BLIMMP-Explorer: A Visual Decomposition of Bayesian Evidence in Metabolic Module Inference

Neha Sontakke*
Department of Computer Science
University of Arizona

Joshua A. Levine†
Department of Computer Science
University of Arizona

Travis J. Wheeler‡
Department of Pharmacy Practice and Science
University of Arizona

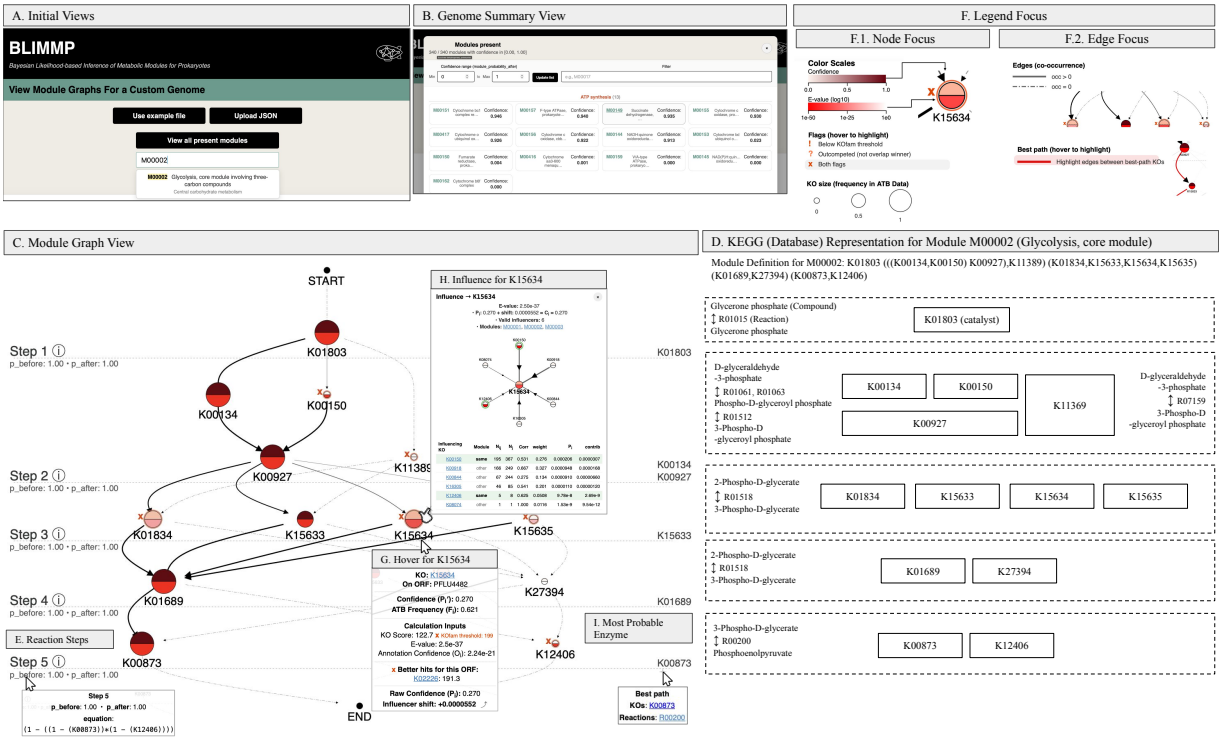

Figure 1: The BLIMMP-Explorer interface. (A) Initial view with module search and file upload. (B) Module browser displaying inferred modules grouped by functional category with confidence scores. (C) Module graph view rendering a metabolic module as a directed acyclic graph. (D) KEGG database representation of the same module, showing the Boolean structure of sequential steps, alternative enzymes, and enzyme complexes. (E) Layered reaction steps with enzymes satisfying the same reaction step sharing rank, separated by dashed lines. (F) Legend: nodes are split circles encoding posterior confidence (top) and sequence match strength (bottom); node size encodes reference genome frequency; edge thickness encodes conditional co-occurrence. (G) Node hover panel showing full evidence decomposition including sequence score, E-value, raw and posterior probability, and shift contributed by associates. (H) Influence diagram showing the probability shift contributed by associated enzymes to a selected node; green outlines mark associates from the current module. (I) Best-path highlight tracing the highest-probability route through the module.

