
A Prototype Biovisualization of Gene Expression Changes in *Mycobacterium tuberculosis* Metabolic Pathways

Major Qualifying Project

Written By:

ADRIAN ORSZULAK

Advisor:

DR. LANE HARRISON



A Major Qualifying Project
WORCESTER POLYTECHNIC INSTITUTE

Submitted to the Faculty of the Worcester Polytechnic
Institute in partial fulfillment of the requirements for the
Degree of Bachelor of Science in Bioinformatics and
Computational Biology.

AUGUST 25, 2021 - APRIL 28, 2022

This report represents the work of one or more WPI undergraduate students submitted to the faculty as evidence for completion of a degree requirement. WPI routinely publishes these reports on the web without editorial or peer review.

ABSTRACT

Metabolic pathways are complex, multi-structured organizations of proteins and metabolites that promote the continuation and survival of a living system. Through regulation of the proteins that compose these pathways, a robust system enables greater survival in a variety of stresses and conditions. *Mycobacterium tuberculosis*, the causative agent of tuberculosis, is one such organism. This makes treatment of the pathogen difficult. Generating and using biovisualizations is important for culminating a greater understanding of *M. tuberculosis* and its mechanisms of regulating gene expression. While there are an abundance of biovisualizations mapping metabolic pathways and protein-protein interactions, there is little work involving the ability to visualize changes in protein expression and impacts of regulations on systems. Additionally, many of these visualizations and other existing visualization tools are based on reference pathways, and are not specific to *M. tuberculosis*. In this paper, a prototype biovisualization of *M. tuberculosis* metabolic pathways was constructed consisting of three panels: one visualizing metabolic pathways, one displaying protein-protein interactions, and one showing changes in gene expression from a user uploaded file. The tasks and design components were made based on interviews with *M. tuberculosis* researchers. This work highlights a prospective idea for future metabolic biovisualization work.

ACKNOWLEDGEMENTS

I would like to thank my Bioinformatics and Computational Biology advisor Dr. Lane Harrison. His support, wealth of knowledge, and guidance during the course of this project was invaluable.

I would also like to thank my Biology and Biotechnology advisor Dr. Scarlet Shell. Her support during the project and feedback about the final prototype biovisualization were incredibly helpful.

I would like to thank and extend my appreciation to Hilson Shrestha, Yiren Ding, and all graduate students part of the Worcester Polytechnic Institute data visualization community. Their help in troubleshooting errors and problems during the creation of the prototype biovisualization was invaluable.

I would like to extend my thanks and appreciation for all interviews during the goals and tasks analysis portion of this MQP. Your time, ideas, and feedback were all greatly appreciated.

Finally, I would also like to thank members of the Shell Lab. Your feedback during lab meetings and presentations over the course of the project's development was very useful.

TABLE OF CONTENTS

	Page
List of Figures	v
List of Tables	vi
1 Introduction	1
2 Background	4
2.1 Importance of Visualizing Metabolic Pathways and Gene Expression Changes . . .	4
2.2 Goals and Task Taxonomy Regarding Metabolic Pathway Visualizations	5
2.3 Existing Metabolic Pathway Databases for <i>Mycobacterium tuberculosis</i>	6
2.4 Existing Metabolic Pathway Visualizations of <i>Mycobacterium tuberculosis</i>	8
2.5 Existing Metabolic Pathway Visualization Tools	11
3 Experimental Setup	15
3.1 Compiling the Dataset for the Prototype Metabolic Pathway Biovisualization . . .	15
3.2 Understanding and Compiling the Important Goals and Tasks for the Prototype Metabolic Pathway Biovisualization	16
3.3 Creating and Improving on the Prototype Metabolic Pathway Biovisualization . .	16
4 Results	18
4.1 Building an Understanding Regarding the Important Tasks and Goals for a Proto- type Metabolic Pathway Biovisualization of <i>Mycobacterium tuberculosis</i>	18
4.1.1 Key Actions	18
4.1.2 Important Data Features	19
4.1.3 Displaying the Data	20
4.2 Compiling a Dataset for the Prototype Metabolic Pathway Biovisualization of <i>Mycobacterium tuberculosis</i>	21
4.3 The Prototype Metabolic Pathway Biovisualization of <i>Mycobacterium tuberculosis</i>	22
4.3.1 The Metabolic Pathway Panel	23
4.3.2 The Protein-Protein Interaction Panel	24
4.3.3 The Regulation Data Panel	26

TABLE OF CONTENTS

5	Discussion	29
5.1	Data Limitations and Future Applications	29
5.2	Interview Limitations and Future Applications	30
5.3	Biovisualization Limitations and Future Applications	31
6	Conclusions	33
	Appendix A: Interview Questions	i
	Appendix B: Informed Oral Consent Form	iii
	Bibliography	iv

LIST OF FIGURES

FIGURE	Page
2.1 A summary of the features present and not present in existing databases of <i>M. tuberculosis</i> containing metabolic pathway information	8
2.2 A summary of the features present and not present in existing biovisualizations of <i>M. tuberculosis</i> metabolic pathways	10
2.3 A summary of the features present and not present in existing tools that visualize and inform researchers about metabolic pathways.	14
4.1 A representation of the data involved in each panel of the biovisualization.	21
4.2 A sample image of the metabolic pathway panel	24
4.3 A sample image of the protein-protein interaction panel	25
4.4 A sample image of the regulation data panel	27
4.5 A sample image of the file upload mechanism tied to the regulation data panel	27
4.6 A sample image of the impacts of the file upload on the metabolic pathway panel . . .	28

LIST OF TABLES

TABLE	Page
3.1 Researchers Interviewed	16

INTRODUCTION

Researchers in the biological sciences have access to vast amounts of data regarding specific systems and models for study. Metabolism is and the study of metabolic pathways is one such area of interest. Metabolic pathways form the basis for understanding how organisms function and survive and respond to environmental changes[1]. These metabolic pathways are outlined by many concurrent and sequential reactions that occur within cells. These biochemical networks are difficult to visualize because of the vast amount of associated data and the impacts small changes have on the entire pathway or set of pathways[1, 2]. Metabolism and the study of metabolic pathways is, thus, one area of biological intrigue in which visualizations have the greatest need and impact.

To support this intricate web of reactions, most if not all proteins are regulated in their activity in some way. To better understand widespread changes in global gene expression, many researchers use RNA sequencing (RNA-seq) to note changes in expression in order to understand how an organism or system reacts to changes in the environment. Using visualizations, these biological researchers can examine and study the extensive amount of metabolic data and RNA-seq data to visualize where changes and adaptations are occurring. This leads to an enhanced understanding and knowledge to drive further research.

One such research community is the one that studies *Mycobacterium tuberculosis*. *Mycobacterium tuberculosis* is the pathogen that causes tuberculosis in humans [3]. With 1.3 million deaths attributed to *Mycobacterium tuberculosis* in 2020 alone, it is one of the leading causes of death by an infectious disease in the world [3]. There is an extensive amount of work and research being devoted to understanding this pathogen, its mechanism of persistence, and how to treat it [3]. Given the importance of metabolism and its regulation, visualizing such data would

prove useful for these researchers as they work to better study and find means to combat this pathogen. Using visualizations enhances the researchers understanding and to draw inferences regarding their experimental systems and models for a greater push of scientific understanding.

Dang, et al. (2017) and Murray, et al. (2016) have addressed important task taxonomies and design elements for creating protein-protein interaction and metabolic pathway visualizations. These findings have not been examined for researchers studying *Mycobacterium tuberculosis* specifically. Additionally, there is an extensive amount of bioinformatics data and information regarding metabolic pathways that has been generated and are commonly used. Many of these commonly used resources are limited in their representations of specific organisms, lack of gene expression feature change, or limited by some other external factors (Figure 2.1 and figure 2.2 and Figure 2.3). These are key components for many researchers who are trying to elucidate the role of a particular protein or note differences between two or more strains.

Building on the work of Murray, et al. (2016) and Dang, et al. (2017), interviews with researchers studying *M. tuberculosis* and its model organism *Mycobacterium smegmatis* were conducted to build a better understanding of the tasks important for a prototype metabolic visualization of *M. tuberculosis*. The responses were recorded and specific tasks that a visualization should accomplish were identified. In regard to the prototype metabolic visualization, the final prototype was created in React.js using d3.js, node.js, and express.js. The final visualization consists of three panels. The first panel visualizes one pathway selected by the user as either a network of connecting proteins or a more traditional view of protein to molecule to protein connections. Transcription factors and other pathway connections are toggleable. The second panel is displayed once a protein from the first is selected, a protein's protein-protein interactions are displayed and interaction designated by a color legend. The third and final panel allows a user to upload a file containing a protein or transcription factor and their associated fold change from RNA-seq data. This data is compiled into tables displayed on the panel and is incorporates visually within the first, pathway panel. Broadly, the aim of this work was two-fold. The first goal was to compile a list of the important tasks for a metabolic pathway visualization of *M. tuberculosis* from researchers working to understand the pathogen. The second goal was to create a prototype visualization of the metabolic pathway visualization of *M. tuberculosis*.

