



Automating bibliometric analyses using Taverna scientific workflows: A tutorial on integrating Web Services



Arzu Tugce Guler^a, Cathelijm J.F. Waaijer^b, Yassene Mohammed^a, Magnus Palmblad^{a,*}

^a Center for Proteomics and Metabolomics, Leiden University Medical Center, Leiden, The Netherlands

^b Centre for Science and Technology Studies, Faculty of Social and Behavioural Sciences, Leiden University, Leiden, The Netherlands

ARTICLE INFO

Article history:

Received 14 January 2016

Received in revised form 11 May 2016

Accepted 13 May 2016

Available online 22 June 2016

Keywords:

Bibliometrics

Europe PMC

PubMed

Taverna

Scientific workflows

R

Biomolecular interactions

ABSTRACT

Quantitative analysis of the scientific literature is a frequent task in bibliometrics. Several large online resources collect and disseminate bibliographic information, paving the way for broad analyses and statistics. The Europe PubMed Central (PMC) and its Web Services is one of these resources, providing a rich platform to retrieve information and metadata on scientific publications. However, a complete bibliometric analysis that involves gathering information and deriving statistics on an author, topic, or country is laborious when consuming Web Services on the command-line or using low level automation. In contrast, scientific workflow managers can integrate different types of software tools to automate multi-step processes. The Taverna workflow engine is a popular open-source scientific workflow manager, giving easy access to available Web Services. In this tutorial, we demonstrate how to design scientific workflows for bibliometric analyses in Taverna by integrating Europe PubMed Central Web Services and statistical analysis tools. To our knowledge, this is also the first time scientific workflow managers have been used to perform bibliometric analyses using these Web Services.

© 2016 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

1. Introduction

As science becomes more data intensive, access to data and the process of generating meaningful information from them become the main vehicle in the scientific process. In this process, the primary challenge is moving from generated or retrieved data to information. As in most fields, typical bibliometric analysis workflows require several discrete steps, each employing different software tools. Frameworks that allow users to efficiently but easily connect data access points to information generation play a key role here. However, it is not always straightforward to use a generic framework or design custom workflows every time a new analysis protocol is to be implemented. In the absence of a framework, users have to manually connect the inputs and outputs of individual steps through the entire analysis. This risks introducing errors and makes analyses difficult to reproduce, especially for other researchers.

Scientific workflow managers integrate several processing units to automate a data analysis procedure. They are field-independent, so analysis on data from any field, including bibliometrics, can be automated. Scientific workflows typically

* Corresponding author.

E-mail addresses: a.t.guler@lumc.nl (A.T. Guler), cathelijmwaaijer@gmail.com (C.J.F. Waaijer), y.mohammed@lumc.nl (Y. Mohammed), m.palmblad@lumc.nl (M. Palmblad).

have inputs and outputs, where series of operations are performed on the inputs in order to produce the outputs. Thus various atomic processing units can be assembled to produce an analysis protocol that can run without manual intervention (de Bruin, Deelder, & Palmblad, 2012). On the other hand, reusability and reproducibility are also important for *in silico* experiments, facilitating collaboration and combining efforts. These are promoted by online scientific workflow repositories such as myExperiment (Goble et al., 2010). However, deciphering the hierarchical composition of a workflow, its control and connections could be difficult in a larger-scale workflow (Lu and Zhang, 2009). Taking a modular approach and defining the scope of each module in the workflow eases this process. Most of the freely available scientific workflow managers have a graphical user interface that helps to visualize the overall protocol, both when designing and when executing the workflow. Galaxy (Goecks, Nekrutenko, & Taylor, 2010), KNIME (Berthold et al., 2008) and Taverna (Oinn et al., 2004) are popular examples of such scientific workflow managers that also allow modular design. Automating an analysis consisting of several steps, such as in bibliometrics, using scientific workflow managers makes the process less laborious and decreases the risk of human errors. Scientific workflow managers follow a different paradigm than interactive software tools, such as the domain-specific (or perhaps domain-limited) BibExcel (Persson, 2016), Publish or Perish (Harzing, 2007) and Sci2 (Sci2 Team, 2009) though Sci2 certainly provides some aspects of the modularity and tool integration of the workflow managers.

We have previously presented how scientific workflows can be used to solve simple bibliometrics problems, using Taverna Workbench (Guler, Waaijer, & Palmblad, 2016). Like any other scientific workflow manager, Taverna enables the user to integrate different types of components. What makes Taverna very useful for bibliometrics is that it already provides custom support for a number of tools and services that are easily adopted for performing such analyses, e.g., R tools and XPath, Beanshell and WSDL services. The programming language R is primarily developed for statistical computing and visualization. Specific R plug-ins or packages expands its capabilities to machine learning, text mining and natural language processing (Feinerer, Hornik, & Meyer, 2008; Hornik, 2015). The XPath service is a user-friendly tool for creating XPath queries to parse XML documents by simply selecting nodes from an XML tree with a few mouse clicks. This is highly convenient, as most bibliometric databases can export information in XML format. For tabular formats, the Spreadsheet import service provides a similarly minimalistic tool for parsing tables. For general tasks, Beanshell services allow inclusion of scripts using a Java-like language. Last but not least, integrated support for Web Services allows Taverna workflows to directly communicate with remote databases using WSDL queries (Wolstencroft et al., 2013). As most Web Services use XML as the preferred message format, the Taverna XPath service is typically used to parse the results returned from Web Service calls.

An important functional aspect of Taverna is that iterations over individual processes or parts of the workflows are done implicitly by list handling. This feature provides great flexibility if a process or a sub-workflow has more than one input port. The user can specify whether the inputs are subjected to a “cross product” (all list elements in one input against all list elements in the other input) or a “dot product” (element-wise), or for processors with more than two inputs a combination of both; all while being able to define the order and precedence of the workflow operations on these input lists. A core set of built-in features and services provides basic list handling, such as flattening, merging a list to a string and removing duplicates.