Abstract

We present BLIMMP-Explorer, an interactive visualization system for BLIMMP, a Bayesian tool we developed for inferring metabolic module presence from prokaryotic genomes. Probabilistic tools for metabolic module inference produce layered evidence that tabular outputs cannot clearly communicate. BLIMMP-Explorer renders each module as a gene-centric directed acyclic graph with split-circle nodes encoding two independent evidence values, step-level probability annotations, and flag-based indicators of annotation quality. An influence panel decomposes the probability shift contributed by associated enzymes across a global co-occurrence network. We demonstrate the system on *Pseudomonas fluorescens* SBW25 through two use cases: (i) recovering a pathway in an incomplete genome and (ii) identifying a false positive by reviewing annotation flags and competing enzyme annotations. The visualization enables reasoning that would otherwise require manual cross-referencing of multiple output files.

Keywords: Metabolic module visualization, Bayesian inference visualization, gene-centric graph layout, probabilistic annotation, metagenomics

*e-mail: nsontakke@arizona.edu
†e-mail: josh@cs.arizona.edu
‡e-mail: twheeler@arizona.edu

1 Introduction

Microbial communities shape ecosystems by breaking down organic and inorganic compounds, cycling nutrients, and producing and secreting metabolites [6]. These processes are carried out by *metabolic*

modules, each comprising a series of biochemical reactions that convert one compound into another, catalyzed by enzymes encoded in the organism's genome [12]. Each reaction in a module constitutes a *step*, and a step may be fulfilled by any of several alternative enzymes or may require multiple enzymes working together as a complex. Identifying which modules an organism encodes, and how modules are distributed across the members of a community, reveals the metabolic capabilities of individual microbes and the division of labor among them. This understanding underpins efforts to enhance beneficial activities, engineer new metabolic capabilities, and mitigate harmful ones. Soil microbiome analyses, for instance, inform strategies for building ecosystem resilience following drought and for designing biofertilizers that promote plant growth and nutrient cycling [15].

Inferring metabolic function from environmental genomes is challenging. Understudied bacteria can use uncharacterized proteins or alternative routes that current metabolic pathway protein catalogs do not recognize, known protein families may be too diverged in a given species to identify confidently, and fragmented genome assemblies may leave responsible genes unobserved. These factors yield partial or misleading metabolic inference [4].

BLIMMP (Bayesian Likelihood-based Inference of Metabolic Modules for Prokaryotes) infers metabolic modules even when genome data is incomplete or uncertain (manuscript in prep). Enzyme presence is determined by matching genome sequences against known protein families [1], which are often represented by profile hidden Markov models [7]. BLIMMP combines three sources of evidence. The first is based on the strength of enzyme sequence matches in the recovered genome. The second is based on prior expectations derived from the frequency with which each enzyme and module appears across a large reference database of complete genomes [9]. The third is based on an influence network: enzyme *j* is an *associate* of enzyme *i* if *i* is present in at least 25% of the reference genomes that contain *j*. Associates boost confidence through guilt-by-association [14, 21], with each contribution weighted by their co-occurrence frequency. Since enzymes participate in multiple modules, these connections are global, reaching beyond a singular module. Together these produce a posterior probability for each enzyme, each reaction step, and each module. BLIMMP infers the presence of metabolic modules in environmentally sourced bacterial genomes with state-of-the-art accuracy.