Thus, the contributions of this project are as follow:

- A compiled data set of two pathways and one collected set of proteins grouped together by similar function with a number of dimensions created specifically for this visualization
- A list of notable tasks relevant to a metabolic pathway visualization of *M. tuberculosis* from researchers studying *Mycobacterium tuberculosis*
- A prototype biovisualization of the metabolic pathways of *Mycobacterium tuberculosis* using d3.js that includes:

-
- The ability to visualize two pathways and one collected set of proteins grouped together by similar function
 - The ability to visualize the protein-protein interactions of each protein in the visualized pathway
 - The ability to upload a file containing a list of proteins or transcription factors and their associated fold changes of expression generated from RNA-seq data that is then incorporated onto the pathways being visualized

2.1 Importance of Visualizing Metabolic Pathways and Gene Expression Changes

Metabolic pathways are complex, biological systems that are central in the survival of a variety of organisms. Each pathway consists of several different interacting proteins and metabolites. Researchers organize these components into pathways to better understand their relation and map their impacts on living systems. These pathways are multivariate and contain numerous associated datapoints, inherent through the nature of these living systems and expressed in the multivariate statistical analysis that is taken to analyze them[1, 4, 5]. When studying the impacts of different protein-coding genes, groups of protein-coding genes, or their gene products, expression is often varied to determine their role in their specific living system. These changes may have great implications on the system and its survival by impacting one or more metabolic pathways. These changes in protein and metabolite concentrations are known to easily perturb and affect metabolic systems[2]. In the context of gene expression changes, living systems are easily and potentially significantly affected by those changes. Understanding those changes, in a neighborhood and global context, is important to researchers. In the context of *M. tuberculosis*, the changes to its metabolism are central in the bacteria's ability to cause disease, survive various stressors, and become tolerant to drug exposure[6]. Given the multivariate nature, number of interactions, and significance of changes that are present in a metabolic pathway, visualizing metabolic pathways and their effects will aid researchers in their efforts to better understand the systems they study. For researchers studying *M. tuberculosis*, a visualization tool that can map gene expression changes to the metabolic pathways of the organism would understand the

various aspects of *M. tuberculosis* pathogenesis, persistence, and survivability.

2.2 Goals and Task Taxonomy Regarding Metabolic Pathway Visualizations

It is important to know what members of the scientific community and those seeking to use this visualization in order to understand the goals and tasks of a visualization. In a general approach, Lee, et al. (2006) describes a number of tasks for different types of network visualizations. In an attribute-based network, nodes should have specific values that allow for users to act on one or more of the following tasks: filtering the nodes, computing data information, determining a range, or characterize the distribution of nodes for and over a given dataset[7]. For metabolic pathway biovisualizations, it could be argued that noting correlations could also be vital for attribute-based nodes, using Lee, et al. (2006) as a reference. Links should be specific and defined by some relationship between one or more nodes[7]. This general view is, as mentioned, not specific to biovisualizations or metabolic pathway biovisualizations.

A couple of studies have previously examined a list of task taxonomies related to metabolic pathway visualizations and protein-protein interaction visualizations. In the former, Murray, et al. (2016) interviewed many different biomedical researchers to create a list of visualization task taxonomy for metabolic pathway biovisualizations. Highlighted from their analysis are three categories of tasks: attribute associated, relationship associated, and data curation associated[8]. Attribute associated tasks focused on the multivariate nature of the data, the ability to compare data between pathways, denoting experimental evidence for each entry, and noting the level of confidence between relationships[8]. Relationship associated tasks highlighted an understanding of the relationships and interactions between different pathway components in terms of the reaction, directionality, grouping, casualty, and feedback[8]. Finally, data curation tasks focus on modifying and updating the dataset in order to match new findings, discoveries, and clarification of protein functions and interactions in order to maintain accuracy in the biovisualization[8]. Current implementations or methods for implementing each of these tasks was provided[8].

In the latter, protein-protein interaction, Dang, et al. (2017) highlighted some key tasks for a solely protein-protein interaction visualization in cancer cells. Though specific to cancerous cells, these insights can extend to protein-protein interaction visualizations related to biomedical research data. These tasks focused on presenting accurate and relevant data that allowed users to investigate protein interactions and protein neighborhoods. Additionally, pointing out relevant and conflicting literature was a key task identified in their work[9]. Overall, these previous works serve as a great resource for building visualizations incorporating both metabolic pathways and

protein-protein interactions. No specific insight has been examined in researchers studying *M. tuberculosis* or other bacterial species. Investigating these individuals could provide insights into tasks, goals, and features specific to this and other related communities of researchers and their research goals.

2.3 Existing Metabolic Pathway Databases for *Mycobacterium tuberculosis*

There are a number of databases that curate and store metabolic information for *M. tuberculosis*. When discussing proteins, the UniProtKB database, a collection of protein entries for most organisms, is commonly used[10, 11]. Each protein entry lists a number of structural, functional, and taxonomic data associated with the protein of interest[10, 11]. This makes identifying proteins easy. The database has many annotations to other databases. For example, protein-protein interactions for each entry direct the user to the STRING network[10, 11]. There are not many features specific to visualizing pathways or noting gene expression changes with this database.

The Systems Bio MtB database is a large-scale, multi-omics curated dataset that seeks to generate a predicted network for gene regulation in *M. tuberculosis*, especially for host-pathogen interactions[12]. Each gene is given an entry that denotes its sequence data, regulation and co-regulation overview, and existing quantitative proteomics data [12]. Annotations to other databases are also included for metabolic pathway and protein-protein interaction data[12]. The most notable feature of this database is its integration of regulation data and information between transcription factors and proteins. Users are able to better define where changes of gene expression can occur and which specific regulation elements, whether that be proteins or transcription factors, are likely to be causing any noted changes[12]. No user data is able to be uploaded to entries within the database.

The Pathosystems Resource Integration Center (PATRIC), is a database specifically created for collecting, organizing, and integrating data for use in biomedical research against bacterial infectious diseases[13]. Taking a more global approach, this database highlights large-scale transcriptomics, protein-protein interactions, sequence analysis, protein family grouping, and metabolic pathways[13]. Using experimental evidence, interactions between proteins and their effects are noted. Large-scale transcriptomics data highlights changes in expression indicated through the literature. No direct changes in expression can be manipulated or inputted to any of the sites directly. The metabolic pathway collection uses what appear to be hand-drawn reference pathways designed by the Kyoto Encyclopedia of Genes and Genomes Database (KEGG) onto which proteins known to be found in *M. tuberculosis* are mapped[13]. Proteins can then be

2.3. EXISTING METABOLIC PATHWAY DATABASES FOR *MYCOBACTERIUM TUBERCULOSIS*

selected, but the goal and main function of this database appears to be storing and presenting the data to the user.

The BioCyc database is the most extensive of the databases examined, in regard to data and features. Proteins from an extensive number of organisms, *M. tuberculosis* included, can be investigated for sequence, function, and structural data[14]. In conjunction with the database, BioCyc includes a metabolic map for a number of organisms[14]. This map organizes reactions into specific pathways where applicable[14]. As opposed to other visualizations discussed, the reference metabolic map is based on that particular organism rather than one, general reference metabolic pathway where applicable[14]. In this particular network, nodes of varying shapes represent different components of a reaction whereas links represent the reactions themselves[14]. A zoom feature is used to examine the metabolic map[14]. Using the Omics tool, users are able to upload a file, specifying gene or protein name and a numerical value, to see the degree of gene expression change for metabolic experiments particularly[14, 15]. Limiting a user's investigation, however, is a subscription to use the database[14]. The MetaCyc database, a derivative of the BioCyc database, has many features specific to metabolism and metabolic pathways[16]. From this database, protein entries can be examined to obtain sequence, function, and structural data about a specific protein. This database does not feature a metabolic map or similar functionality.

In summary, these databases are useful and robust collections of curated data that help researchers easily reference that available data. Many of the investigated databases have some degree of specificity to *M. tuberculosis* and are useful to easily find proteins of interest. Two, specifically BioCyc and PATRIC, present metabolic pathways in a visualized way to some degree[13, 14]. In regard to BioCyc, this database incorporates a metabolic map tool that allows users the ability to note experimental gene expression changes from user inputted data, all based on a *M. tuberculosis* reference pathway[14]. BioCyc, though a useful tool that does address metabolic gene expression changes, is often limited in use for its not complete representation of *M. tuberculosis* metabolic pathways and subscription based service[14]. Aside from BioCyc, other investigated databases are limited in their representation of protein-protein interactions and gene expression changes. These conclusions are summarized in Figure 2.1. Thus, facilitating the creation of a prototype biovisualization to denote differences in gene expression in the *M. tuberculosis* would serve to benefit the larger tuberculosis research community. Additionally, creating a prototype biovisualization would present the novel idea incorporating a metabolic pathway visualization with the ability to denote gene expression changes and protein-protein interactions specific to *M. tuberculosis*.

	Visualize Pathway	Specific to <i>M. tuberculosis</i>	Identify Particular Proteins	Allow Uploading Data	Indicate Impacts of Gene Expression Changes	Protein-Protein Interactions
Systems Bio MtB						
BioCyc						
MetaCyc						
PATRIC						
UniProtKB						

Figure 2.1: A summary of the features present and not present in existing databases of *M. tuberculosis* containing metabolic pathway information. The “Visualize the Pathway” feature refers to the database containing a visualization of a metabolic pathway. The “Specific to *M. tuberculosis*” feature refers to the database containing data specific and inherent to *M. tuberculosis* and not simply orthologous information. The “Identify Particular Proteins” feature refers to the database allowing users to find and search for proteins or pathways. The “Allow Uploading Data” feature refers to the database allowing users to enter data in some form, regardless if that data is specific to gene expression datasets or simply a list of proteins to highlight. The “Indicate Impacts of Gene Expression Changes” feature refers to the database providing information related to gene expression changes, with or without user file uploading. The “Protein-Protein Interactions” feature refers to the database indicating interactions between proteins and the nature of those interactions. Features that are present in a specific biovisualization tool are noted in green. Features that are present in a specific biovisualization tool to some capacity are noted in brown. Features that are not present in a specific biovisualization tool are white. Any component listed as annotation to another database is not included as the database containing the feature or containing the feature to some capacity.