Here we present a tutorial on how to use Taverna to build workflows that interact with the Europe PubMed Central Web Services. In principle, Taverna could interact with any Web Service that provide a SOAP or RESTful interface. The reason we are demonstrating the integration of Web Services in Taverna using PubMed rather than Scopus® or Web of Science™ (Falagas, Pitsouni, Malietzis, & Pappas, 2008) is that, among these three, PubMed is currently the only that provides a free Web Service interface. PubMed is also the most used bibliographic resource in the life sciences. In this tutorial, we show how to retrieve information using Web Services, how to parse this information, and how to use the various built-in Taverna services and processors to calculate and visualize the results. In principle, the same approach could be taken using other resources, provided that the user has access to them. We also made an example Taverna interface for connecting to the Thomson Reuters Web of Science™ Web Services and made this available on myExperiment (<http://www.myexperiment.org/workflows/4705.html>). We have built and tested the workflows in Taverna Workbench Bioinformatics 2.5.0, but in principle the workflows should run in any flavor of Taverna Workbench version 2.4.0 or later. For instructions on how to download and install Taverna, see www.taverna.org.uk. For Rshells to be executable in Taverna, R, Rserve and required R packages must be installed and deployed (Williams, 2014).

2. Getting started: connecting to Europe PMC Web Services

Europe PubMed Central, or PMC (<http://europepmc.org>) is one of the leading databases for peer-reviewed life science literature, providing access to 30.4 million abstracts and 3.3 million full-text articles and metadata (December 14, 2015). The goal of Europe PMC is to “build open, full-text scientific literature resources and support innovation by engaging users, enabling contributors and integrating related research data” (The Europe PMC Consortium, 2015). This is achieved by providing access through a user-friendly Web interface, FTP, and SOAP and RESTful Web Service APIs. Here we will use the latter from within Taverna workflows. This is done as follows. First, the Europe PMC SOAP-based Web Services are imported into Taverna using “Import new services” in the Design pane using the WSDL <http://www.ebi.ac.uk/europepmc/webservices/soap?wsdl>. The available Web Services should now be listed as available in Taverna services menu. The 55-page Europe PMC SOAP Web Service Reference Guide (Europe PMC, 2015) describes all details of the API to these newly imported services. Although strongly recommended, it is not absolutely necessary to read the entire manual before starting to integrate Europe PMC Web Services from within Taverna. A Web Service component is simply added to a workflow by dragging it from the

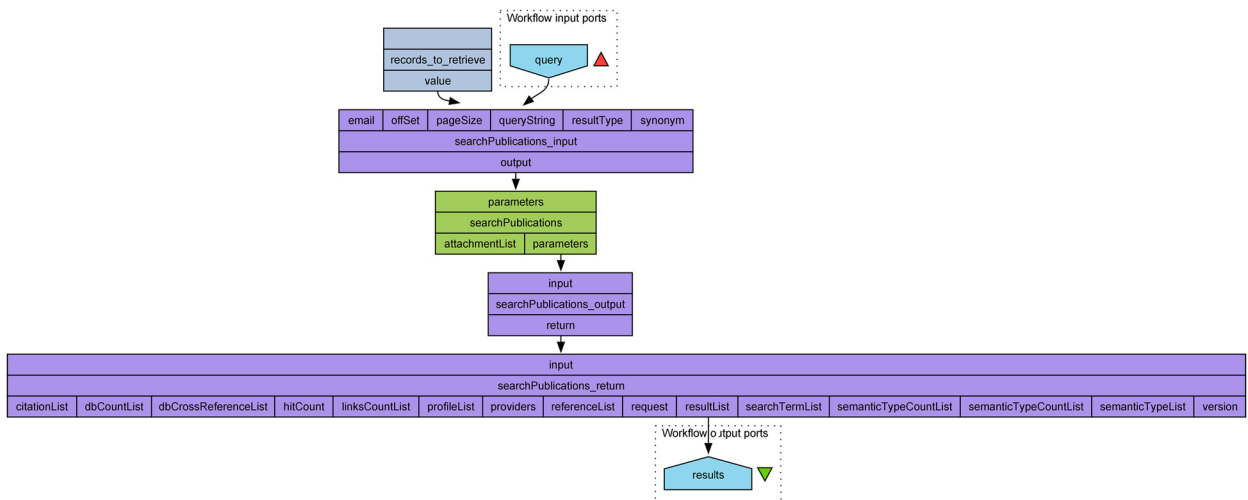


Fig. 1. Basic workflow to access Europe PMC Web Services.

service menu and dropping it into the workflow whiteboard. To expose the component's inputs and outputs, we add XML splitters. These are found in the component Edit menu. For example, the *searchPublications* service currently has six input ports: *email*, *offset*, *pageSize*, *queryString*, *resultType* and *synonym*. Of these, only the *queryString* is mandatory. This string corresponds to what one would normally enter in the search field on the Europe PMC website. The *email* address registers the user with Europe PMC, the *pageSize* the number of entries to be retrieved in one page, the *offset* refers to which page of size *pageSize* to retrieve, *synonym* whether to expand the query using the MeSH and UniProt synonyms. The *resultType* is used to limit the retrieval of results to the data we want. It has three settings: *idlist*, *lite* and *core*. If we only want the PubMed IDs (PMIDs) for subsequent queries, *idlist* would be sufficient. The *lite* results contain key metadata such as the author list and basic bibliographic information, and *core* all metadata, including abstracts and full journal details. The full article, if in Europe PMC, is retrieved using another service, *getFulltextXML*. The workflow in Fig. 1 illustrates the use of the *searchPublications* service with its input and output XML splitters. The workflow performs a single search similar to using a Web browser and the Europe PMC Website. The *results* is an XML tree with the first 100 results of the Europe PMC search defined by *query* where the number 100 is defined by *records_to_retrieve* constant. This workflow is available on myExperiment (www.myexperiment.org/workflows/4780.html). In Fig. 1, all input and output ports of every workflow components are shown. In subsequent workflows, the ports details are hidden for simplicity. However, these can easily be displayed in Taverna Workbench.

The Europe PMC results are retrieved in XML, and the extraction of the precise information we want are done by further XML output splitters or XPath services in Taverna. An XPath is a query written in the XPath language for selecting elements and attributes in an XML document. XPath allows postfix conditional statements within square brackets. For example, to restrict the results to PMIDs of cited papers (having a *citedByCount* larger than 0), the XPath `resultList/result[citedByCount > 0]/pubYear` could be used on the output of the workflow in Fig. 1 to retrieve the publication year (*pubYear*) for cited papers only. The XPath service in Taverna provides a configuration pane to automatically generate simple XPath expressions, which the user can then customize, for example by adding conditionals, or combining several expressions. The results of the Web Service and XML parsers can be passed to other workflow components either as text or as XML. A second workflow attaching an XPath statement to the workflow in Fig. 1 is also available on myExperiment (www.myexperiment.org/workflows/4779.html). The output of the workflow, after parsing the *searchPublications* output *resultList* with the XPath above, is a Taverna list of the publication years of the cited articles among the first 100 retrieved articles matching the search query *query*. As mentioned in the introduction, Taverna does iteration implicitly using lists. If a component for performing a certain task is given a list as input, the task will be performed on all elements in that list.