Producing a probabilistic inference is one challenge; understanding it another. A module is called present in a genome if at least one enzyme can be found to catalyse every one of its reaction steps. However, BLIMMP can also call a module present even when not all of its enzymes have been recovered from the genome, based on support from correlated enzymes mixed with prior expectations. We present BLIMMP-Explorer to make the inference process legible to domain scientists through coordinated views designed around three core tasks:

- **Gene-centric layout (DO1, T1):** Genes and reactions have many-to-many relationships: an enzyme can catalyze multiple reactions, and a reaction can be catalyzed by multiple enzymes. Since enzymes determine whether a reaction occurs, identifying which ones drive a module call is essential to interpreting it. BLIMMP-Explorer places enzymes as nodes in a layered graph, with position encoding reaction step membership and edges connecting enzymes satisfying consecutive steps, making each enzyme's contribution to module directly inspectable.
- **Path structure (DO3, T2, T3):** Module presence is path-dependent: a module is present only if a complete route through its Boolean expression is satisfied. BLIMMP-Explorer preserves this topology in the graph layout, making it possible to trace the highest-probability path, identify gaps, and understand which reaction steps remain unsupported.
- **Evidence decomposition (DO2, DO4, T4, T5):** BLIMMP combines sequence confidence, prior expectation, and associate influence. These sources carry qualitatively different biological meaning. BLIMMP-Explorer keeps these sources separable and inspectable through coordinated linked views. A split-circle node displays sequence evidence and posterior probability on each enzyme. An influence panel decomposes each enzyme's posterior into named associate contributions drawn from a global co-occurrence network spanning all modules in which that enzyme participates.

2 Related Work

2.1 Metabolic Module Visualization

Metabolic modules are sets of enzymes that carry out a specific biochemical function. KEGG Modules abstract metabolism into minimal functional units, distilled by curators from the densely interconnected KEGG pathway maps into Boolean combinations of reaction steps with alternative completion paths [12]. We use KEGG module definitions as the basis for the structure of the graph but do not rely on KEGG's own diagrams (Figure 1D), which are static, reaction-centric images with no probabilistic attributes.

Existing tools for metabolic inference visualization collapse module structure in different ways. Reducing each module to a presence-absence call in tabular outputs is common. Heatmaps have been used to represent module completeness scores, colored by reactions or shared enzymes for pathway comparisons [5, 18]. Curated diagrams of biogeochemical cycles render a fixed set of modules with enzymes on edges and compounds on nodes [24], but these are static images with no per-genome evidence or probabilistic content.

The absence of network-based visualization for metabolic module predictions has been identified as a gap [13, 16]. Prior work has approached this gap at the pathway level, drawing whole pathways as compound-reaction graphs with transcriptomics data on reaction nodes [16] or metabolic concentration bargraphs on compound nodes [13]. At this scale, layouts become dense and hard to read, and shared enzymes across modules introduce cycles and duplications that further complicate analysis [2]; multi-scale clustering has been proposed as a remedy, but collapsing reactions into pathway clusters hides the module structure needed to trace paths or inspect enzyme evidence [10].

2.2 Visualizing Bayesian Inference on Graphs

The underlying model for BLIMMP-Explorer combines multiple evidence sources to assign probability of metabolic module presence. Visualizing probability on graph structures is a challenge that has been studied across many domains (including metabolic inference). Encoding multiple probabilistic variables on nodes and edges risks conflation of visual channels; using color for one attribute restricts its use for others [11]. Existing approaches (i) focus on graph structure by encoding probability to edge existence [17];(ii) animating graph realizations via sampling to communicate topological variation [23]; or (iii) focusing on the model-building process through prior checks, posterior predictive checks, and MCMC diagnostics [8]. None of these approaches address the specific structure of metabolic module inference, where evidence comes from multiple distinct sources that each carry biological meaning.

The closest precedent to BLIMMP-Explorer is iVUN [22], a visual analytics system for biochemical reaction networks in which nodes are metabolites, edges are reactions, and edge color is animated to show reaction flux over time, demonstrating that probabilistic attributes can be placed directly on a biological graph and made inspectable through coordinated views. BLIMMP-Explorer differs from iVUN in two respects. First, using nodes to represent enzymes lets per-enzyme evidence attach directly to the visualization. Second, iVUN's uncertainty arises from a single source producing one aggregate distribution per edge, whereas BLIMMP-Explorer represents three distinct evidence sources, each carrying different biological meaning. We address this through coordinated linked views at in-

creasing levels of detail. The enzyme node encodes two evidence values simultaneously, the graph layout preserves the Boolean path structure connecting genes to reactions within a module, and hovering or clicking any element opens additional evidence views on demand.

