Published in final edited form as:

Proteomics. 2010 May ; 10(9): 1880–1885. doi:10.1002/pmic.200900723.

Pathway Palette: a rich internet application for peptide-, protein- and network-oriented analysis of mass spectrometry data

Manor Askenazi^{1,2,3,4,†}, Shaojuan Li^{1,2,†}, Saurav Singh^{1,2}, and Jarrod A. Marto^{1,2,3,**}

¹ Departments of Cancer Biology, Harvard Medical School, Boston, MA 02115

² Blais Proteomics Center, Dana-Farber Cancer Institute, Harvard Medical School, Boston, MA 02115

³ Department of Biological Chemistry and Molecular Pharmacology, Harvard Medical School, Boston, MA 02115

⁴ Department of Biological Chemistry, The Hebrew University of Jerusalem, Israel

Abstract

Recent improvements in proteomics technologies have collectively yielded data sets that far exceed the capabilities of typical low-throughput interpretation strategies. Unfortunately, tools designed to leverage the “peptide-centric” content of mass spectrometry-based proteomics lag the current rate of data production. Here we describe Pathway Palette, a freely accessible internet application that enables researchers to easily transition from peptides to biological pathways, while simultaneously retaining the qualitative and quantitative aspects of the underlying mass spectrometry data.

Keywords

Rich Internet Application; Mass Spectrometry proteomics; network analysis; protein-protein interaction; quantitative proteomics

Mass spectrometry is now well established as the technique of choice for characterization of proteins derived from a plethora of model systems in biomedical research. Moreover, collective advances in sample preparation, enrichment, and fractionation, amongst other methods[1] now support proteomics-based experiments that are designed to monitor changes in protein expression and post-translational modification state as a function of biological perturbation. In fact, mass spectrometry-based proteomics is now well integrated with hypothesis-driven research within many laboratories. In addition to sheer quantity, the diversity of relevant information which needs to be integrated in order to glean biological insight from the data poses an enduring challenge to practitioners[2]. A typical example of this problem, which has garnered much attention in recent years, is mass spectrometry-based analysis of protein-protein interactions derived from both targeted and large scale affinity purification experiments. This approach provides an ideal means to characterize multi-

^{**}To whom correspondence should be addressed: Jarrod A. Marto, Department of Cancer Biology, Dana-Farber Cancer Institute, 44 Binney Street, Smith 1158A, Boston, MA, 02115-6084, USA, Phone: (617) 632-3150 (office), Fax: (617) 582-7737, jarrod_marto@dfci.harvard.edu.

[†]These authors contributed equally to this work.

Pathway Palette can be accessed at: <http://blaispathways.dfci.harvard.edu>

COMPETING INTERESTS STATEMENT

The authors declare that they have no competing financial interests.

component protein complexes formed within a physiologically relevant context, e.g., *in-vivo* or *in-vitro*; nevertheless, reduction of these data to biological pathways and networks remains tedious and challenging. Despite these difficulties, a “network diagram” representation of protein-protein interaction data has become so ubiquitous that it is a *de facto* standard display for results from such large scale studies[3][4][5]. As a result, proteomics researchers naturally gravitate towards tools that allow extension of protein lists to meaningful biological annotation. Unfortunately, software tools that currently offer a network perspective tend to focus at the gene level and typically accept either a gene list or one of the few accession lists commonly used in microarray and high throughput sequencing platforms. This approach is problematic across a wide range of mass spectrometry-based proteomics experiments. For example, various *ad hoc* and community-derived standards[6] for data reporting suggest that protein identifications based on a single peptide sequence require higher stringency as compared to identifications based on multiple peptides. Unfortunately, simple metrics such as the number of peptides identified per protein are lost in gene-centric tools. As a more complex example, consider the case in which protein identification and quantification are augmented with data for post-translational modification. In this case, quantification data may vary across peptides otherwise assigned to the same protein. Here again, pathway tools that rely solely on gene ID fail to capture the full information content of mass spectrometry-based proteomics data. As a result of these and other limitations, practitioners typically use proteomic-specific toolsets at the early phases of their data analyses and then switch to more general-purpose software in an effort to place their primary data and observations into a biologically meaningful context. While some tools[7] have tried to bridge the gap, their functionality is typically limited to a physical-chemical analysis of the inferred proteins, which, while being potentially useful, e.g. for optimization of the experimental protocols, is far less relevant for interpretation of the data in the context of biological pathways. With these considerations in mind, we set out to develop Pathway Palette (Figures 1 and 2), a freely accessible, rich internet application that provides a canvas on which researchers can iteratively explore biological networks, while retaining all peptide-centric annotation of the source proteomics data.