3. Publication records and citation networks

From these simple first steps, and using the same types of components, we will now construct more advanced workflows exploring the full power of the Europe PMC Web Services and Taverna. We do this using the notions of embedding and extensibility of scientific workflows. A simple workflow can be embedded in more complex workflows. Existing workflows, shared in the myExperiment community, can be accessed from the myExperiment pane in Taverna and modified or extended according to the user's needs.

The well-known Thomson Reuters Web of Science™ search provides a link to a "Citation Report" with two histograms, one over the number of published items in each year and one over the citations for these items in each year, based on the search results. In addition, the Citation Report provides simple statistics, such as average citations and the *h*-index for these




Fig. 2. A workflow to generate simple statistics of the citation related to a specific author.

search results (the *h*-index may be most relevant for a single author name search, but is calculated and reported for any set of publications). We can produce similar histograms based on the Europe PMC database using a Taverna workflow. For this, it is necessary to use two Web Services, *searchPublications* as before, and *getCitations* to get the publication year of papers citing the papers returned by the *searchPublications* query (for example on an author name). Fig. 2 shows such a workflow, which is also available on myExperiment (www.myexperiment.org/workflows/4720.html). The workflow uses two Europe PMC Web Services: *searchPublications* and *getCitations* to generate publication statistics in the form of a “citation report” for a particular author.

The workflow in Fig. 2 takes as input the full name of an author. This argument is passed to a Java BeanShell *build_query* that constructs the specific query “auth:\” “+author_name+”\” sort_date:y”. Combined with a value 1000 for the number of *articles_to_retrieve*, this will request the 1000 most recent publications (sorted by date) for the author or authors matching *author_name*. The list of PMIDs returned by *searchPublications* is used as input to *getCitations*, which returns a list of lists of publication years for the papers citing the papers returned by *searchPublications*. In this workflow, XPath is directly applied on the Web Services results. This skips the two XML output splitters and simplifies the visual appearance of the workflow. Whether to use output splitters and short XPath, or longer XPath directly on the Web Service output, is mostly a matter of taste. The XPath extracting the *pubYear* for the citing articles produces a list of lists of publication years as output. In order to make a combined histogram over all citations to all papers from the author, we flatten this list of lists of publication years to a single list using the built-in *Flatten_List* local service. This single list of publication years is then passed to an Rshell component *draw_histogram* as data of the (semantic) type “integer vector”, as specified in the input port to this workflow component. The integer type in R exists to pass data to programs written in strongly typed languages that expects them, and so that integer data can be represented “exactly and compactly” (Becker, Chambers, & Wilks, 1988). In this workflow, the publication years could just as well be passed as “numeric” vectors. The Rshell is very simple and uses the *hist()* function (Venables and Ripley, 2002) to generate the two histograms. For authors having a unique identifier, such as an ORCID, the *build_query* can be changed to “authorid:\” “+author_id+”\” sort_date:y”.

An output of this workflow for the author “Jonas Bergquist” (Professor Jonas Bergquist, Department of Chemistry—Biomedical Centre, Uppsala University, Sweden) is shown in Fig. 3. An extended version of this workflow is available on myExperiment (www.myexperiment.org/workflows/4781.html) that combines the publications and citations records in a two-dimensional heatmap showing the delay, increase and decrease of citations for papers over time. The work-




Fig. 3. Citation report for an author ("Jonas Bergquist") generated by the workflow in Fig. 2.

flow can easily be extended to accept a list of authors rather than a single author, generating either combined statistics or individual citation reports for each author in the list.

Suppose instead we are interested in who is cited by whom or citing the work of a particular researcher and how these authors in turn cite each other. To put it more simply: we would like to construct and visualize a co-citation network based on one researcher. Any network consists of multiple items (vertices or nodes) and their underlying relationships (edges). In the case of our co-citation network, the vertices are the single researcher and the authors cited by or citing this particular researcher. The edges are all the citations to and from the researchers in the co-citation network.

Constructing and visualizing this co-citation graph requires a slightly more elaborate workflow. Dividing up the task into smaller ones, we first look up all publications for the author using *searchPublications*. For each publication in the returned list, we then look up the references and citations in parallel using *getReferences* and *getCitations* respectively. These three requests returns all vertices in our co-citation graph, but does not retrieve citations, or edges, between papers by other authors. To retrieve these we combine all the vertices and call *getReferences* and *getCitations* again, once for each vertex. If the reference or citation is already represented by a vertex in the graph, we add a new edge to or from that vertex. To make the analysis more interesting, we can also have the workflow keep track of the author's own papers and self-citations. The best way to do this is by defining attributes to the edges and vertices in the co-citation graph. The workflow in Fig. 4 does this by generating a description of the graph in Pajek (Batagelj & Mrvar, 1998) format using the BeanShell *combine_and_make_Pajek_file*. In most workflows, one would normally strive to use stream data between components or use simple tabular of XML file formats. When dealing with graphs, however, it is sensible to use a common format for defining graphs, such as GraphML (Eiglsperger, Brandes, Lerner, & Pich, 2013), GML (Himsolt, 1997), LGL (Adai, Date, Wieland, & Marcotte, 2004) or Pajek. The workflow in Fig. 4 finds all papers citing and cited in articles published by an author, and all citations between them. The workflow generates a citation network graph that is captured by and displayed inside Taverna. The workflow also generates a Pajek file incorporating the information on self-citations using different edge attributes (color) and labels the vertices differently for the author (last name and publication year) than for the other vertices (PubMed ID).