2.4 Existing Metabolic Pathway Visualizations of *Mycobacterium tuberculosis*

There are a handful of different, dedicated visualizations of metabolic pathways that already exist and include information and data for *M. tuberculosis*. The KEGG database contains one such visualization[17–19]. Using a base reference map of metabolic pathways, the KEGG visualization highlights proteins or orthologous proteins of a specific organism that construct the metabolic pathways for that organism or particular isolate of that organism[17–19]. The user is then able to highlight specific pathways of interest and specific proteins[17–19]. Clicking on a node or link in the pathway directs the user to another page with more details regarding a particular protein or metabolite[17–19]. The KEGG biovisualization is robust and extensive in its representation of metabolic pathways for a number of organisms, but lacks specificity as a result[17–19]. Certain

2.4. EXISTING METABOLIC PATHWAY VISUALIZATIONS OF *MYCOBACTERIUM TUBERCULOSIS*

proteins and protein collections specific to *M. tuberculosis* are absent from the base reference map and are not represented. Additionally, while the pathway does allow users to upload data, this upload feature is restricted to highlighting specific proteins in the reference protein only[17–19]. No changes nor notes of gene expression changes can be submitted alongside a protein in question.

Another such biovisualization is found in the BRENDA database[20]. BRENDA uses its extensive source of data and a reference map of the metabolic pathways in order to construct the metabolism of many different organisms[20]. A number of different strains and isolates of *M. tuberculosis* are represented within this metabolic pathways biovisualization[20]. Using a zoom feature, a user is able to get a greater insight into each specific pathway, being able to denote proteins, molecules, and some cofactors involved in the many reactions of a metabolic pathway[20]. Once an organism is selected, the visualization generates a coverage score and highlights proteins or orthologous proteins within the reference pathways that are found in the organism of interest[20]. Clicking on a protein node directs the user to another page with more information specific to that protein of interest[20]. While extensive, the BRENDA biovisualization trades specificity for scope per its intended goal, and is not specific to *M. tuberculosis* or any particular organism[20]. Thus, certain protein or protein collections are not present in the visualization. The upload feature allows users to highlight specific proteins only[20].

While not specific to displaying the metabolic pathways of any particular organism, the STRING and STITCH visualizations provide an interesting viewpoint into the metabolism of an organism of interest through protein-protein interactions and protein-molecule interactions respectively[21–23]. In STRING, a user is able to enter a protein of interest within a particular organism to generate a network of protein-protein interactions at varying levels of confidence[22, 23]. In STITCH, a user is able to enter a protein of interest within a particular organism to generate a network of protein-protein interactions and protein-molecule interactions at varying levels of confidence[21]. Both visualizations generate these networks as representations of the confidence of these interactions rather than the specificity[21–23]. In the STRING visualization, users are able to identify physical interactions, the formation of a complex, and functional interactions, some unspecified interaction, between proteins[22, 23]. In the STITCH visualization, users are able to note the molecular action and binding affinity of protein-protein and protein-molecule interactions[21]. Inherent in their representation and robust dataset, these visualizations are able to show metabolic pathways through the interactions of proteins, but not in a style commonly associated with metabolic pathway representations[21–23]. Furthermore, this dataset allows the visualization to be robust in its representation of many different proteins for a variety of organisms without being made specific for any one, including *M. tuberculosis*[21–23]. As a representation of interactions between proteins and molecules respectively, STRING and STITCH have some degree of indicating gene expression changes using outside information, but

	Visualize Pathway	Specific to <i>M. tuberculosis</i>	Identify Particular Proteins	Allow Uploading Data	Indicate Impacts of Gene Expression Changes	Protein-Protein Interactions
KEGG						
BRENDA						
STRING						
STITCH						

Figure 2.2: A summary of the features present and not present in existing biovisualizations of *M. tuberculosis* metabolic pathways. The “Visualize the Pathway” feature refers to whether the visualization visualizes a metabolic pathway. The “Specific to *M. tuberculosis*” feature refers to the visualization containing data specific and inherent to *M. tuberculosis* and not simply orthologous information. The “Identify Particular Proteins” feature refers to the visualization allowing users to find and search for proteins or pathways. The “Allow Uploading Data” feature refers to the visualization allowing users to enter data in some form, regardless if that data is specific to gene expression datasets or simply a list of proteins to highlight. The “Indicate Impacts of Gene Expression Changes” feature refers to the visualization providing information related to gene expression changes, with or without user file uploading. The “Protein-Protein Interactions” feature refers to the visualization indicating interactions between proteins and the nature of those interactions. Features that are present in a specific biovisualization tool are noted in green. Features that are present in a specific biovisualization tool to some capacity are noted in brown. Features that are not present in a specific biovisualization tool are white. Any component listed as annotation to another database or visualization is not included as the visualization containing the feature or containing the feature to some capacity.

have no means to represent those changes within the visualization[21–23]. Likewise, no ability to upload data is present[21–23].

In summary, these biovisualizations are extensive and robust in the information and detail that they present regarding the metabolism and protein-protein interactions of a variety of organisms, including *M. tuberculosis*. The metabolic visualizations created in KEGG and BRENDA are based on a reference pathway that, while effective at displaying the majority of metabolism in *M. tuberculosis*, is not specific to that or any organism[17–20]. Thus, certain proteins and pathways specific to the organism are not able to be found within this visualization. While STRING and STITCH do include many proteins specific to *M. tuberculosis*, their goal is not to indicate the whole metabolic pathways of these proteins[21–23]. STRING and STITCH are able to indicate impacts of gene expression changes to some capacity, but the goal of these biovisualizations does not allow the user to draw inferences about how these gene expression changes impact the pathway as a whole or other pathways[21–23]. Finally, aside from the ability to highlight a set of proteins if those proteins are found in the KEGG and BRENDA biovisualizations, these visualizations do

not allow the user to upload data or datasets denoting gene expression changes[17–20]. These conclusions are succinctly detailed in Figure 2.2. Thus, creating a prototype biovisualization of *M. tuberculosis* metabolic pathways would address a novel idea of incorporating and visualizing gene expression changes in a pathway alongside protein-protein interaction data.

2.5 Existing Metabolic Pathway Visualization Tools

While not necessarily specific to *M. tuberculosis*, a number of visualization tools have been developed to better map, understand, and draw inferences from metabolic visualization data. A select number of other visualization tools referenced in Murray, et al. (2016) and review of the literature were chosen to understand the capabilities of these tools for use with *M. tuberculosis*. VisANT is one such visualization tool that utilizes networks to represent large scale datasets for researchers to draw inferences regarding the nature of the interactions and gene expression ratios[24–26]. Interactions with the node can reveal further increasing complexity between the interactions present[24–26]. Edgers represent different experimental methods and evidence through use of a color visual channel[24–26]. Many annotations allow the user to easily find other entries and further information more easily, especially when paired with the Predictitome database[24–26]. Layout and graph operations are determined from user input on various filters[24–26]. While not displaying metabolic pathways explicitly, the inherent displaying of interaction networks shows the metabolic relationship of proteins[24–26]. No explicit feature for denoting gene expression changes is found within this visualization tool[24–27].

Cytoscape is another visualization tool originally designed for assessing and analyzing metabolic data[28]. Using the various layout tools, different networks with different degrees of details can be shown[28]. This ranges from displaying pathways based from a public source database, protein interaction networks, and protein-protein interactions[28]. Filters are present in order to reduce or increase the magnitude of data displayed at any one time[28]. Furthermore, the users are able to upload gene expression data and draw inferences about the changes in protein expression. There is no specific data to *M. tuberculosis* regarding metabolomics data, but access to specific *M. tuberculosis* sequence and genome data is present[28].

PTools is the most extensive and robust of the available metabolic visualization tools discussed. This tool is used in conjunction with and found as a part of the BioCyc database, as aforementioned[14, 15]. Given this relationship, PTools reconstructs a metabolic pathway using BioCyc for a variety of organisms, including *M. tuberculosis*[15]. As noted with BioCyc, the representation is not entirely complete in regard to *M. tuberculosis*[14]. Using this tool, specific proteins, molecules, and reactions can be found within organized pathways, where applicable, and

in intracellular or extracellular space as they would exist in a living system[15]. A highlighting feature allows users to identify proteins easily within the visualization tool[15]. Furthermore, an omics aspect to the visualization tool permits the user to upload gene expression data and note changes in expression on the visualized pathway[15]. A scale is used to be able to denote differences in expression changes[15]. As previously mentioned with regard to BioCyc, this Omics tool does require a subscription for continued use[14].

Another metabolism visualization tool, BiologicalNetworks, incorporates the ability to upload and analyze microarray gene expression data[27]. The BiologicalNetworks visualization tool visualizes metabolic pathways, protein-protein interactions, and protein-DNA interactions through a variety of networks[27]. Through filtering features and hierarchical organization of networks within the tool, users are able to specify what is exactly visualized[27]. Using those visualized networks, users are able to find and search for proteins, pathways, common targets, common regulators of gene expression, and intersections with other pathways[27]. Furthermore, this visualization tool supports the upload of user gene expression data to reveal the impact of gene expression changes from microarray data in a static or dynamic display [27]. Capabilities to analyze and compare datasets of gene expression changes are included in this tool allow users to better understand and define the relevance of their gene expression datasets[27]. The BiologicalNetworks tool has no inherent specificity to *M. tuberculosis* provided its sources of data integration[27].

