 2021, Terzoli *et al.* 2024).

4 Discussion

phylobar adapts the focus-plus-context principle to compare microbiome sample composition across taxonomic groups and levels of resolution. The package interactively links hierarchical overviews with barplots of sample composition in selected subtrees. This approach solves a longstanding limitation of stacked barplots when applied to omics data visualization, where the number of features exceeds what can be meaningfully represented with a qualitative color palette. The display requires only two inputs, allowing for its broad applicability. The package documentation includes nine vignettes: five data analysis walkthroughs (including those that reproduce the microbiome and single-cell examples presented above), and four guides on preparing package inputs and refining visual outputs. For future work, we note that specialized interactions could be tailored to specific experimental designs, like the longitudinal data in (Kodikara and Cao 2024) or pseudotime data in (Kotliar and Colubri 2021). We also note that phylobar can be adapted to hierarchies that are learned from sample abundance or expression profiles (Klimovskaia *et al.* 2019, Li *et al.* 2025), rather than specified in advance. Code and documentation are available at <https://mkdiro-O.github.io/phylobar>.

Author contributions

Megan Kuo (Conceptualization [equal], Methodology [equal], Software [lead], Validation [equal], Visualization [lead], Writing—original draft [lead], Writing—review & editing [equal]), Kim-Anh

Lê Cao (Conceptualization [equal], Funding acquisition [equal], Methodology [equal], Supervision [equal], Writing—review & editing [equal]), Saritha Kodikara (Conceptualization [equal], Methodology [equal], Software [equal], Validation [equal], Visualization [equal], Writing—original draft [equal], Writing—review & editing [equal]), Jiadong Mao (Conceptualization [equal], Methodology [equal], Software [equal], Validation [equal], Visualization [equal], Writing—original draft [equal], Writing—review & editing [equal]), and Kris Sankaran (Conceptualization [equal], Funding acquisition [equal], Methodology [equal], Software [equal], Validation [equal], Visualization [equal], Writing—original draft [equal], Writing—review & editing [equal])

Supplementary material

Supplementary material is available at *Bioinformatics* online.

Conflicts of interest

None declared.

Funding

KS and MK were supported by the National Institute of General Medical Sciences of the National Institutes of Health under award number R01GM152744.

KALC, JM and SK were supported by the National Health and Medical Research Council (NHMRC) Investigator Grant (GNT2025648).

Data availability

All data used in this study are available at the following DOIs: <https://doi.org/10.5281/zenodo.17477876>, <https://doi.org/10.5281/zenodo.18791960>

References

- Bostock M, Ogievetsky V, Heer J. D³ data-driven documents. *IEEE Trans Vis Comput Graph* 2011;**17**:2301–9.
- Cuisiniere T, Santos M. Stackbarextended: a user-friendly stacked bar-plot representation incorporating phylogenetic information and microbiota differential abundance analysis. *F1000Res* 2024;**13**:914. <https://doi.org/10.12688/f1000research.151662.1>
- DiGiulio DB, Callahan BJ, McMurdie PJ *et al*. Temporal and spatial variation of the human microbiota during pregnancy. *Proc Natl Acad Sci USA* 2015;**112**:11060–5. <https://doi.org/10.1073/pnas.1502875112>
- Fukuyama J, Sankaran K, Symul L. Multiscale analysis of count data through topic alignment. *Biostatistics* 2023;**24**:1045–65.
- G

Xia Y, Sun J. Applied microbiome statistics: Correlation, Association, Interaction and Composition, 1st ed. Boca Raton, FL, USA: Chapman and Hall/CRC, 2024. <https://doi.org/10.1201/9781003121572>

Yang C, Chen J, Zhang X *et al.* Comprehensive statistical and visualization methods for microbiome and multi-omics data. 2025. <https://zenodo.org/records/14881492>