The Pajek content created by *combine_and_make_Pajek_file* and written to file by *Write.Text.File* is read by the Rshell *draw_graph* using the *igraph* R package (Csárdi and Nepusz, 2006). The *igraph* package contains functions for reading and writing graphs in several formats, including those mentioned here. The outputs of the workflow are a simplified graph in Sugiyama, Tagawa, and Toda (1981) layout created by *igraph* using *simplify()* and *layout.sugiyama()*, and the corresponding Pajek file created by *write_graph()* after simplification. A static but visual representation of the graph is captured by Taverna as well as written to a PDF file. To interactively explore and analyze the graph, the Pajek file can be opened in Pajek or a tool such as the VOSviewer (van Eck & Waltman, 2010), which are both tools for the analysis and visualization of (bibliometric) networks. The Pajek file created by the Taverna workflow (run November 30, 2015) was opened in VOSviewer 1.6.3, showing the largest set of connected items (3782) out of the 3851 vertices in this citation graph (Fig. 5; can also be opened as an interactive Java application by clicking on the <https://www.elsevier.com/books-and-journals/content-innovation/interactive-network-viewer>). The clustering was performed with clustering resolution 0.05 and minimum cluster size 50. The author's own papers are annotated with first author, last name and year, other papers with PMID. The publication record in the example above, including citations, can be visualized as a several highly interconnected and overlapping clusters, the largest of which (red) is on proteomics. In the center of this large cluster is a highly cited review by Aebersold and Mann on mass spectrometry-based proteomics (Aebersold

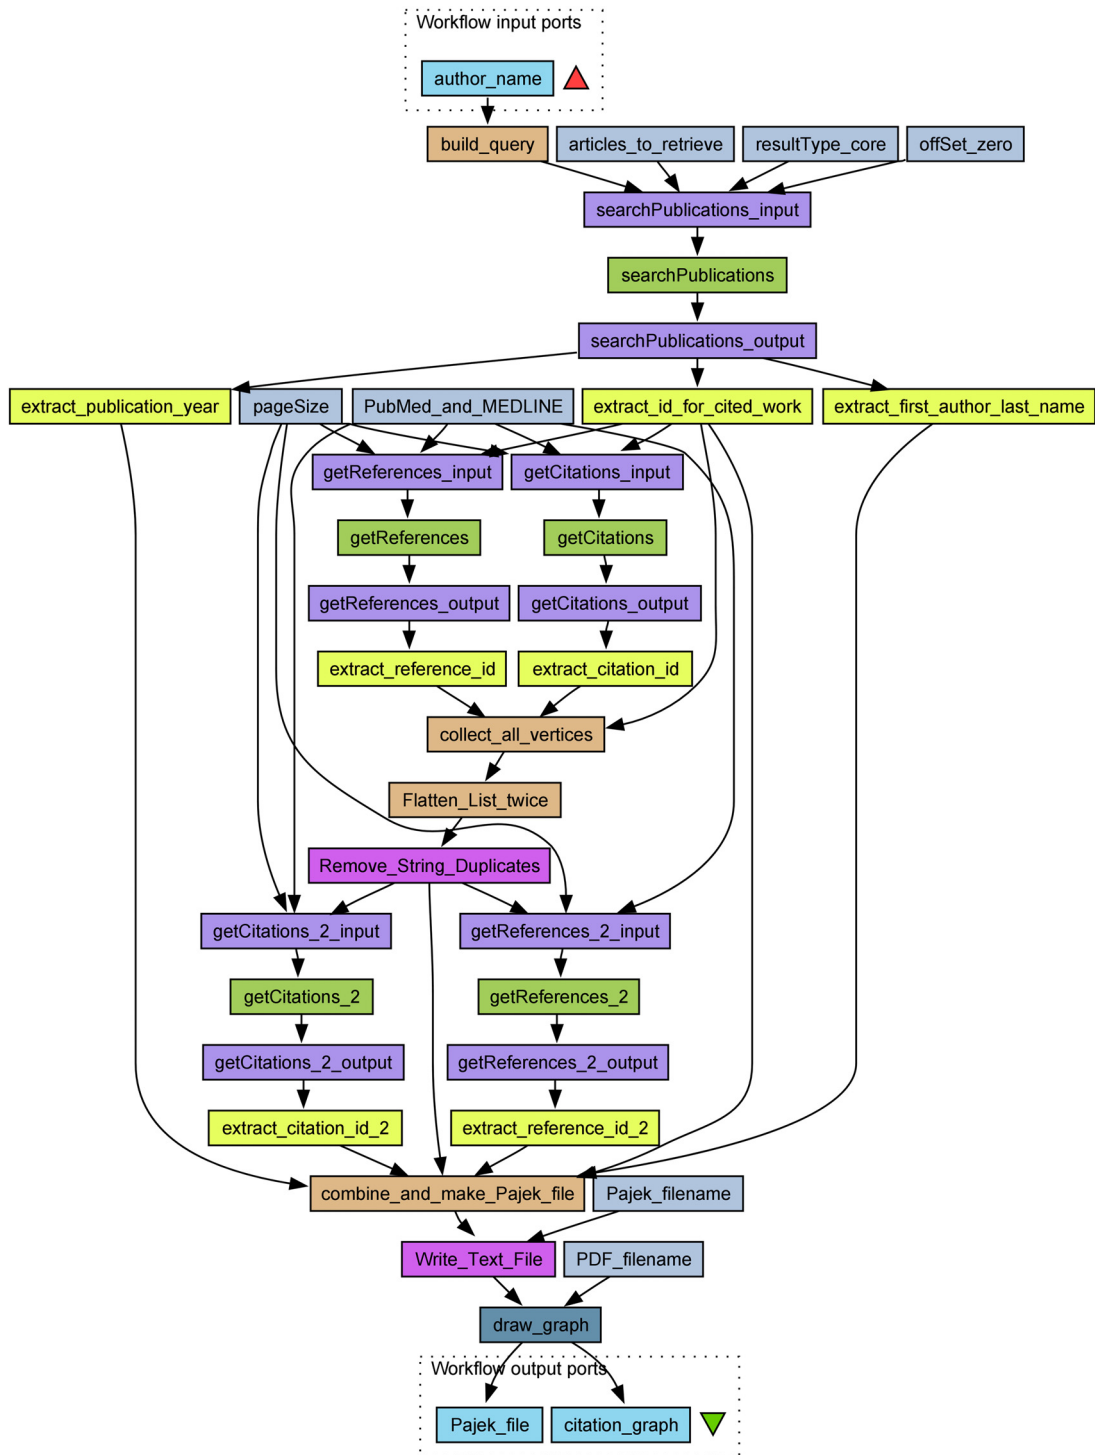


Fig. 4. A scientific workflow for generating an author citation network.