Design Philosophy

Pathway Palette is designed to (i) provide seamless integration of experimental results with existing, publically available data resources, (ii) enable iterative personalized knowledge refinement without a requirement for user registration or tracking of specific user history, and (iii) efficiently divide computational tasks between client- and server-side resources. Because Pathway Palette is built upon a foundation of database repositories that are keyed by a wide range of annotations, including gene, protein, disease, and pathway, there is already a fundamental benefit in simply making this information available through a convenient, unified interface. However, our broader objective is for researchers to use Pathway Palette as a navigational tool to iteratively build new knowledge, typically in the form of new gene lists, annotations, or novel connections between otherwise unrelated molecules, in biological pathways and networks.

An implicit requirement to establish an enthusiastic and significant user base for Pathway Palette is the efficient management of protein lists. Although seemingly a trivial task, we have made two design choices related to list management which we believe will encourage broad community participation and ultimately publication of analyses that leverage Pathway Palette functionality. First, use of Pathway Palette does not require login credentials nor is user history tracked within or between sessions. In addition, simple text files can be used as input. As a result, users may anonymously re-visit previous analyses in an iterative manner, with their history and strategy stored client-side within human-readable files. Second, Pathway Palette treats user-created lists as first class annotation sources. This means that

users can upload lists, created from primary data, directly into the very same analytic framework used to assess significance of biological annotations such as GO[8], KEGG[9], and OMIM[10]. This enables users to quickly detect over-representation of lists created in previous experiments, and thereby validate that new knowledge has been generated. We believe that Pathway Palette is the first freely and anonymously accessible proteomics tool that provides this form of cyclical knowledge refinement.

Software Architecture

Pathway Palette is delivered to the user in the form of a rich internet application deployed through a combination of a Flex/Flash frontend and a backend built on PHP augmented with Python scripts, as well as a dedicated graph layout server implemented as a java servlet (<http://www.yworks.com>) from an Apache/Tomcat/Zend server (<http://www.apache.org/>, <http://www.zend.com>) (Figure 3). Applet-based solutions which compute graph layouts on the client-side were not considered since protein-interaction graphs can be prohibitively large and are not necessarily amenable to layout on a modern netbook or legacy desktop. In addition, it is well recognized that Flash has a much higher installation base in the modern internet browser ecosystem (<http://riastats.com/>, <http://www.statowl.com/>) as compared to Sun's Java JDK/JRE (<http://www.java.com/>).

Data Input

As a tool intended to provide primary support for proteomics research, Pathway Palette enables direct upload of data elements that are unique to mass spectrometry. In principle, these input data could be based on community-defined standards such as mzIdentML (<http://www.psidev.info/index.php?q=node/319>), however, since the associated definitions have not yet stabilized (<http://www.psidev.info/index.php?q=node/408>), we opted in the first deployment of Pathway Palette to use human-readable text and spreadsheets for data input. Two primary modes of input are available to the user:

1. The most unstructured form of input is a list. The list can either be a peptide list, a protein list or a gene list. This mode also enables users to cut and paste directly from their usual data sources into the Pathway Palette browser input fields.
2. Users may also upload structured spreadsheets generated by multiplierz, our recently described, mzAPI-based framework for mass spectrometry data analysis[11] (<http://blais.dfci.harvard.edu/multiplierz>). These spreadsheets can also be used in mode one above, through copy and paste of peptide or protein data into the website.