3 METHODS

3.1 Data Input

BLIMMP-Explorer accepts a JSON file produced by the BLIMMP command-line pipeline. The file encodes information for a single input genome, including observational sequence evidence, module-level calculations, and probabilistic inference across 340 prokaryotic modules from the KEGG database.

3.2 Module Browser

After file parsing, the module browser displays all metabolic modules inferred for the analyzed genome, grouped by functional categories (Figure 1B). A genome contains hundreds of modules, each with some confidence assigned by BLIMMP; the user can specify a confidence range to filter modules per category. Each module is represented as a tile displaying the module name from KEGG, along with BLIMMP-assigned confidence score. Hovering reveals the full KEGG title; clicking opens the module graph.

3.3 Module Graph View

The module graph (Figure 1C) is the central contribution of the system, rendering each metabolic module as an interactive directed acyclic graph using a Sugiyama-style layout [20]: enzymes are nodes positioned by reaction step, with earlier steps placed higher and left-to-right reading order following the module definition (Figure 1D), and edges connect enzymes that can satisfy consecutive steps. Branching alternative enzymes within a step share a rank, keeping the layout non-overlapping even as alternatives grow. This hierarchical approach extends naturally to modules that combine multiple constituent modules into a single definition, such as M00617, which can be explored directly at https://blimmp.org.

Graph Generation KEGG module definitions are Boolean expressions (top of Figure 1D) encoding sequential steps, alternative enzymes, and complexes (Figure 1D). We developed a custom context-free grammar and recursive-descent parser using the Lark library [19] to parse these expressions into an abstract syntax tree representing the module's logical structure. Each module definition is bracketed by a virtual start and sink node marking entry and exit of the metabolic route. A depth-first traversal enumerates all valid paths from start to sink, and their union forms the graph adjacency list. To disambiguate enzymes that appear in multiple steps we append a step suffix to prevent cycles in the DAG. Hierarchical layout assigns each node a layer based on its minimum distance from the sink across all enumerated paths, so enzymes satisfying the same step share a rank (Figure 1E). The resulting adjacency list, node layers, and paths are serialized as JSON and rendered in the browser using the Dagre layout library [3].

Visual Encoding Each **node** (Figure 1F.1. & C) is a split circle encoding two values. The bottom half encodes the strength of the enzyme match on the genome, based on the score of the HMMER alignment between the encoded protein and the Kofam pHMM (red signals high confidence).

The top half is the BLIMMP-assigned presence probability, obtained by first adjusting the sequence-match confidence using the enzyme's background frequency across reference genomes, then further adjusting that estimate using support from co-occurring associate enzymes. A red top and white bottom indicates a gene rescued by prior or network support despite weak sequence evidence. And a white top and red bottom signals strong sequence evidence discounted by low prior expectation (see Supplementary Material for a worked example of the relevant calculations). Node size encodes the enzyme's frequency in the reference genome set: larger nodes appear more commonly across reference genomes and carry a higher prior probability of presence. A legend (Figure 1F) maps each visual encoding to its meaning. Hovering over the legend highlights all nodes of that type in the graph view. A researcher can isolate detail by node size category, by flag type, or by best-path membership.

Node flags indicate annotation quality. Enzyme families often share evolutionary history, meaning that a genome sequence may score meaningfully against multiple HMM profiles for related families. KofamKOALA defines a reliability threshold for each enzyme family: a match above the threshold indicates that the sequence can be confidently assigned to that family rather than a related one [1]. The `!` marks an enzyme whose sequence match falls below this threshold, indicating that while the sequence shows similarity to the family, the match is not strong enough to rule out membership in a related family. The `?` marks an enzyme outcompeted by a higher-scoring hit on the same sequence. `x` indicates a dubious hit due to both. These flags make annotation quality visible without requiring the researcher to consult sequence-level analysis.