& Mann, 2003, PMID 12634793). The dark blue cluster covers work in psychophysiology and neuroscience, excluding proteomics but including new methods for analysis of cerebrospinal fluid (e.g., Wetterhall et al., 2010, Dahlin et al., 2012). In addition to this core of work in proteomics and neuroscience, we see a few protuberances representing collaborations with researchers in different disciplines, such example veterinary science applications (Holst et al., 2015, light brown) and surface

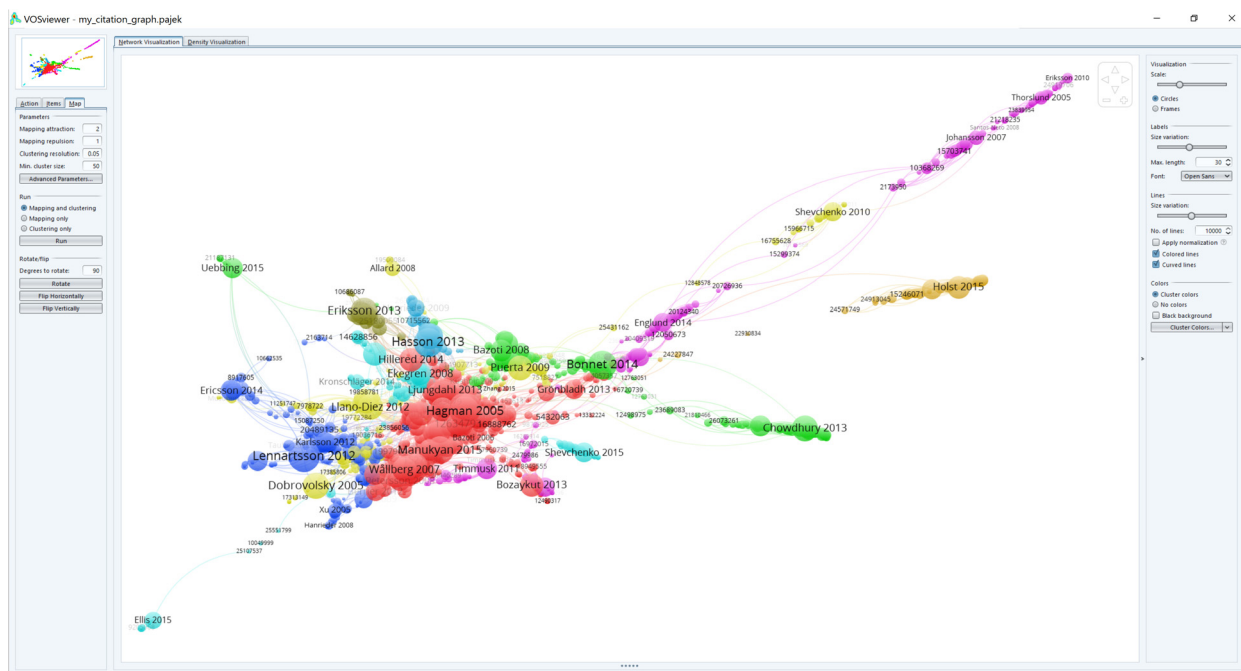


Fig. 5. Visualization of the citation network for “Jonas Bergquist” using the VOSviewer (Fig 5; can also be opened as an interactive Java application by clicking on the http://www.vosviewer.com/vosviewer.php?map=http://cpm.lumc.nl/export/bibliometrics/Jonas_Bergquist_map.txt).

chemistry techniques (Thorslund et al., 2005, Eriksson et al., 2010, magenta). The full VOSviewer map is also included as supplemental information.

Another way to view the research topics of a particular author is to count words and noun phrases in the titles and abstracts, visualizing the results as a graph or tag cloud. A workflow using the *searchPublications* Web Service and the R packages *tm* (Feinerer et al., 2008) for text mining and *wordcloud* for visualization is also available on myExperiment (<http://www.myexperiment.org/workflows/4775.html>).

4. Biomolecular interactions

Resources such as UniProt (Magrane and UniProt Consortium, 2011), IntAct (Orchard et al., 2014) and the RCSB PDB (Berman et al., 2000) provide a wealth of curated information on proteins, their functions, interactions and structures. Importantly, they always cite the original source of the information, which most commonly is a peer-reviewed scientific publication. Europe PMC is also cross-referenced to these and several other databases. The *getDatabaseLinks* Web Service is used to access the UniProtKB, IntAct or PDB records associated with an article. This may seem like a trivial service, but is in fact a programmatic access that allows us to explore the scientific literature, not only for bibliometrics, but also to investigate what the publications are about, e.g., the chemical compounds, genes, proteins, diseases or biological species. For example, consider a researcher who is interested in protein *P* and would like to find all proteins mentioned in connection with this protein in the scientific literature in a specific context. This context could be a molecular interaction, being part of the same protein complex, one protein activating the other, etcetera. Most researchers would use databases such as IntAct or UniProtKB to search for this information under the entry for protein *P*. But suppose the researcher wants to look a bit more broadly at what has been reported in the scientific literature but not yet annotated in UniProtKB, IntAct, or any other database as a specific type of protein–protein interaction. This can be accomplished using *searchPublications* and *getDatabaseLinks* in the same workflow (Fig. 6). The workflow looks up the proteins in UniProtKB most frequently co-occurring in the literature with a query protein and in a specified context, e.g., type of protein–protein interaction or disease. The workflow then builds a network with the proteins as nodes and the weights of the edges corresponding to the number of co-occurrences in the literature.

For simplicity, the input and output port splitters are embedded with the Web Services as Taverna components in the workflow in Fig. 6. The workflow builds a query string from user provided input to search for a particular UniProt identifier in the context of a certain phrase appearing in the title or abstract. The list of retrieved article identifiers (PMIDs) is then passed to *getDatabaseLinks*, which, like *getCitations*, returns a list of lists of all UniProt identifiers co-occurring in those publications. These may be very few, or number in the thousands for large proteomics studies. In general, we would expect a co-occurrence in a publication with few linked UniProt IDs to be more relevant than a co-occurrence in a list of several thousand proteins. The results can be weighted using the returned *dbCountList*, or by limiting the number of UniProt IDs




Fig. 6. Workflow to generate a protein–protein network based on co-occurrence of UniProtKB accessions.