Direct upload of the latter provides convenient and comprehensive support for quantification data, meaning that in subsequent network analyses (see below) users can readily inspect all peptide identification and quantification evidence for each protein node. For either input mode described above, Pathway Palette will accept peptide sequences, protein identifiers from RefSeq[12] (e.g. gi|62362414) or SGD[13] (e.g. YHR023W), and genes symbols based on the NCBI's Entrez Gene[14] (currently the system supports the Human, Mouse and Yeast proteomes). In cases where the input consists of peptide sequences, users must specify whether the system should generate the largest set of proteins consistent with the input list, or use a greedy algorithm to calculate a minimal set. These options constitute the two extreme interpretations of the provided peptide list: either every protein containing any of the given peptides is listed, or the system iteratively selects a protein which accounts for the largest number of peptides, reports that protein and then eliminates the constituent peptides from subsequent iterations.

Navigating the Pathway Palette Canvas

The Protein Palette web application enables functional analysis of mass spectrometry data within the context of a protein network. In this section we will review key features of the browser interface and underlying system which directly support this goal.

Protein-Protein Interaction as a Primary Perspective

Upon data upload, Pathway Palette generates a protein-protein interaction (PPI) network based on information curated by BioGRID[15]. Furthermore we reduced the 25 interaction types in the BioGRID ontology into 4 major subsets: (i) Low-Throughput techniques (green), (ii) HTP/Complex detection, i.e., pull-downs (blue), (iii) HTP/Pairwise, i.e., yeast-two-hybrid (red) and (iv) Genetic interactions (not selected in default settings). Other datasources that are made available include: (v) NCBI Entrez database (PPI entries from: HPRD[16], BioGRID and BIND[17]), (vi) interologs[18] (putative interactions through homology), as well as (vii) the top 10% of the STRING[19] predicted PPI database. The system also supports input of user-defined PPI data as a simple tab-delimited list of Entrez Gene ID pairs. Pathway Palette enables the user to filter, prioritize, and color-code edges based on the underlying data source (rather than displaying all data sources for a given edge, Pathway Palette colors edges based on the highest priority data source containing the edge in question). In the case of peptide sequence input, the resulting nodes are colored (i) green if they represent a gene uniquely identified by peptide evidence, (ii) red if they can be removed from the plot without loss of peptide evidence, or (iii) blue if they are not uniquely supported by the peptide evidence.

The user can extend the network further by interactively selecting nodes and incorporating known interactors not present in the original dataset. Pathway Palette provides graph-theoretic operations which help manage exploration of the network, including: (i) extension of the network to include new nodes based on various PPI data sources, (ii) removal of leaf nodes, (iii) retrieval of missing edges, (iv) shortest path calculations, (v) Steiner Tree generation[20], among others. The resulting network is user-editable, from the manual addition and removal of nodes, through a graph layout algorithm, to manipulation of individual node shape, color, outline etc. The resulting graphs can be stored as GraphML files[21] on the user's desktop (for later upload) or printed in publication quality (vector graphics compatible) output. We recognize that much of the functionality of Pathway Palette is available in Cytoscape[22], a platform renowned for its advanced features and plugin architecture. While locally hosted software environments are suitable for many tasks and research environments, we chose to deploy Pathway Palette from a central server to alleviate administrative overhead for casual or other users who lack the necessary skills or infrastructure to support local instances of Cytoscape or similar platforms.

While the PPI network view constitutes the central data view in Pathway Palette, two additional representations are available which rely directly on mass spectrometry as the primary data source:

1. **Peptide view:** allows the user to inspect the exact peptides upon which a given protein (and hence gene) identification is based. The protein sequence is shown with the areas covered by peptide reads highlighted and underlined. By clicking on a highlighted region, all relevant peptides covering the region are shown ("coverage drill-down").
2. **Quant-Network:** At any point in the analysis, by navigating to the quantitation tab, the user can color proteins in the PPI network by fold change based on the uploaded spreadsheet data. Ratio values can be used to select proteins and can

therefore serve as a basis for user-defined protein-lists and custom annotations (see Figure 2).