The GeneVis visualization tool takes a different approach to understand genetic regulatory networks in prokaryotic organisms using a simulation[29]. Using a circle to represent a bacterial chromosome, smaller circular nodes representing protein-coding genes[29]. About each protein-coding gene node, activity state of a gene and whether regulatory proteins are bound to the gene are found[29]. Using this approach, regulatory and activity dynamics of these protein-coding genes are visualized[29]. The degree of detail and exact genes, with their associated protein products, are able to be toggled through the use of different views and lenses [29]. Interactions between gene-coding proteins and protein products amass to display a hierarchical system of regulation based on the chromosomal ring[29]. The impacts of gene expression shown are limited to effects of selected proteins products on genes already within the visualization as opposed to a more global approach [29]. The exact nature of a global application is not discussed. While designed for prokaryotic organisms, *Escherichia coli* is the organism of interest explored within the initial publication of this visualization tool[29]. No known work has been done using data from *M. tuberculosis* with GeneVis.

The final metabolic visualization tool discussed in this report is VANTED. The VANTED tool allows researchers to visualize biochemical data regarding proteins and metabolites in the context of metabolic pathways[30]. Using the KEGG pathways, another independent pathway

file, or a pathway built by the tool itself from the data, the visualization tool is able to display the metabolic pathways of an organism [30]. On the final constructed pathway, user experimental data can be uploaded and mapped to the proteins or metabolites of interest[30]. For each specified experiment or condition, the metabolite or protein concentration is displayed at each node which allows for easy comparison within a pathway[30]. Plotting expression across a similar timeframe can also be done within a pathway using a neuronal-network algorithm: the self-organizing map (SOM)[30]. Finally, the tool allows users to perform a correlation analysis between metabolite expression and protein expression [30]. Though there is a lack of specificity with regard to the *M. tuberculosis* metabolic pathway, VANTED may construct its own pathway if the KEGG pathway is not complete[17–19, 30]. The tool is able to represent these pathways with greater specificity to *M. tuberculosis* should such data be uploaded and used.

In summary, existing visualization tools are quite robust in their data integration and functionality. Alongside visualizing metabolic pathways to some degree, the majority of these tools are able to show indications of gene expression changes from a user created dataset. By proxy, identifying particular proteins is an inherent and basic function within all these tools. The use of different filters and lenses specific to each tool make searching easy and allow users to draw conclusions and inferences while still maintaining the multivariate nature of the data. Additionally, all visualization tools discussed provide some degree of statistical analysis to them. These visualization tools, however, are not specific to *M. tuberculosis*, providing a niche for another visualization tool designed for the *M. tuberculosis* research community. These conclusions are summarized in Figure 2.3. Through this project, integrating the aspects of a metabolic pathway visualization, a protein-protein interaction visualization, and noting changes of gene expression for the *M. tuberculosis* research community is the main goal. Many of these features and design choices, especially with regard to filtering and expressing data, should be carried into a biovisualization that this project seeks to design and create.

	Visualize Pathway	Specific to <i>M. tuberculosis</i>	Identify Particular Proteins	Allow Uploading Data	Indicate Impacts of Gene Expression Changes	Protein-Protein Interactions
VisANT						
Cytoscape						
PTools						
BiologicalNetworks						
GeneVis						
VANTED						

Figure 2.3: A summary of the features present and not present in existing tools that visualize and inform researchers about metabolic pathways. The “Visualize the Pathway” feature refers to the tool containing a visualization of a metabolic pathway. The “Specific to *M. tuberculosis*” feature refers to the tool containing data specific and inherent to *M. tuberculosis* and not simply orthologous information. The “Identify Particular Proteins” feature refers to the tool allowing users to find and search for proteins or pathways. The “Allow Uploading Data” feature refers to the tool allowing users to enter data in some form, regardless if that data is specific to gene expression datasets or simply a list of proteins to highlight. The “Indicate Impacts of Gene Expression Changes” feature refers to the tool providing information related to gene expression changes, with or without user file uploading. The “Protein-Protein Interactions” feature refers to the tool indicating interactions between proteins and the nature of those interactions. Features that are present in a specific biovisualization tool are noted in green. Features that are present in a specific biovisualization tool to some capacity are noted in brown. Features that are not present in a specific biovisualization tool are white. Any component listed as annotation to another database or visualization is not included as the visualization containing the feature or containing the feature to some capacity.

EXPERIMENTAL SETUP

3.1 Compiling the Dataset for the Prototype Metabolic Pathway Biovisualization

The pathway and other related data was gathered from a number of databases. The KEGG, PATRIC, and Mycobrowser databases were used to obtain the protein name, substrate and product details, protein connections, branch point status, and pathway connections[13, 17–19, 31]. No hypothetical proteins were included in this dataset. Protein-protein interaction data was gathered using the STRING and STITCH databases[21–23]. The MTB Network Portal database was used to obtain transcription factor information and to determine the gene name of each protein[12]. Only those interactions with the highest confidence were included. This data was compiled into a csv file for use in the final prototype visualization. The final dataset can be found in the GitHub repository alongside the final prototype visualization (<https://github.com/ao-joker/V2-BCB-MQP-BioVis>). All data recorded was from the *M. tuberculosis* H37Rv strain.

Researcher Title and Research Areas
Title: Bioinformatics Graduate Research Assistant Research: <i>M. tuberculosis</i> , <i>M. smegmatis</i> , mRNA degradation, Machine learning
Title: Biology and Biotechnology Graduate Research Assistant Research: <i>M. smegmatis</i> , RNA-binding proteins
Title: Post-Baccalaureate Research Assistant Research: <i>M. smegmatis</i> , esxB, ribosomal proteins
Title: Undergraduate MQP Research Assistant Research: <i>M. smegmatis</i> , mRNA half-life, mRNA concentration

Table 3.1: Researchers Interviewed. A description of the researchers who were interviewed.

3.2 Understanding and Compiling the Important Goals and Tasks for the Prototype Metabolic Pathway Biovisualization

To understand the important tasks for this prototype metabolic pathway visualization, four interviews were conducted with researchers of the *M. tuberculosis* community (Table 1). Prior to conducting the interviews, oral consent was provided. These researchers (Table 1) were asked to think about what a visual image or description of a metabolic pathway is. Then, they were asked to utilize some existing metabolic pathway visualizations and note both effective and less effective aspects or functions. At the end, interviewees were asked to consider important features and tasks after exploring existing visualizations. The final questions used in all interviews can be found in Appendix A. The goal of these interviews were to understand what aspects are great in existing data visualizations, understand what elements are lacking in existing data visualizations, to identify what are the goals for a data visualization of *M. tuberculosis* metabolic pathways would be, to note tasks that a biovisualization should accomplish, and to define some task abstractions from the interviewee’s response. The responses were then recorded for use when building this prototype metabolic pathway visualization. Responses were then organized by order of occurrence to denote the importance of a task or feature.

3.3 Creating and Improving on the Prototype Metabolic Pathway Biovisualization

The bulk of this prototype visualization was built using JavaScript and CSS. The d3.js version 6 and node.js libraries were used to handle the structuring and display of data seen in all three constructed panels. Express.js and axios library were utilized in the file input mechanism. The entire visualization is hosted on a React.js server. As the final prototype was

3.3. CREATING AND IMPROVING ON THE PROTOTYPE METABOLIC PATHWAY BIOVISUALIZATION

being constructed, conversations regarding the type of information and exact detail choices were had with both advisors. The completed prototype can be found at the following GitHub repository:<https://github.com/ao-joker/V2-BCB-MQP-BioVis>.

4.1 Building an Understanding Regarding the Important Tasks and Goals for a Prototype Metabolic Pathway Biovisualization of *Mycobacterium tuberculosis*

Interviews with *M.tuberculosis* researchers (Table 1) were performed in order to better understand both the tasks and goals of a prototype metabolic pathway biovisualization of *Mycobacterium tuberculosis*. The insights of these interviews can be denoted into 3 distinct categories: key actions, important data features, displaying the data. The goals of the biovisualization focused on displaying a metabolic pathway without overwhelming the user and adding a mechanism to receive inputted data which would then be visualized alongside the pathway to be able to draw inferences about larger changes of gene expression. This information was important in structuring and building the prototype metabolic pathway biovisualization.

4.1.1 Key Actions

There were three classes of key actions that the interviewees found important: discovering data, searching options, and uploading data. Firstly, in regard to discovering data, many interviewees expressed desire to take the data presented from a visualization and make new inferences about gene expression changes. For example, one researcher noted the importance of having all kinds of data present in a given visualization for ease of access. Metabolic pathway data is inherently multivariate in nature[1, 4, 5, 8]. Using interactivity to display data per the user's

4.1. BUILDING AN UNDERSTANDING REGARDING THE IMPORTANT TASKS AND GOALS FOR A PROTOTYPE METABOLIC PATHWAY BIOVISUALIZATION OF *MYCOBACTERIUM TUBERCULOSIS*

interest, such as in the use of “on-click” mechanisms and tooltips, would be useful to show all the data in a meaningful way without being too overwhelming. Additionally, doing so will provide a sense of discovery of the data and information.