Fig. 7. Workflow results visualized in Cytoscape with BINGO Gene Ontology Annotations.

retrieved for each PMID to a small number to reduce the influence of proteomic studies. For example, a *searchPublications* query for UniProt ID P29083, or the Transcription factor IIE alpha subunit, with the phrase “complex” in the title or abstract returns a list of PMIDs for 9 publications (November 30, 2015). Passing this list of PMIDs to *getDatabaseLinks* and specifying a *pageSize* of 10 to retrieve at most 10 identifiers per PMID produces a list of lists with a total of 62 UniProt IDs, of which 53 are unique. The workflow in Fig. 6 then counts the frequencies of these UniProt IDs and sorts them in descending order using `sort(table(UniProt_IDS), decreasing=TRUE)` in the Rshell *count.frequencies*. The protein most frequently occurring in these lists is the query protein itself (5 occurrences). The runner-up is unsurprisingly UniProt ID P29084 or the beta subunit of the Transcription factor IIE with 3 appearances. Three other UniProt identifiers occur twice and the remainder once. Raising the *pageSize* limit to the maximum allowed 1000 returns 2948 identifiers, 2550 of which are unique. Two sublists from two large-scale proteomics reports (Choudhary et al., 2009) reached the maximum of 1000 UniProt IDs, reporting 2932 and 5159 identifiers respectively. The query protein is again in the top (13 occurrences), but the beta subunit is now only in 95th place, still with only 3 co-occurrences.

The network produced by the workflow in Fig. 6 can be further analyzed in Cytoscape, a common tool for network visualization and analysis in bioinformatics (Shannon et al., 2003). The workflow output can be opened either directly as GML in Cytoscape. Fig. 7 shows Cytoscape 3.3.0 with the output from the workflow in Fig. 6 on Apolipoprotein A-I (UniProt ID P02647) and “complex” as before with the “Edge-weighted Spring Embedded” Cytoscape layout. The cluster of proteins associated with Apolipoprotein A-I was analyzed for enrichment of Gene Ontology biological processes by BINGO 3.0.3 (Maere, Heymans, & Kuiper, 2005). In Fig. 7, proteins frequently co-occurring with Apolipoprotein A-I and being involved in “macromolecular complex remodeling” (as well as “protein–lipid complex remodeling” and “plasma lipoprotein particle remodeling”) are highlighted in yellow. Again, these results are not surprising given that Apolipoprotein A-I is the dominant protein component of high density lipoprotein (the “good cholesterol”) in plasma.

The relevance of a co-occurrence (of genes or proteins) decreases with the number of co-occurring genes or proteins in a particular publication, as long lists are the results of broad proteomics studies rather than specific experiments probing the interactions of a particular protein or dissecting a specific protein–protein complex. But to take all proteins into account, a simple trick to retrieve an arbitrary number of results from any Europe PMC Web Service in a Taverna workflow is to supply the Web Service call with a sufficiently long list of *offset* values, counting from zero. This list can be created inside the *build.query* BeanShell to hide the details from the workflow view. For example, adding a simple piece of code defining a new output *offsets* as `int [] offsets=new int [] {0, 1, 2, 3, 4}`; to *build.query* and connecting the output port *offsets* to the *offset* input port of the Web Service will retrieve at most 5 pages of *pageSize* results from a Europe PMC Web Service (5000 results with the maximum *pageSize* of 1000). This approach is generally fine for literature on genes and proteins, but the Europe PMC Web Services, or Web Services generally, are not intended for piecemeal retrieval of millions of records (or the entire Europe PMC). This can better be done using the FTP access. UniProt is just one of many molecular databases linked with Europe PMC. The API to access these database links is the same for all molecular databases, all using the *getDatabaseLinks* Web Service.

5. Discussion

Most scientific workflow managers interact with the user through intuitive and visual interfaces. Like all software, Taverna also has some peculiarities. For example, to execute R scripts, a connection to RServe must first be established. The advantage is that this can be on a remote server just as easily as on the local machine, something that may be useful for computationally demanding tasks. Most workflow managers also support multiple scripting languages and types of workflow components. While this brings a lot of flexibility and power, it also makes it more difficult for those not familiar with all of these languages or services to understand the details of heterogeneous workflows. In this tutorial, we have used only Beanshell, Rshell, XPath and WSDL components. For simplicity, and because they were not needed, we did not include any local tools or shells, JSONPaths or REST services in these workflows. However, we have also uploaded a REST equivalent of the Fig. 2 workflow to myExperiment (www.myexperiment.org/workflows/4869.html).

Taverna is flexible, and can be used to organize the running of locally installed software, arrange a series of R scripts, shuffle data between external Web Services, or any combination thereof. Unlike KNIME, Taverna is free both as in 'speech' (open source) and as in 'beer' (gratis). Taverna's emphasis on Web Services makes it a perfect partner to bibliometric resources such as Europe PMC. The Taverna codebase is in Java, whereas Galaxy's is in Python. This is also reflected in the default scripting language in the workflow managers (Java in Taverna, Python in Galaxy). The programming paradigm is shared between all workflow managers however, and there have even been efforts to enact Taverna workflows through Galaxy (Karasavvas et al., 2012).

Documentation is important, in particular for sharing or collaborative development of workflows. All elements (processors, data links, inputs and outputs) in Taverna workflows can be annotated individually. These annotations follow the components when imported from one workflow to another and are found under the "Details" tab in the Service panel in Taverna Workbench. Components and connections only have a generic "Description" field whereas inputs and outputs also have an "Example" field that can be used as a default value when executing the workflow. The workflow itself has "Author" and "Title" fields, in addition to a description. Workflows can be shared on myExperiment, as we have done. When uploading a Taverna workflow, myExperiment attempts to extract workflow metadata such as title and description directly from these annotations. This works for Taverna, Galaxy and several other workflow managers. myExperiment also provide basic version control and allow users to comment on and discuss workflows. All this information can then be used to find workflows using a keyword search on the myExperiment website. Currently (May 2016); there are 3752 workflows shared on myExperiment; so it is not practical to browse all workflows to find the one closest to what one needs (or the best starting point for one's workflow).

The examples in this paper do not use any nested structures which are otherwise common in large workflows. The legibility of complex workflows such as that in Fig. 4 may be improved by boxing the Web Service calls, hiding the details of the input/output splitters and XPath and allowing the user to first grasp the overall logic of the workflow.

6. Conclusions

Bibliometric analyses often involve several steps that are carried out in different software tools. This requires much manual orchestration from one software tool to the other, which makes the process labor intensive and error prone. Scientific workflow managers, which are increasingly being used in other data intensive fields but have not yet seen widespread usage in bibliometrics, are useful tools to connect these different data retrieval and computational steps in an automated way. One such workflow manager is Taverna. In this study, we argue the direct support of Web Services, XML parsers and R in Taverna workflows make Taverna particularly useful for bibliometrics. With R comes direct access to a great number of powerful software packages such as *igraph*, *wordcloud* and *rworldmap* (South, 2011) for visualization, *tm* and *openNLP* for text mining and natural language processing. One limitation of using a scientific workflow manager such as Taverna, is that they are not meant for interactive exploration of large datasets. For this, it is more sensible to use domain-specific tools such as Pajek or VOSviewer for scientometrics, or Cytoscape for bioinformatics, as we demonstrated here.