In addition to network visualization and exploration, Pathway Palette supports functional annotation of protein networks. Enrichment analysis (represented graphically as a bar plot of negative log p-values) can be generated for KEGG, OMIM, and GO annotations (slim version[23]). In addition to these standard annotation sources, users can upload previously curated protein lists which can effectively be used as a private form of annotation, accessible through the same tools as the standard annotation sources. Finally, Pathway Palette includes additional functions for comparison of protein lists as well as other miscellaneous operations that are beyond the scope of this brief report. Detailed descriptions of the functionality provided by the system as well as detailed tutorials are all provided on the accompanying web-page: <http://blaispathways.dfci.harvard.edu/tutorial>

Pathway Palette is a rich, web-based resource which provides researchers with a convenient means to place mass spectrometry-based proteomics data within the context of biological pathways and networks. Support for anonymous user access along with server-side deployment of computationally intensive operations will lower barriers for community participation and ensure that the scale of graph analysis and network visualization is not limited by client-side computational resources. We will continue to engineer improvements to the website, driven by our own research needs in addition to user suggestions for future functionality.

Acknowledgments

The authors thank Eric Smith for preparation of figures, and also members of the Marto Lab for valuable discussions and testing of Pathway Palette. This work was supported by the Dana-Farber Cancer Institute and the National Human Genome Research Institute (P50HG004233).

Abbreviations

PPI Protein-Protein Interaction

References

1. Aebersold R, Mann M. Mass spectrometry-based proteomics. *Nature*. 2003; 422:198–207. [PubMed: 12634793]
2. Patterson SD. Data analysis-the Achilles heel of proteomics. *Nature Biotechnology*. 2003; 21:221.
3. Ideker T, Thorsson V, Ranish JA, Christmas R, et al. Integrated Genomic and Proteomic Analyses of a Systematically Perturbed Metabolic Network. *Science*. 2001; 292:929–934. [PubMed: 11340206]
4. Gavin A, Bosche M, Krause R, Grandi P, et al. Functional organization of the yeast proteome by systematic analysis of protein complexes. *Nature*. 2002; 415:141–147. [PubMed: 11805826]
5. Yan W, Lee H, Yi E, Reiss D, et al. System-based proteomic analysis of the interferon response in human liver cells. *Genome Biology*. 2004; 5:R54. [PubMed: 15287976]
6. Carr S, Aebersold R, Baldwin M, Burlingame A, et al. The Need for Guidelines in Publication of Peptide and Protein Identification Data: Working Group On Publication Guidelines For Peptide And Protein Identification Data. *Mol Cell Proteomics*. 2004; 3:531–533. [PubMed: 15075378]
7. Park D, Kim B, Cho S, Park S, et al. MassNet: a functional annotation service for protein mass spectrometry data. *Nucleic Acids Research*. 2008; 36:W491–495. [PubMed: 18448467]
8. Ashburner M, Ball CA, Blake JA, Botstein D, et al. Gene ontology: tool for the unification of biology. The Gene Ontology Consortium. *Nature Genetics*. 2000; 25:25–29. [PubMed: 10802651]
9. Kanehisa M, Goto S. KEGG: kyoto encyclopedia of genes and genomes. *Nucleic Acids Research*. 2000; 28:27–30. [PubMed: 10592173]