Secondly, the ability to search and find data of interest was highly important to the interviewees. When provided sample metabolic pathway biovisualizations, interviewees often called attention to the degree of ease in finding particular data of interest. A biovisualization of a metabolic pathway should present some ability to find data of interest, whether it be a more global view focused on pathways or a more specific view focused on individual proteins. Many noted the limited reach and ability of existing biovisualizations to both find identifiable pathways and proteins in those pathways. Databases do have this well-received feature, but are more text and information heavy by nature. A prototype metabolic pathway biovisualization should work, thus, to combine the ability to search per pathway and per protein. With a greater variability in searching, researchers will be able to use a metabolic pathway biovisualization most effectively. Current implementations of this task would include the ability to locate, browse, and specifically search for the data in question.

The final action noted was the ability to input data to note changes in gene expression. As aforementioned, many researchers, including the interviewees, are interested in understanding how changes in expression of particular proteins may be able to impact a biological system. These changes in the expression of genes is a large area of interest that fuels research in *M. tuberculosis*. No visualizations known or used by the interviewees allowed data to be implemented and compared to the present metabolic pathway network. All interviewees expressed interest in implementing the ability to insert some data to understand how it impacts metabolic pathways. The use of different visual channels would especially be vital in this role, as users would want to notably see those differences. Current implementation would likely involve the use of an upload mechanism that then impacts the visualization as a whole, providing data and editing the visualization in some way discernable and useful to the user. Organizing the provided data in a table would be an easy solution that would be both familiar to the user and easy to reference when identifying specific changes with the greater visualization as a whole.

4.1.2 Important Data Features

Metabolic pathways, by their complexity, are inherently multivariate in nature. All the interviewees and researchers as a whole are well aware of this notion. Thus, when compiling a dataset to represent the *M. tuberculosis* metabolic pathway, there is a variety of data that needs to be conveyed in order to form a meaningful and useful visualization. Interviewees noted that a metabolic pathway should include not only protein to protein connections in a

pathway, but the molecules that serve as substrates and products. Doing so would provide users a greater sense of familiarity in the pathway structure. Furthermore, molecules were noted by some interviewees to be important to the understanding impacts of gene expression changes. Additionally, the interviewees noted the importance of including gene name alongside the protein name. With different identifiers, including all possible varieties will make the visualization more apt in the key actions of discovery and search. As final note, ensuring the data within the visualization was accurate and supported with high confidence was noted as essential for any good visualization. One interviewee specifically added that potentially linking proteins and their interaction to academic papers would provide an easy way to validate the visualization in high confidence and find the original literature. Given the importance of accuracy and peer-review in the scientific community, this is a logical and important feature that should be inherent to all data. Current ideas and implementations would present the confidence of a particular result alongside a particular data point. Consequently, only including data with high confidence could also address the aspect of validity and accuracy of the data. Linking the original and other potential literature, as in the case of database annotations, would address the thoughts of the aforementioned interviewee. More specifically about the pathway itself, the inclusion of connections between molecules and other pathways could prove strenuous on many devices, as in the case with the BRENDA pathway[20]. To condense, current implementation ideas could focus on providing different layouts to allow the user to investigate different aspects and perspectives of the pathway.

4.1.3 Displaying the Data

As in all visualizations, displaying the data is itself an intricate design choice. Apart from important data, this area delves into how the user would see and interact with the data benign visualized. Interviewees all noted that metabolic pathway data is inherently extensive in nature, with there being countless values and items to represent at one time. In one specific example, an interviewee indicated that displaying all data at one would be overwhelming, impeding the ability to draw inferences and conclusions for a visualization. With current techniques, any tabular and otherwise structured data should remain uncondensed and easy to identify. Furthermore, using tooltips and other features focused on interactivity would make the visualization more appealing and prevent the user from being overwhelmed with the vast amount of information. Additionally, implementing a layout system to show different kinds of data at once is a viable option and would allow users to adjust their perspective on their own discretion. On a final note, using the visual channels of size and color will be essential to discerning all the different kinds of data. Most interviewees agreed on this statement. Implementing color and size channels will aid the user in quickly and readily identifying differences in data values and features of the visualization.

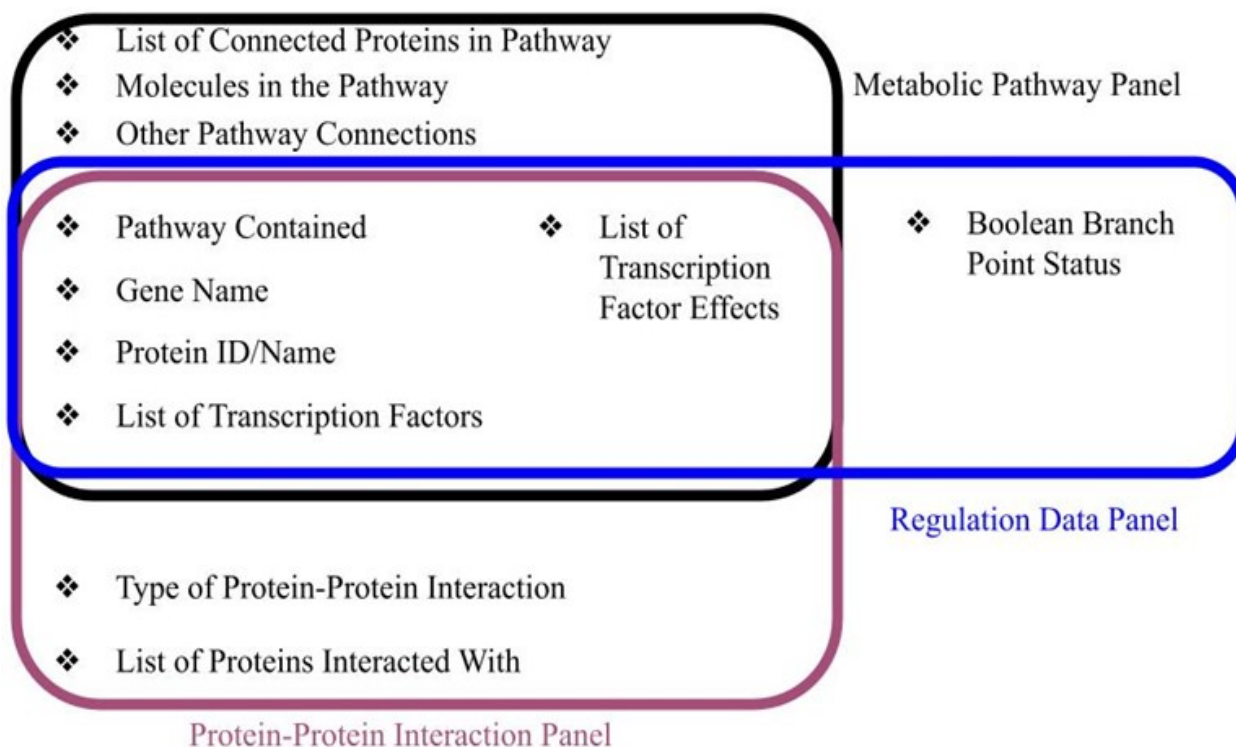


Figure 4.1: A representation of the data involved in each panel of the biovisualization. The black bordered data are used in the metabolic pathway panel. The pink-red bordered data are used in the protein-protein interaction panel. The blue bordered data are used in the regulation data panel.

4.2 Compiling a Dataset for the Prototype Metabolic Pathway Biovisualization of *Mycobacterium tuberculosis*

In order to create a visualization of *M. tuberculosis* metabolic pathway, the data for this pathway needed to be sought and organized. For the purposes of this prototype, data was compiled for the Glycolysis/Gluconeogenesis and Citrate Cycle pathways and a collection of proteins involved with T-receptor signaling grouped together for similar function. All are specific to *M.tuberculosis*. The final dataset consists of the following data categories: protein name, protein ID, pathway, proteins connected to in the pathway, molecules connected to by substrate and product relation, pathways connected to, a list of transcription factors that impact expression of a protein (if applicable), the impact of the transcription factor (if applicable), boolean status as a branch point, high confidence proteins that interact with a given protein in some way, and the nature of the interaction of those high confidence proteins. The importance of this data to each panel of the biovisualization can be found in Figure 4.1.


Protein name provides the gene name whereas protein ID denotes the name of the protein and enzyme commission (EC) number. Proteins connected to in the pathway refers to the organization of the proteins in which one protein takes in a substrate and outputs a product that a connected protein uses as a substrate in its reaction. The layout of the metabolic pathway would only show these proteins and not include molecules. Molecules connected to by substrate and product relation refers to creating a layout where both the molecules and proteins accounting for the change of substrates and molecules are included along with the proteins. The category “Pathways connected to” refers to proteins that lead into other proteins not included in the current pathway of interest. The transcription factor list and the impact of the transcription factor, denoted by a positive or negative one, indicates if other proteins impact the expression of a particular protein in a positive or negative way. Branch point status refers to proteins that can make two different products from a given substrate. Given limited information on the branch point status, this value was characterized as a boolean so as to present this information to the user. Finally, protein interactions were included in a list with a separate, corresponding list concerning the nature of those protein interactions in order to construct a protein-protein interaction network. This data indicated the type of interaction as specific as possible given the information available. Only proteins with high confidence were included. All the data here was used in the construction of the final visualized metabolic pathway, displaying a protein-protein interaction map, and indicating changes in gene expression by file upload. In order to compile this dataset, a number of databases were used to compile all the necessary data. No one specific database had all the necessary data needed for this visualization. Thus, the dataset was entirely compiled by hand, scoping through a number of databases (Found in experimental setup). Given this method of construction, the addition of papers to link users to the original literature was not included.