Each **reaction step** is separated by a dashed line (Figure 1E). Hovering reveals step probability before and after associate influence, making the network's contribution visible at the step level. The hover also shows the highest-probability enzyme for that step, with a link to the corresponding KEGG entry (Figure 1I).

Hovering over any node reveals the full evidence decomposition: enzyme identifier, matched gene sequence, posterior probability after associate influence, background frequency in the prior, sequence match score, E-value, annotation confidence, raw probability before influence, and the total shift contributed by associates (Figure 1G).

Edges between nodes encode the conditional co-occurrence of adjacent steps computed from the prior reference genomes. Edge thickness scales with co-occurrence frequency (thick edge indicates two enzymes are frequently observed together across reference genomes). A dashed edge indicates a possible enzymatic transition that has no observed co-occurrence (Figure 1F.2).

Influence Panel Clicking on any node in the module graph opens the influence panel (Figure 1H). The panel displays the probability shift of an enzyme contributed by valid reaction associates (other KOfams that co-occur with the clicked KOfam in more than 25% of reference genomes). The Supplementary Material works through this combination for Case Study 1 (M00022) using real values for a specific enzyme call. Associates are not limited to the current module but span a global network of KEGG modules in which the selected enzyme participates. A network diagram shows all influencing enzymes as nodes with edge thickness proportional to its contribution weight; green outlines mark associates from the current module while unmarked nodes indicate influence from associates found in some other module. A linked table lists each associate's co-occurrence count, conditional correlation, adjusted weight, and probability contribution. The influence panel further explains the split-circle node when an enzyme's posterior diverges from its sequence evidence, along with the associates' contribution.

3.4 Implementation Details

Implementation is a web application requiring no installation. The front end is built with D3.js for graph rendering and interactive visualization, with Dagre providing hierarchical graph layout. The graph generation pipeline is implemented in Python using the Lark parsing library and is available as a standalone command-line tool. The visualization is accessible at `https://blimmp.org`; the BLIMMP inference pipeline will be released as open source upon publication.

4 RESULTS AND DISCUSSION

We demonstrate the explorer using the *Pseudomonas fluorescens* SBW25 genome, a well-characterized soil bacterium whose metabolic capabilities are known, allowing biological expectations to serve as ground truth for the visualization. This application was chosen based on domain expertise of the senior author; analysis and

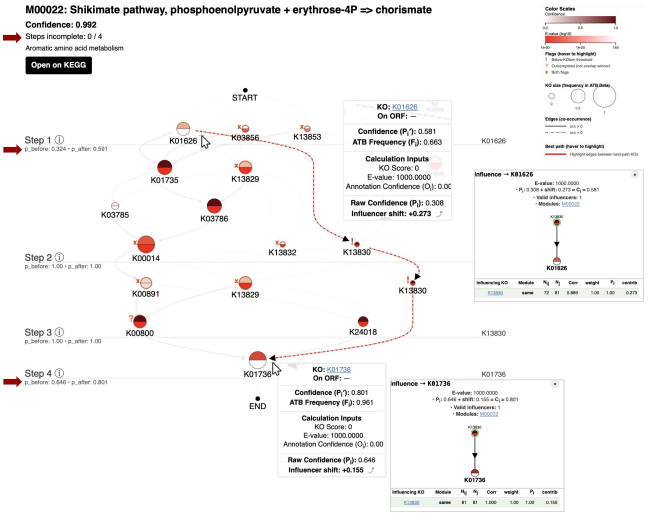

Figure 2: Case 1: Shikimate pathway (M00022) in a 70% complete *P. fluorescens* SBW25 assembly. Red arrows guide the reader to the two reaction steps with low confidence. Hover and influence diagrams are overlaid for the most probable enzymes catalysing those steps, K01626 and K01736.

use cases were developed with input from soil microbiome experts at the Pacific Northwest National Laboratory (PNNL).