10. Hamosh A, Scott AF, Amberger J, Bocchini C, et al. Online Mendelian Inheritance in Man (OMIM), a knowledgebase of human genes and genetic disorders. *Nucleic Acids Research*. 2002; 30:52–55. [PubMed: 11752252]
11. Parikh J, Askenazi M, Ficarro S, Cashorali T, et al. multiplierz: an extensible API based desktop environment for proteomics data analysis. *BMC Bioinformatics*. 2009; 10:364. [PubMed: 19874609]
12. Pruitt KD, Tatusova T, Maglott DR. NCBI reference sequences (RefSeq): a curated non-redundant sequence database of genomes, transcripts and proteins. *Nucl Acids Res*. 2007; 35:D61–65. [PubMed: 17130148]
13. Cherry JM, Adler C, Ball C, Chervitz SA, et al. SGD: *Saccharomyces* Genome Database. *Nucleic Acids Research*. 1998; 26:73–79. [PubMed: 9399804]
14. Maglott D, Ostell J, Pruitt KD, Tatusova T. Entrez Gene: gene-centered information at NCBI. *Nucl Acids Res*. 2007; 35:D26–31. [PubMed: 17148475]
15. Stark C, Breitkreutz B, Reguly T, Boucher L, et al. BioGRID: a general repository for interaction datasets. *Nucleic Acids Research*. 2006; 34:D535–539. [PubMed: 16381927]
16. Mishra GR, Suresh M, Kumaran K, Kannabiran N, et al. Human protein reference database--2006 update. *Nucleic Acids Research*. 2006; 34:D411–414. [PubMed: 16381900]
17. Bader GD, Betel D, Hogue CWV. BIND: the Biomolecular Interaction Network Database. *Nucleic Acids Research*. 2003; 31:248–250. [PubMed: 12519993]
18. Yu H, Luscombe NM, Lu HX, Zhu X, et al. Annotation transfer between genomes: protein-protein interologs and protein-DNA regulogs. *Genome Research*. 2004; 14:1107–1118. [PubMed: 15173116]
19. von Mering C, Huynen M, Jaeggi D, Schmidt S, et al. STRING: a database of predicted functional associations between proteins. *Nucleic Acids Research*. 2003; 31:258–261. [PubMed: 12519996]
20. Scott MS, Perkins T, Bunnell S, Pepin F, et al. Identifying Regulatory Subnetworks for a Set of Genes. *Mol Cell Proteomics*. 2005; 4:683–692. [PubMed: 15722371]
21. Brandes U, Eiglsperger M, Herman I, Himsolt M, Marshall M. Graph Drawing. 2002:109–112.
22. Cline MS, Smoot M, Cerami E, Kuchinsky A, et al. Integration of biological networks and gene expression data using Cytoscape. *Nature Protocols*. 2007; 2:2366–2382.
23. Barrell D, Dimmer E, Huntley RP, Binns D, et al. The GOA database in 2009--an integrated Gene Ontology Annotation resource. *Nucleic Acids Research*. 2009; 37:D396–403. [PubMed: 18957448]

Figure 1.

Upon data upload, Pathway Palette creates a protein interaction network as the primary view (center panel). Users can click on specific nodes to obtain information rich “gene cards” (bottom left), and explore detected peptides and protein sequence coverage (bottom right). In addition, researchers can modify the network topology through, for example, addition of edges based on known protein-protein interactions (top left) and elimination of leaf nodes from the expanded graph (top center). At any point, users can query the current network view for enriched annotations, including generation of a significance bar plot (top right).

**NODE SELECTION BASED ON
QUANTITATIVE DATA
(e.g. iTRAQ RATIOS)**

**CUSTOM ANNOTATION BASED
ON USER-GENERATED GENE LISTS**

**USER-GENERATED GENE LISTS
BASED ON CURRENT SELECTION**

Figure 2.

Data for relative quantification, in this case based on iTRAQ, may be uploaded from multiplier^z-formatted spreadsheets and visualized as a color gradient. Users can then select nodes based on discrete ratio boundaries expressed as $\log_2$ ratios (top left). Selected nodes, based on ratio, topological features, etc., can be captured as a gene list (bottom left). Individual or sets of user-defined gene lists can then be downloaded as a human-readable text files. These gene lists can in turn be used as custom annotations (top right), which can then be analyzed for enrichment, in a manner analogous to that used with the standard annotation sources (e.g. GO, KEGG, and OMIM).

Figure 3.

Pathway Palette features a robust web architecture that is driven server-side (top) by an Oracle database and powerful graph agent for execution of computationally intensive operations that are delivered to the client (bottom) through a rich Adobe Flex front-end. Users can upload data in a variety of formats (bottom left), including peptide sequences, protein accession lists, Entrez Gene Symbols, and quantitative proteomics data from multiplierz-formatted spreadsheets. Users can also upload custom gene lists defined in previous Pathway Palette sessions (bottom right), thus enabling an intuitive, and anonymous, knowledge refinement cycle. Finally, Pathway Palette generates vector-compatible output to facilitate production of publication quality protein network graphics.