4.3 The Prototype Metabolic Pathway Biovisualization of *Mycobacterium tuberculosis*

Guided by the tasks and goals analysis and using the data collected, a prototype metabolic pathway visualization of *M. tuberculosis* was constructed. The final, prototype visualization consisted of three, linked panels that comprise the one larger visualization. These three panels are aptly named the metabolic pathway panel, the protein-protein interaction panel, and the regulation data panel. The final prototype visualization, its code, and the datasets used can be found at this GitHub repository: <https://github.com/ao-joker/V2-BCB-MQP-BioVis>. The panel colors are not a design choice, but rather a designation choice.

4.3.1 The Metabolic Pathway Panel

As implied for the title of this panel, this portion of the visualization focuses on displaying the actual pathway. For this visualization, only one pathway is displayed at a time. A drop down menu allows the user to alternate and change the pathway being shown. The pathway here is constructed using a force-directed network. With a particular pathway, radio buttons allow the user to alternate the pathway layout. One layout shows only connecting proteins, where each node represents a protein and the edges connect proteins in the metabolic pathway. Most reactions are in equilibrium and this visualization implies equilibrium between reactions. While this is not true for many cases, all are left undirected. The second layout includes the molecules, substrates and products, and proteins in the pathway. Additional nodes, smaller and gold in color, represent the molecules in the pathway. Edges represent the progression of the pathway as with the first layout. Regardless of the layout, each node always includes the name of the molecule or the gene name of the protein in question. A third radio button allows the user to toggle transcription factors as part of the pathway. These transcription factors are represented as smaller, red nodes with the gene name of the transcription factor positioned next to the node similarly to the proteins and molecules. This function mainly lets the user see which transcription factors impact which proteins in the pathway to draw inferences from changes in gene expression if applicable. Edges for transcription factors are acting on the connected node and are designated as undirected. A fourth and final radio button allows the user to toggle other pathway connections. If a protein is connected to another protein found in a different pathway, then a node to present that pathway is drawn. Other pathways are represented as larger, orange nodes with the name of the pathway positioned next to the node similar to the other nodes described. This allows the user to draw further inference and denote where impacts of gene expression changes one pathway may cascade to other pathways. Hovering over a particular protein reveals a tooltip with further information about the protein. This includes the gene name, protein name, and EC identification number. Clicking on a protein node will draw the protein-protein interaction network for that protein in the protein-protein interaction panel. The function of clicking on a node is subsequently reserved for creating protein-protein interaction networks. Hovering over a molecule node will just indicate the molecule's name. Hovering over a transcription factor node will display the transcription factor name and the effect it has on its connected protein. Hovering over a pathway node will display the pathway itself and indicate users to explore that pathway, if applicable, in the biovisualization. Clicking on a molecule node, transcription factor node, or pathway node will produce an alert that informs the user that no protein-protein interaction network exists for this particular node. Figure 4.2 shows a sample image the metabolic pathway panel described.



4.3. THE PROTOTYPE METABOLIC PATHWAY BIOVISUALIZATION OF *MYCOBACTERIUM TUBERCULOSIS*

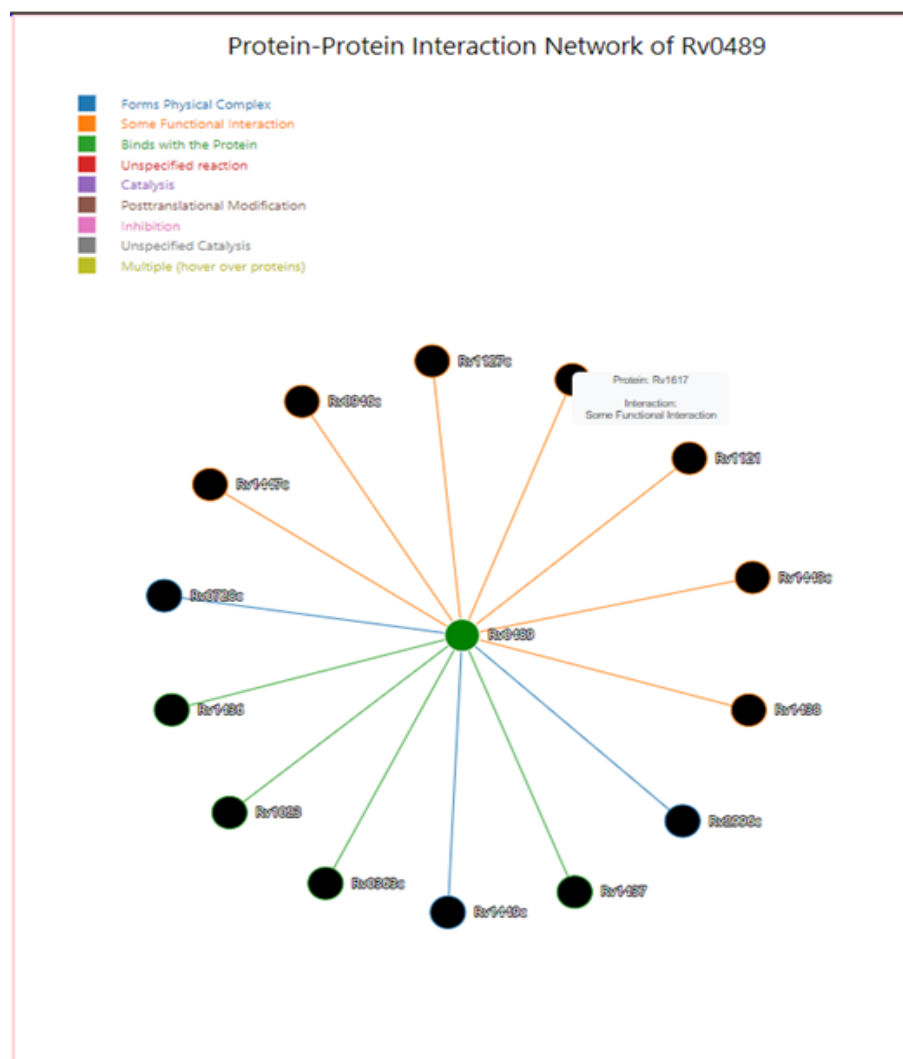


Figure 4.3: A sample image of the protein-protein interaction panel. This sample image displays the protein-protein interaction panel within the prototype biovisualization. The protein-protein interaction network of the protein product of Rv0489. A legend to decipher legend color with the specific type of interaction is found within the panel. The Rv1617 node is hovered over to display the tooltip.

4.3.3 The Regulation Data Panel

This panel displays data that the user inputs for noting changes in gene expression. Attached to this panel, the user is able to upload a csv file from the client side to the server. The csv file in question specifies two parameters: protein name and fold change, as would be obtained from data such as RNA sequencing. When uploaded, the panel will display a maximum of two data tables. The first table contains the protein name, pathway or pathways the protein is located in, whether the resulting regulation is positive or negative, and the exact degree of fold change taken from the data inputted by the user. The second table contains the transcription factor name, pathway or pathways the protein is located in, whether the resulting regulation is positive or negative, and the exact degree of fold change taken from the data inputted by the user. Additionally, using this uploaded data, the metabolic pathway panel would be modified. Any protein or transcription factor listed in the selected pathway would result in that specific node being outlined in green for upregulation or red for downregulation. If the fold change was zero or not included, no change to the protein or transcription factor node would be seen. Currently, if a different pathway should be examined, a user would have to re-upload the dataset for regulation again. Another limitation of the design is that only data part of the final data set compiled is included in the tables mentioned and the changes in gene expression. Figure 4.4, Figure 4.5, and Figure 4.6 displays a set of sample images that display an active regulation data panel, the mechanism for file upload, and the changes a file upload has on the metabolic pathway panel respectively.

4.3. THE PROTOTYPE METABOLIC PATHWAY BIOVISUALIZATION OF *MYCOBACTERIUM TUBERCULOSIS*

Regulation Data from testReg3.csv

<u>Protein</u>	<u>Pathway</u>	<u>Branch Point</u>	<u>Regulation Type</u>	<u>Fold Change</u>	<u>P value</u>
Rv1436	Glycolysis/Gluconeogenesis	FALSE	Downregulation	-1.5	0.03
Rv2702	Glycolysis/Gluconeogenesis	FALSE	Downregulation	-0.4	0.4
Rv1023	Glycolysis/Gluconeogenesis	FALSE	Upregulation	2	0.1

<u>Transcription Factor</u>	<u>Pathway</u>	<u>Branch Point</u>	<u>Regulation Type</u>	<u>Fold Change</u>	<u>P value</u>
Rv0767c	N/A	N/A	Upregulation	1.2	0.2

Figure 4.4: A sample image of the regulation data panel. The sample image shows the panel after a sample dataset, testReg3, is uploaded to the prototype biovisualization. Two tables, one from proteins and one for transcription factors, are displayed that contain data from the uploaded file and from dataset the prototype biovisualization uses.



Figure 4.5: A sample image of the file upload mechanism tied to the regulation data panel. A file is selected and is then uploaded using the respective button.




Figure 4.6: A sample image of the impacts of the file upload on the metabolic pathway panel. Displayed is a portion of the Glycolysis / Gluconeogenesis pathway in the metabolic pathway panel after uploading a sample gene expression dataset. The panel is much larger in the prototype biovisualization. The uploaded dataset is testReg3. Upregulated genes proteins are outlined in green whereas downregulated proteins are outlined in red. Exact references regarding the fold change and significance are done by referring to the tables in the regulation data panel.