4.1 Case 1: Recovering a Present Pathway in an Incomplete Genome

Genomes assembled from metagenomic sequencing are often incomplete. Rare community members may be sequenced too sparsely for full reconstruction, leaving some protein-coding genes unobserved even when the organism is present. To simulate this case, we randomly removed protein coding sequences from the complete SBW25 genome to produce a 70% complete assembly. We examined the shikimate pathway (M00022), an essential biosynthetic route whose presence is certain in SBW25. BLIMMP correctly infers the pathway as present with a confidence of 0.87, and we use this case to demonstrate how the visualization supports probing the evidence behind a probabilistic inference.

In the module graph (Figure 2) Steps 1 and 4 start with low step confidence (0.32 and 0.64), and increase after influence diffusion to 0.59 and 0.80. The best-path highlight traces K01626 (Step 1), K13830 (Steps 2+3), and K01736 (Step 4). The missing enzymes appear large with white bottoms and lightly colored top halves, indicating no sequence match but a fairly high prior probability. Hovering shows K01626 at 66.3% and K01736 at 96.1% background frequency across reference genomes. Clicking K01736 reveals K13830 as the sole influencer, co-occurring in all 81 reference genomes where K13830 was observed and confidently detected in this genome; this contributes a shift of +0.155, yielding a posterior of 0.801. Without the visualization, a researcher sees only the final scores, with no way to see that the entire lift above the raw annotation evidence traces to a single, perfectly co-occurring associate.

4.2 Case 2: Identifying a False Positive Using Annotation Flags

This case shows how a high-confidence call can mask unreliable annotation evidence. M00345, part of the Ribulose Monophosphate pathway, is not found in SBW25, yet scores a high confidence value (0.817). The graph shows every node at Steps 1 and 2 highlighted with flag outlines. The best-path legend highlights path K13831 (Step 1), K00850 (Step 2), and K01624 (Step 3). K01624 shows a red top and red bottom with no flags, indicating a strong, reliable sequence match. However, K13831 carries a ! flag: its sequence

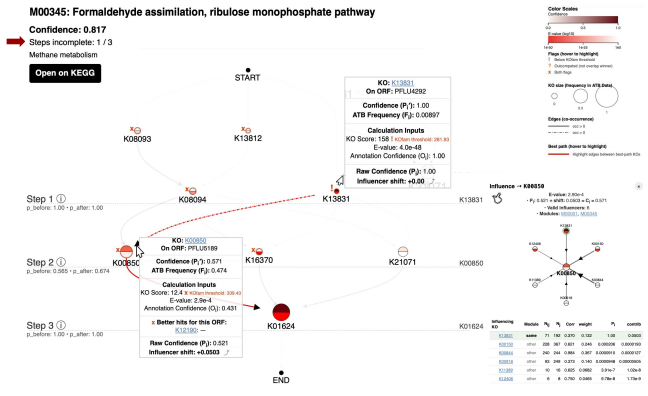

Figure 3: Case 2: RuMP pathway (M00345) in *P. fluorescens* SBW25. Hovering over flagged nodes reveals KofamKOALA threshold values for K13831 and K00850, exposing unreliable sequence matches that underlie the high module confidence score.

scores 158 against a KofamKOALA reliability threshold of 281.93. The score indicates a true similarity to the K13831 family, but falls below the threshold required to confidently distinguish it from related enzyme families. Step 2 presents an even weaker case: K00850 carries both ! and ? flags with a pale top half. Hovering over K00850 shows a score of 12.4 against a threshold of 339.43: statistically significant but unreliable similarity to the K00850 family. K00850 receives a minor probability shift (+0.05); its sole contributing influencer is K13831 from step 1. Its other six cross-module associates show strong co-occurrence (up to 0.984) but near-zero confidence in SBW25 and contribute nothing. The ? flag reveals that the same genome sequence matches another enzyme, K12190, at a score of 12.8, a marginally better hit from a related family, outcompeting K00850 on the same sequence. The 0.817 score implies a well-supported module, but the flags show every best-path node in the critical steps is below threshold or outcompeted, reflecting domain sharing rather than specific functional assignment.