In addition, software such as Taverna supplies repeatability and reusability to bibliometrics analyses. For example, all workflows discussed in this paper can be found in the myExperiment group for Bibliometrics and Scientometrics (www.myexperiment.org/groups/1278.html) for anyone to open and run from within Taverna, using the exactly the same or any other input parameters to define the query. Other Taverna workflows in the Bibliometrics and Scientometrics group on myExperiment use *rworldmap* to map differences in the geographic distribution of author affiliations between two PubMed search results. Such workflows can for example look at geographical patterns of research on particular diseases, or geographical bias in different journals.

The bibliometric analyses illustrated in this tutorial are exemplary of the kind of analyses we do in our research and here focused on the use of the Europe PMC Web Services. In a previous paper (Guler, Waaijer, & Palmblad, 2016) we have used Taverna for other types of bibliometric analyses, such as geographic and temporal analyses of publication patterns, word usage and co-citation analysis. Here we have shown how to access Europe PMC through a Web Service API and how to perform bibliometric analyses using the Taverna scientific workflow manager, but, more importantly, how to combine the two.

Conflict of interest

The authors declare they have no competing interests or conflicts of interest.

Author contributions

Conceived and designed the analysis: Arzu Tugce Guler, Cathelijn J. F. Waaijer, Yassene Mohammed and Magnus Palmblad.
Collected the data: Arzu Tugce Guler, Yassene Mohammed and Magnus Palmblad.

Contributed data or analysis tools: Arzu Tugce Guler, Yassene Mohammed and Magnus Palmblad.

Performed the analysis: Arzu Tugce Guler, Yassene Mohammed and Magnus Palmblad.

Wrote the paper: Arzu Tugce Guler, Cathelijn J. F. Waaijer, Yassene Mohammed and Magnus Palmblad.

Acknowledgements

The authors acknowledge Dr. Nees Jan van Eck at CWTS for advice on the VOSviewer and Professor Jonas Bergquist, Uppsala University for allowing us to use his publication record to demonstrate the workflows. The authors also wish to thank Michele Ide-Smith and Vid Vartak at the European Bioinformatics Institute for answering our questions regarding Europe PMC data.

Appendix A. Supplementary data

Supplementary data associated with this article can be found, in the online version, at <http://dx.doi.org/10.1016/j.joi.2016.05.002>.

References

- Aday, A. T., Date, S. V., Wieland, S., & Marcotte, E. M. (2004). LGL: creating a map of protein function with an algorithm for visualizing very large biological networks. *Journal of Molecular Biology*, 340, 179–190. <http://dx.doi.org/10.1016/j.jmb.2004.04.047>
- Aebersold, R., & Mann, M. (2003). Mass spectrometry-based proteomics. *Nature*, 422(6928), 198–207.
- Batagelj, V., & Mrvar, A. (1998). Pajek—program for large network analysis. *Connections*, 21(2), 47–57.
- Becker, R. A., Chambers, J. M., & Wilks, A. R. (1988). *The new S language: a programming environment for data analysis and graphics*. Monterey, CA, USA: Wadsworth and Brooks/Cole. <https://stat.ethz.ch/R-manual/R-devel/library/base/html/integer.html>
- Berman, H. M., Westbrook, J., Feng, Z., Gilliland, G., Bhat, T. N., Weissig, H., . . . & Bourne, P. E. (2000). The Protein Data Bank. *Nucleic Acids Research*, 28(1), 235–242. <http://dx.doi.org/10.1093/nar/28.1.235>
- Berthold, M., Cebron, N., Dill, F., Gabriel, T., Kötter, T., Meinl, T., . . . & Wiswedel, B. (2008). KNIME: The Konstanz Information Miner. In C. Preisach, H. Burkhardt, L. Schmidt-Thieme, & R. Decker (Eds.), *Data analysis, machine learning and applications SE–38* (pp. 319–326). Berlin, Heidelberg: Springer. http://doi.org/10.1007/978-3-540-78246-9_38
- Choudhary, C., Kumar, C., Gnad, F., Nielsen, M. L., Rehman, M., Walther, T. C., . . . & Mann, M. (2009). Lysine acetylation targets protein complexes and co-regulates major cellular functions. *Science (New York, N.Y.)*, 325(5942), 834–840. <http://dx.doi.org/10.1126/science.1175371>
- Csárdi, G., & Nepusz, T. (2006). The igraph software package for complex network research. *InterJournal Complex Systems*, 1695. Retrieved from <http://igraph.sf.net>
- Dahlin, A. P., Hjort, K., Hillered, L., Sjödin, M. O., Bergquist, J., & Wetterhall, M. (2012). Multiplexed quantification of proteins adsorbed to surface-modified and non-modified microdialysis membranes. *Analytical & Bioanalytical Chemistry*, 402(6), 2057–2067. <http://dx.doi.org/10.1007/s00216-011-5614-y>
- de Bruin, J. S., Ceelder, A. M., & Palmblad, M. (2012). Scientific workflow management in proteomics. *Molecular & Cellular Proteomics*, 11. <http://dx.doi.org/10.1074/mcp.M111.010595>
- Eiglsperger, M., Brandes, U., Lerner, J., & Pich, C. (2013). *Graph Markup Language (GraphML)*. In *Handbook of graph drawing and visualization*. pp. 517–541. CRC Press.
- Europe PMC. (2015). *EBI Europe PMC SOAP Web Service 4. 4*, (November). Retrieved from http://europepmc.org/docs/EBI.Europe_PMC.Web_Service.Reference.pdf
- Eriksson, A., Bergquist, J., Edwards, K., Hagfeldt, A., Malmström, D., & Agmo Hernández, V. (2010). Optimized protocol for on-target phosphopeptide enrichment prior to matrix-assisted laser desorption-ionization mass spectrometry using mesoporous titanium dioxide. *Analytical Chemistry*, 82(11), 4577–4583.
- The Europe PMC Consortium. (2015). Europe PMC: a full-text literature database for the life sciences and platform for innovation. *Nucleic Acids Research*, 43(D1), D1042–D1048. <http://dx.doi.org/10.1093/nar/gku1061>
- Falagas, M. E., Pitsouni, E. I., Malietzis, G. A., & Pappas, G. (2008). Comparison of PubMed, Scopus, Web of Science, and Google Scholar: strengths and weaknesses. *FASEB Journal*, 22(2), 338–342. <http://dx.doi.org/10.1096/fj.07-9492lsf>
- Feinerer, I., Hornik, K., & Meyer, D. (2008). Text mining infrastructure in R. *Journal Of Statistical Software*, 25(5), 1–54. <http://doi.org/citeulike-article-id:2842334>
- Goble, C. A., Bhagat, J., Alekseyevs, S., Cruickshank, D., Michaelides, D., Newman, D., . . . & de Roure, D. (2010). myExperiment: a repository and social network for the sharing of bioinformatics workflows. *Nucleic Acids Research*, 38(May), 677–682. <http://dx.doi.org/10.1093/nar/gkq429>
- Goecks, J., Nekrutenko, A., & Taylor, J. (2010). Galaxy: a comprehensive approach for supporting accessible, reproducible, and transparent computational research in the life sciences. *Genome Biology*, 11, R86. <http://dx.doi.org/10.1186/gb-2010-11-8-r86>
- Guler, A. T., Waaijer, C. J. F., & Palmblad, M. (2016). Scientific workflows for bibliometrics. *Scientometrics*, 107(2), 385–398. <http://dx.doi.org/10.1007/s11192-016-1885-6>
- Harzing, A. W. (2007). *Publish or Perish..* <http://www.harzing.com/pop.htm> Accessed January 2016
- Himsolt, M. (1997). GML: a portable Graph File Format. In *Technical report (University of Passau, 94030 Passau, Germany)*. Retrieved from <http://www.uni-passau.de/index.php?id=17297>
- Holst, B. S., Kushnir, M. M., & Bergquist, J. (2015). Liquid chromatography-tandem mass spectrometry (LC-MS/MS) for analysis of endogenous steroids in the luteal phase and early pregnancy in dogs: a pilot study. *Veterinary Clinical Pathology*, 44(4), 552–558. <http://dx.doi.org/10.1111/vcp.12301>
- Hornik, K. (2015). *Apache OpenNLP Tools Interface R package version 0. 2-4..* Retrieved from <http://CRAN.R-project.org/package=openNLP>
- Karasavvas, K., Wolstencroft, K., Mina, E., Cruickshank, D., Williams, A., De Roure, D., . . . & Roos, M. (2012). Opening new gateways to workflows for life scientists. *Studies in Health Technology and Informatics*, 175, 131–141.