DISCUSSION

Presented in this paper is a prototype biovisualization of *M. tuberculosis* metabolic pathways. Using a variety of databases, a dataset was compiled that incorporated data specific to the aspects of this visualization. A set of interviews with researchers in the *M. tuberculosis* community were conducted to build an understanding of the goals and tasks that would be important for building a metabolic pathways biovisualization. With this data and direction, a prototype biovisualization of *M. tuberculosis* metabolic pathways was constructed. This visualization consists of three distinct panels: one that showcases one specific pathway, one that indicates protein-protein interactions, and one that allows users to input gene expression data and use it to draw inferences from the first panel. The final prototype is, as the name implies, a representation of what a final product would look like. There are some noteworthy limitations with the dataset and the interviews that should be noted. Likewise, there are a number of aspects of the biovisualization that should be addressed or added in order to improve the user experience with the biovisualization as a whole. Despite these limitations, this prototype presents a tool for visualizing gene expression datasets alongside protein-protein interactions and metabolic pathways specific to *M. tuberculosis* from which researchers can use and draw inferences about the effect of gene expression changes on the overall system.

5.1 Data Limitations and Future Applications

As is true for any set of compiled or collected data, there are limitations to the one created in this paper. In an effort to create the dataset for a few pathways, the data was taken from many

different databases that were of high confidence and accompanied with relevant citations. The original literature was not checked to ensure the validity of the data if there was any discrepancy. Furthermore, the amount of available data from these databases was limited. There were a few instances where access to a database and its information was locked unless a subscription was purchased. While going to the original literature is another option, this process is time intensive and can also be prevented without a subscription as well. Using a number of databases and collaborating with researchers who have access to these locked databases, especially if the data is important to the visualization, could prove useful to obtaining the necessary data. One notable case can be found in the protein-protein interaction data and panel. This panel displays interactions of a protein of interest with high confidence. In the construction of this dataset, only data of high confidence from primarily STRING and STITCH was selected. This limits the ability of the users to explore deeper protein neighborhoods as with these two databases and as highlighted in Dang, et al. (2017). Selecting high confidence data in the construction of the dataset limits this ability in favor of greater empirical evidence.

In terms of methodology, the data collected was taken and investigated by hand. Any and all data with high confidence was recorded individually and exported as a csv formatted file for use in the visualization. Part of this effort was to understand what data should be included as an important step in creating the base dataset the visualization draws from. In order to expand the use of the visualization, however, more pathways will need to be added to this dataset. There may be a more efficient method that can be used to comb through one or multiple databases at a time and collect the necessary data. Doing so will save time in building the dataset that can be allocated to improving the visualization.

5.2 Interview Limitations and Future Applications

As noted, previous work such as in Murray et al, (2016) and Dang, et al. (2017), have identified many key goals and tasks that are important to researchers in general. To gain a greater insight into the specific needs and tasks for a biovisualization of *M. tuberculosis* metabolic pathways, interviews were conducted with researchers in the *M. tuberculosis* community. In order to get a better understanding of the goals and tasks that are specific to this group of researchers, more interviews should be performed. Adding questions specific to improving the current prototype biovisualization would be useful to understand areas not highlighted in the results or discussion that can be improved for both efficiency and user experience.

5.3 Biovisualization Limitations and Future Applications

Limitations of the biovisualization are related to both implementation and design choice. Starting with the former, the biovisualization is defined as a prototype, implying further development and improvement should be completed. One aspect of this development that should be of key interest is related to the uploading of data in the regulation data panel. In its current implementation, uploading a file will only indicate changes of gene expression related to the current pathway layout visualized in the metabolic pathway panel. This limits the ability of the user to actively investigate, discover, and draw inferences about their data on metabolic pathways as a whole. Furthermore, if the layout is changed and the user reuploads the same data set, no gene expression changes will be shown as the information has already been uploaded. Thus, finding a way to remove datasets from the program and to hold onto their information between layouts is an important next step in the development of this biovisualization.

When uploading data to the regulation panel, the user is expected to reference the regulation data panel to understand the degree of change of a protein or transcription factor's expression. While a small set of data is manageable, larger datasets will make this task less efficient and manageable if the user would like to easily denote changes in expression between two or more proteins. Implementing a colored scale would be easier for the user and make denoting differences in expression more efficient. Additionally, the current set-up of the regulation data panel involves incorporating the uploaded data in a set of tables for the user to easily reference. As more data is included in uploaded files, the panel will become filled with long lists of tables. To better facilitate this process, making the regulation data panel svg scrollable would easily facilitate the additional data without losing the ability to see the information.

In regard to design choice, there are three key aspects that limit the presentation of understanding of the data: accessibility to the colorblind, the use of force-directed networks, and directionality of the edges in those networks. Firstly, the biovisualization does not take into account colorblindness. While panel backgrounds will change, as they simply exist for designation, node coloration and the legend in the protein-protein interaction panel will be important for denoting different items. Identifying and using colors that are color-blind friendly will be an important step for improving the biovisualization for a wider audience. This idea should be at the forefront with other recommendations aforementioned. Secondly, force-directed graphs are limiting in their ability to convey meaningful structure. Many pathways are drawn to convey more of an ordered progression that makes recognizing and identifying the pathway much easier. Glycolysis and Gluconeogenesis, for example, are often depicted as a linear chain of reactions. Force-directed networks place nodes at a certain distance from each other by some force metric and length distance that is irrespective of commonly drawn organizations used by researchers.

Using this network structure was found to be the best method for implementing many different protein data points together at the cost of this lost structure. While the different layouts highlight the organized flow, looking at other network structures and algorithms may improve the user experience and help them draw better inferences from the visualization as a whole. Thirdly, the current biovisualization has edges which are undirected. Future improvements should add direction, especially for reactions where equilibrium cannot be assumed and is not the case. Likewise, adding direction for transcription factor action would make visualizing the action simpler.

CONCLUSIONS

Metabolic networks are complex systems that are always in flux because of changes in gene expression within a living system. To better understand these living systems, visualizing these changes will allow researchers to better understand how these changes in these pathways impact growth, survival, and persistence. In this paper, a prototype metabolic pathway visualization for *M. tuberculosis* was constructed based upon the goals and tasks from published work and interviews with researchers in the *M. tuberculosis* community. While nothing inherently new in design and task implementation, such as a new network design, was constructed with this prototype, this visualization combines metabolic pathway data and protein-protein interaction with the ability to upload gene expression datasets to allow researchers to visualize and draw inferences from their data. Furthermore, this prototype creates a visualization tool with a specific focus on *M. tuberculosis* for use in that research community. It is the goal of this paper to showcase these ideas in order to be able to implement them in a more integrated and robust visualization.

APPENDIX A: INTERVIEW QUESTIONS

1. Can you please tell me about your research and work?
2. Have you ever used any databases or visualizations specific to biology generally or *M. tuberculosis* specifically? What have you used them for? (Continues if applicable)
 - a) With regard to this visualization, what aspects of the visualization is useful for understanding *M. tuberculosis* metabolic pathways?
 - b) Is there anything that is extraneous, ineffective or missing, regarding the visual or the data present?
3. Using your knowledge of the metabolic pathway (and the visualization you mentioned if applicable), what are the key elements when visualizing metabolic or protein data?
 - a) Could you list some particular actions that are associated with those key components?
4. *Shows the KEGG database OR from answer to question 2 and provides time for exploration*
 - a) What are your thoughts when first exploring this visualization? *If needed* Examples would include the design, the data, the interface, etc.
 - b) With regard to this visualization, what aspects of the visualization is useful for understanding the *M. tuberculosis* metabolic pathway?
 - c) Is there anything that is extraneous, ineffective or missing, regarding the visual or the data present?
5. *Shows the STRING PPI visualization and provides time for exploration* (Using a goal-like find confidence and most strong evidence for a functional and physical connection).
 - a) What are your thoughts when first exploring this visualization? *If needed* Examples would include the design, the data, the interface, etc.
 - b) With regard to this visualization, what aspects of the visualization is useful for understanding the *M. tuberculosis* metabolic pathway?
 - c) Is there anything that is extraneous, ineffective or missing, regarding the visual or the data present?
6. In visualizing the *M. tuberculosis* metabolic pathway and given from the examples shown, what are some aspects that you particularly liked or disliked across all the visualizations?
7. After observing some sample visualizations, what would you determine are the most important parts of a visualization? This could be data to include, features, details, etc.

APPENDIX A: INTERVIEW QUESTIONS

8. It is my goal to create a biovisualization that incorporates the *M. tuberculosis* metabolic pathway, a PPI network, and a method to input RNA-seq data to identify changes in protein levels and regulation.
 - a) Given what you have seen and identified, are there other key components that should be noted when comparing and visualizing these three aspects together?
 - b) Are there areas that you believe should receive more or less detail?
 - c) How would a visualization like this be useful in your research and work, if at all?
9. Are there any other details you would like to add or questions you would like to ask?

APPENDIX B: INFORMED ORAL CONSENT FORM

Thank you again for responding and accepting my request for an interview with you. In order to proceed with the interview, I ask that you kindly provide oral consent to this interview.

You, the participant, are being invited to participate in an interview with Adrian Orszulak. This interview is part of a Major Qualifying Project (MQP) to understand the goals and tasks for creating a biovisualization of the *M. tuberculosis* metabolic pathways. In this interview, I am seeking to define what you, the researcher, would define as important goals for a biovisualization of the *M. tuberculosis* metabolic pathways. Using this information, I will be able to list a set of goals for building the said visualization, and identify the corresponding tasks needed to accomplish those goals. The interview will last from 1 hour to 1.5 hours. Your name will not be published, but your specific field of work (such as metabolomics) will be.

Do you have any questions?