4.3 Independent user feedback

To assess real-world utility beyond the case studies above, we gathered informal feedback from our domain expert who used BLIMMP-Explorer independently on genomes of interest. She reported two strengths; the visualization helped her follow the best-path enzyme selection among alternatives more quickly than tabular output and for intermediate-confidence module calls (below 0.8), the annotation flags helped her identify cases where sequence scores were high but did not pass KOfam reliability thresholds. This influences her interpretation of whether a step should be considered present. We view this as encouraging initial evidence; a structured evaluation across multiple users and genomes remains future work.

5 Conclusion

BLIMMP-Explorer makes the reasoning behind probabilistic metabolic module inference inspectable at the level of individual genes, annotation quality, and associated enzyme contributions. The use cases demonstrate interpretive tasks that were previously inaccessible without manual cross-referencing of multiple output files. Future work includes comparative views across multiple genomes, force-directed layouts for the influence network, and guided workflows for researchers without a strong bioinformatics background. We expect this design pattern to extend to other Bayesian models in microbiome research that similarly combine multiple qualitatively distinct evidence sources.

Acknowledgments

We thank Dr. Amy E. Zimmerman, Dr. William C. Nelson, and Dr. Jason E. McDermott from Pacific Northwest National Laboratory (PNNL) for their domain expertise and for providing the Pseudomonas fluorescens SBW25 genome used in the validation experiments.

References

[1] T. Aramaki, R. Blanc-Mathieu, H. Endo, K. Ohkubo, M. Kanehisa, S. Goto, and H. Ogata. KofamKOALA: KEGG ortholog assignment based on profile HMM and adaptive score threshold. *Bioinformatics*, 36(7):2251–2252, Apr. 2020. doi: 10.1093/bioinformatics/btz859

[2] R. Bourqui, L. Cottret, V. Lacroix, et al. Metabolic network visualization eliminating node redundance and preserving metabolic pathways. *BMC Systems Biology*, 1:29, 2007. doi: 10.1186/1752-0509-1-29

[3] dagrejs contributors. Dagre: Directed graph layout for JavaScript. `https://github.com/dagrejs/dagre`, 2025. Accessed: 2025-01-01.

[4] R. Eisenhofer, I. Odriozola, and A. Alberdi. Impact of microbial genome completeness on metagenomic functional inference. *ISME Communications*, 3(1):12, Feb. 2023. doi: 10.1038/s43705-023-00221-z

[5] A. Eren, Ö. Esen, C. Quince, J. Vineis, H. Morrison, M. Sogin, and T. Delmont. Anvi'o: an advanced analysis and visualization platform for 'omics data. *PeerJ*, 3:e1319, Oct. 2015. doi: 10.7717/peerj.1319

[6] P. G. Falkowski et al. The microbial engines that drive Earth's biogeochemical cycles. *Science*, 320:1034–1039, 2008. doi: 10.1126/science.1153213

[7] R. Finn, J. Clements, and S. Eddy. HMMER web server: interactive sequence similarity searching. *Nucleic Acids Research*, 39(Web Server issue):W29–W37, July 2011. doi: 10.1093/nar/gkr367

[8] J. Gabry, D. Simpson, A. Vehtari, M. Betancourt, and A. Gelman. Visualization in Bayesian workflow. *Journal of the Royal Statistical Society Series A: Statistics in Society*, 182(2):389–402, Feb. 2019. doi: 10.1111/rssa.12378

[9] M. Hunt, L. Lima, D. Anderson, et al. AllTheBacteria — all bacterial genomes assembled, available and searchable. *bioRxiv*, 2024. doi: 10.1101/2024.03.08.584059

[10] B. Junker, C. Klukas, and F. Schreiber. VANTED: a system for advanced data analysis and visualization in the context of biological networks. *BMC Bioinformatics*, 7:109, Mar. 2006. doi: 10.1186/1471-2105-7-109

[11] B. Kale, M. Sun, and M. Papka. The state of the art in visualizing dynamic multivariate networks. *Computer Graphics Forum*, 42:471–490, 2023. doi: 10.1111/cgf.14856