- Lu, S., & Zhang, J. (2009). Collaborative scientific workflows. *2009 IEEE international conference on Web Services, ICWS 2009*, 527–534. <http://doi.org/10.1109/ICWS.2009.150>
- Maere, S., Heymans, K., & Kuiper, M. (2005). BiNGO: a Cytoscape plugin to assess overrepresentation of gene ontology categories in biological networks. *Bioinformatics (Oxford, England)*, *21*(16), 3448–3449. <http://dx.doi.org/10.1093/bioinformatics/bti551>
- Magrane, M., & UniProt Consortium. (2011). UniProt Knowledgebase: a hub of integrated protein data. *Database*, 1–13. <http://dx.doi.org/10.1093/database/bar009>
- Oinn, T., Addis, M., Ferris, J., Marvin, D., Senger, M., Greenwood, M., . . . & Li, P. (2004). Taverna: a tool for the composition and enactment of bioinformatics workflows. *Bioinformatics*, *20*(17), 3045–3054. <http://dx.doi.org/10.1093/bioinformatics/bth361>
- Orchard, S., Ammari, M., Aranda, B., Breuza, L., Briganti, L., Broackes-Carter, F., . . . & Hermjakob, H. (2014). The MIntAct project—IntAct as a common curation platform for 11 molecular interaction databases. *Nucleic Acids Research*, *D1*, 358–363. <http://dx.doi.org/10.1093/nar/gkt1115>
- Persson, O. (2016). *Bibexcel*. <http://www.umu.se/onforsk/Bibexcel> Accessed January 2016
- Sci2 Team. (2009). *Science of Science (Sci2) Tool*. Indiana University and SciTech Strategies. <https://sci2.cns.iu.edu>
- Shannon, P., Markiel, A., Ozier, O., Baliga, N. S., Wang, J. T., Ramage, D., . . . & Ideker, T. (2003). Cytoscape: a software environment for integrated models of biomolecular interaction networks. *Genome Research*, *2498–2504*. <http://dx.doi.org/10.1101/gr.1239303.metabolite>
- South, A. (2011). *rworldmap: a new R package for mapping global data*. *The R Journal*, *3*(1), 35–43.
- Sugiyama, K., Tagawa, S., & Toda, M. (1981). *Methods for visual understanding of hierarchical system structures*. *IEEE Transactions on Systems, Man, and Cybernetics*, *SMC-11*(2), 109–125.
- Thorslund, S., Lindberg, P., Andrén, P. E., Nikolajeff, F., & Bergquist, J. (2005). Electrokinetic-driven microfluidic system in poly(dimethylsiloxane) for mass spectrometry detection integrating sample injection, capillary electrophoresis, and electrospray emitter on-chip. *Electrophoresis*, *26*(24), 4674–4683. <http://dx.doi.org/10.1002/elps.200500338>
- van Eck, N. J., & Waltman, L. (2010). Software survey: VOSviewer, a computer program for bibliometric mapping. *Scientometrics*, *84*(2), 523–538. <http://dx.doi.org/10.1007/s11192-009-0146-3>
- Venables, W. N., & Ripley, B. D. (2002). *MASS: modern applied statistics with S*. March. <http://doi.org/10.1198/tech.2003.s33>
- Wetterhall, M., Zuberovic, A., Hanrieder, J., & Bergquist, J. (2010). Assessment of the partitioning capacity of high abundant proteins in human cerebrospinal fluid using affinity and immunoaffinity subtraction spin columns. *Journal of Chromatography B Analytical Technologies in the Biomedical and Life Sciences*, *878*(19), 1519–1530. <http://dx.doi.org/10.1016/j.jchromb.2010.04.003>
- Williams, A. (2014). *Rserve installation and deployment*. <http://dev.mygrid.org.uk/wiki/display/tav250/Rserve+installation+and+deployment> Accessed May 2016
- Wolstencroft, K., Haines, R., Fellows, D., Williams, A., Withers, D., Owen, S., . . . & Goble, C. (2013). Taverna workflow suite: designing and executing workflows of Web Services on the desktop, web or in the cloud. *Nucleic Acids Research*, *41*(May), 557–561. <http://dx.doi.org/10.1093/nar/gkt328>