This process is a voluntary one. If you do not want to answer any particular question, you are not obligated to do so. We will simply move on should you make this known. Additionally, if you no longer feel open to an interview or wish to stop the interview, you have the right to withdraw from the interview at any point.

Do I have your consent to move forward with this interview?

In addition, I would like to ask if I could have your consent to record the interview. You do not need to agree to have the interview recorded. I will act as both interviewer and notetaker if you do not consent to the recording. Do I have your consent to record the interview?

Do you have any questions before we begin?

BIBLIOGRAPHY

- [1] B. Worley and R. Powers, “Multivariate analysis in metabolomics,” *Curr. Metabolomics*, vol. 1, no. 1, pp. 92–107, 2013.
- [2] D. B. Kell, M. Brown, H. M. Davey, W. B. Dunn, I. Spasic, and S. G. Oliver, “Metabolic footprinting and systems biology: the medium is the message,” *Nat. Rev. Microbiol.*, vol. 3, no. 7, pp. 557–565, Jul. 2005.
- [3] World Health Organization, “Tuberculosis (TB),” <https://www.who.int/publications/i/item/9789240037021>, 2021.
- [4] N. Levin, “Multivariate statistics and the enactment of metabolic complexity,” *Soc. Stud. Sci.*, vol. 44, no. 4, pp. 555–578, Aug. 2014.
- [5] S. Ren, A. A. Hinzman, E. L. Kang, R. D. Szczesniak, and L. J. Lu, “Computational and statistical analysis of metabolomics data,” *Metabolomics*, vol. 11, no. 6, pp. 1492–1513, Dec. 2015.
- [6] D. F. Warner, “Mycobacterium tuberculosis metabolism,” *Cold Spring Harb. Perspect. Med.*, vol. 5, no. 4, pp. a021 121–a021 121, Dec. 2014.
- [7] B. Lee, C. Plaisant, C. S. Parr, J.-D. Fekete, and N. Henry, “Task taxonomy for graph visualization,” in *Proceedings of the 2006 AVI workshop on BEyond time and errors novel evaluation methods for information visualization - BELIV '06*. New York, New York, USA: ACM Press, 2006.
- [8] P. Murray, F. McGee, and A. G. Forbes, “A taxonomy of visualization tasks for the analysis of biological pathway data,” *BMC Bioinformatics*, vol. 18, no. Suppl 2, p. 21, Feb. 2017.
- [9] T. Dang, P. Murray, and A. Forbes, “BioLinker: Bottom-up exploration of protein interaction networks,” in *2017 IEEE Pacific Visualization Symposium (PacificVis)*. IEEE, Apr. 2017.
- [10] Y. Wang, Q. Wang, H. Huang, W. Huang, Y. Chen, P. B. McGarvey, C. H. Wu, C. N. Arighi, and UniProt Consortium, “A crowdsourcing open platform for literature curation in UniProt,” *PLoS Biol.*, vol. 19, no. 12, p. e3001464, Dec. 2021.

- [11] UniProt Consortium, “UniProt: the universal protein knowledgebase in 2021,” *Nucleic Acids Res.*, vol. 49, no. D1, pp. D480–D489, Jan. 2021.
- [12] S. Turkarslan, E. J. R. Peterson, T. R. Rustad, K. J. Minch, D. J. Reiss, R. Morrison, S. Ma, N. D. Price, D. R. Sherman, and N. S. Baliga, “A comprehensive map of genome-wide gene regulation in mycobacterium tuberculosis,” *Sci. Data*, vol. 2, no. 1, Dec. 2015.
- [13] J. J. Davis, A. R. Wattam, R. K. Aziz, T. Brettin, R. Butler, R. M. Butler, P. Chlenski, N. Conrad, A. Dickerman, E. M. Dietrich, J. L. Gabbard, S. Gerdes, A. Guard, R. W. Kenyon, D. Machi, C. Mao, D. Murphy-Olson, M. Nguyen, E. K. Nordberg, G. J. Olsen, R. D. Olson, J. C. Overbeek, R. Overbeek, B. Parrello, G. D. Pusch, M. Shukla, C. Thomas, M. VanOeffelen, V. Vonstein, A. S. Warren, F. Xia, D. Xie, H. Yoo, and R. Stevens, “The PATRIC bioinformatics resource center: expanding data and analysis capabilities,” *Nucleic Acids Res.*, vol. 48, no. D1, pp. D606–D612, Jan. 2020.
- [14] P. D. Karp, R. Billington, R. Caspi, C. A. Fulcher, M. Latendresse, A. Kothari, I. M. Keseler, M. Krummenacker, P. E. Midford, Q. Ong, W. K. Ong, S. M. Paley, and P. Subhraveti, “The BioCyc collection of microbial genomes and metabolic pathways,” *Brief. Bioinform.*, vol. 20, no. 4, pp. 1085–1093, Jul. 2019.
- [15] S. Paley, R. Billington, J. Herson, M. Krummenacker, and P. D. Karp, “Pathway tools visualization of organism-scale metabolic networks,” *Metabolites*, vol. 11, no. 2, p. 64, Jan. 2021.
- [16] R. Caspi, R. Billington, I. M. Keseler, A. Kothari, M. Krummenacker, P. E. Midford, W. K. Ong, S. Paley, P. Subhraveti, and P. D. Karp, “Karp the MetaCyc database of metabolic pathways and enzymes - a 2019 update,” 2019.
- [17] M. Kanehisa and S. Goto, “KEGG: kyoto encyclopedia of genes and genomes,” *Nucleic Acids Res.*, vol. 28, no. 1, pp. 27–30, Jan. 2000.
- [18] M. Kanehisa, “Toward understanding the origin and evolution of cellular organisms,” *Protein Sci.*, vol. 28, no. 11, pp. 1947–1951, Nov. 2019.
- [19] M. Kanehisa, M. Furumichi, Y. Sato, M. Ishiguro-Watanabe, and M. Tanabe, “KEGG: integrating viruses and cellular organisms,” *Nucleic Acids Res.*, vol. 49, no. D1, pp. D545–D551, Jan. 2021.
- [20] A. Chang, L. Jeske, S. Ulbrich, J. Hofmann, J. Koblitz, I. Schomburg, M. Neumann-Schaal, D. Jahn, and D. Schomburg, “BRENDA, the ELIXIR core data resource in 2021: new developments and updates,” *Nucleic Acids Res.*, vol. 49, no. D1, pp. D498–D508, Jan. 2021.

- [21] D. Szklarczyk, A. Santos, C. von Mering, L. J. Jensen, P. Bork, and M. Kuhn, “STITCH 5: augmenting protein-chemical interaction networks with tissue and affinity data,” *Nucleic Acids Res.*, vol. 44, no. D1, pp. D380–4, Jan. 2016.
- [22] D. Szklarczyk, A. L. Gable, K. C. Nastou, D. Lyon, R. Kirsch, S. Pyysalo, N. T. Doncheva, M. Legeay, T. Fang, P. Bork, L. J. Jensen, and C. Mering, “January) STRING v11: protein-protein association networks with increased coverage, supporting functional discovery in genome-wide experimental datasets,” *Nucleic Acids Res.*, vol. 47, pp. D607–61 349, 2019.
- [23] —, “) the STRING database in 2021: customizable protein-protein networks, and functional characterization of user-uploaded gene/measurement sets,” *Nucleic Acids Res.*, vol. 49, no. D1, pp. D605–612, 2021.
- [24] Z. Hu, J. Mellor, J. Wu, and C. DeLisi, “VisANT: an online visualization and analysis tool for biological interaction data,” *BMC Bioinformatics*, vol. 5, p. 17, Feb. 2004.
- [25] Z. Hu, E. S. Snitkin, and C. DeLisi, “VisANT: an integrative framework for networks in systems biology,” *Brief. Bioinform.*, vol. 9, no. 4, pp. 317–325, Jul. 2008.
- [26] B. R. Granger, Y.-C. Chang, Y. Wang, C. DeLisi, D. Segrè, and Z. Hu, “Visualization of metabolic interaction networks in microbial communities using VisANT 5.0,” *PLoS Comput. Biol.*, vol. 12, no. 4, p. e1004875, Apr. 2016.
- [27] M. Baitaluk, M. Sedova, A. Ray, and A. Gupta, “BiologicalNetworks: visualization and analysis tool for systems biology,” *Nucleic Acids Res.*, vol. 34, no. Web Server issue, pp. W466–71, Jul. 2006.
- [28] P. Shannon, A. Markiel, O. Ozier, N. S. Baliga, J. T. Wang, D. Ramage, N. Amin, B. Schwikowski, and T. Ideker, “Cytoscape: a software environment for integrated models of biomolecular interaction networks,” *Genome Res.*, vol. 13, no. 11, pp. 2498–2504, Nov. 2003.
- [29] C. Baker, M. Carpendale, P. Prusinkiewicz, and M. Surette, “Genevis: visualization tools for genetic regulatory network dynamics,” in *IEEE Visualization, 2002. VIS 2002.*, 2002, pp. 243–250.
- [30] B. H. Junker, C. Klukas, and F. Schreiber, “VANTED: a system for advanced data analysis and visualization in the context of biological networks,” *BMC Bioinformatics*, vol. 7, no. 1, p. 109, Mar. 2006.
- [31] A. Kapopoulou, J. M. Lew, and S. T. Cole, “The MycoBrowser portal: a comprehensive and manually annotated resource for mycobacterial genomes,” *Tuberculosis (Edinb.)*, vol. 91, no. 1, pp. 8–13, Jan. 2011.