[12] M. Kanehisa and S. Goto. KEGG: kyoto encyclopedia of genes and genomes. *Nucleic Acids Research*, 28(1):27–30, Jan. 2000. doi: 10.1093/nar/28.1.27

[13] A. Kerren, K. Kucher, Y.-F. Li, and F. Schreiber. BioVis explorer: A visual guide for biological data visualization techniques. *PLOS ONE*, 12(11):e0187341, 2017. doi: 10.1371/journal.pone.0187341

[14] P. Kim and N. Price. Genetic co-occurrence network across sequenced microbes. *PLOS Computational Biology*, 7(12):e1002340, Dec. 2011. doi: 10.1371/journal.pcbi.1002340

[15] D. Naylor, R. McClure, and J. Jansson. Trends in microbial community composition and function by soil depth. *Microorganisms*, 10(3):540, Feb. 2022. doi: 10.3390/microorganisms10030540

[16] N. Sato, M. Uematsu, K. Fujimoto, S. Uematsu, and S. Imoto. ggkegg: analysis and visualization of KEGG data utilizing the grammar of graphics. *Bioinformatics*, 39(10):btad622, Oct. 2023. doi: 10.1093/bioinformatics/btad622

[17] C. Schulz et al. Probabilistic graph layout for uncertain network visualization. *IEEE Transactions on Visualization and Computer Graphics*, 23(01):531–540, Jan. 2017. doi: 10.1109/TVCG.2016.2598919

[18] M. Shaffer, M. A. Borton, B. B. McGivern, A. A. Zayed, S. L. La Rosa, L. M. Solden, P. Liu, A. B. Narrowe, J. Rodríguez-Ramos, B. Bolduc, M. C. Gazitúa, R. A. Daly, G. J. Smith, D. R. Vik, P. B. Pope, M. B. Sullivan, S. Roux, and K. C. Wrighton. DRAM for distilling microbial metabolism to automate the curation of microbiome function. *Nucleic Acids Research*, 48(16):8883–8900, Sept. 2020. doi: 10.1093/nar/gkaa621

[19] E. Shinan and contributors. Lark: A parsing toolkit for python. `https://github.com/lark-parser/lark`, 2017. Accessed: 2025-01-01.

[20] K. Sugiyama, S. Tagawa, and M. Toda. Methods for visual understanding of hierarchical system structures. *IEEE Transactions on Systems, Man, and Cybernetics*, (2):109–125, 1981. doi: 10.1109/TSMC.1981.4308636

[21] W. Tian, L. Zhang, M. Taşan, F. Gibbons, O. King, J. Park, Z. Wunderlich, J. Cherry, and F. Roth. Combining guilt-by-association and guilt-by-profiling to predict *Saccharomyces cerevisiae* gene function. *Genome Biology*, 9(Suppl 1):S7, 2008. doi: 10.1186/gb-2008-9-s1-s7

[22] C. Vehlow, J. Hasenauer, A. Kramer, A. Raue, S. Hug, J. Timmer, N. Radde, F. Theis, and D. Weiskopf. iVUN: interactive visualization of uncertain biochemical reaction networks. *BMC Bioinformatics*, 14(Suppl 19):S2, 2013. doi: 10.1186/1471-2105-14-S19-S2

[23] D. Zhang, E. Adar, and J. Hullman. Visualizing uncertainty in probabilistic graphs with network hypothetical outcome plots (NetHOPs). *IEEE Transactions on Visualization and Computer Graphics*, 28(1):443–453, Jan. 2022. doi: 10.1109/TVCG.2021.3114679

[24] Z. Zhou, P. Tran, A. Breister, et al. METABOLIC: high-throughput profiling of microbial genomes for functional traits, metabolism, biogeochemistry, and community-scale functional networks. *Microbiome*, 10:33, 2022. doi: 10.1186/s40168-021-01213-8

Supplemental Material

An extended discussion of Case 1 (M00022, the Shikimate pathway) with evidence-decomposition calculations are available at: `https://osf.io/qwc3x/overview?view_only=1b431ef6f6ad4249b0bf3fd4da420c40`