

# Biblio-MetReS for user-friendly mining of genes and biological processes in scientific documents

Anabel Usie<sup>1,2</sup>, Hiren Karathia<sup>1</sup>, Ivan Teixidó<sup>2</sup>, Rui Alves<sup>1</sup> and Francesc Solsona<sup>2</sup>

<sup>1</sup> Department of Basic Medical Sciences, Edifici Recerca Biomedica I, Universitat de Lleida and IRBLleida, Lleida, Spain

<sup>2</sup> Department of Computer Science, Escola Politècnica Superior and INSPIRES, Universitat de Lleida, Lleida, Spain

## ABSTRACT

One way to initiate the reconstruction of molecular circuits is by using automated text-mining techniques. Developing more efficient methods for such reconstruction is a topic of active research, and those methods are typically included by bioinformaticians in pipelines used to mine and curate large literature datasets. Nevertheless, experimental biologists have a limited number of available user-friendly tools that use text-mining for network reconstruction and require no programming skills to use. One of these tools is Biblio-MetReS. Originally, this tool permitted an on-the-fly analysis of documents contained in a number of web-based literature databases to identify co-occurrence of proteins/genes. This approach ensured results that were always up-to-date with the latest live version of the databases. However, this 'up-to-dateness' came at the cost of large execution times. Here we report an evolution of the application Biblio-MetReS that permits constructing co-occurrence networks for genes, GO processes, Pathways, or any combination of the three types of entities and graphically represent those entities. We show that the performance of Biblio-MetReS in identifying gene co-occurrence is as least as good as that of other comparable applications (STRING and iHOP). In addition, we also show that the identification of GO processes is on par to that reported in the latest BioCreAtIvE challenge. Finally, we also report the implementation of a new strategy that combines on-the-fly analysis of new documents with preprocessed information from documents that were encountered in previous analyses. This combination simultaneously decreases program run time and maintains 'up-to-dateness' of the results. Availability: <http://metres.udl.cat/index.php/downloads>, Contact: [metres.cmb@gmail.com](mailto:metres.cmb@gmail.com).

Submitted 4 December 2013

Accepted 27 January 2014

Published 27 February 2014

Corresponding author

Rui Alves, [ralves@cmb.udl.cat](mailto:ralves@cmb.udl.cat),

Francesc Solsona,

[francesc@diei.udl.cat](mailto:francesc@diei.udl.cat)

Academic editor

Julie Thompson

Additional Information and  
Declarations can be found on  
page 10

DOI [10.7717/peerj.276](https://doi.org/10.7717/peerj.276)

© Copyright

2014 Usie et al.

Distributed under

Creative-Commons CC-BY 3.0

**OPEN ACCESS**

**Subjects** Bioinformatics, Computational Biology, Molecular Biology

**Keywords** Network reconstruction, Systems biology, Literature analysis

## INTRODUCTION

The reconstruction of molecular circuits is an important research goal in the biological sciences. One of the ways to achieve that reconstruction starts with the use of automated text-mining techniques to identify networks of genes that co-occur in scientific

documents. Subsequent human curation of these co-occurrence networks can then lead to accurate circuit reconstruction.

In this process, the automated identification of co-occurrence gene networks is crucial because databases of relevant scientific documents contain many more entries than those that can be manually analyzed (*Alves & Sorribas, 2007; Markowetz & Spang, 2007; Arighi et al., 2013; Krallinger et al., 2013*). A gold standard of these databases, MEDLINE, contains more than  $19 \times 10^6$  records, with 2000–4000 new entries being added each day (*NCBI, 2013*). Extracting biological information from such large databases requires text-mining methods and tools that are able to automatically integrate and summarize useful biological information across the database records.

The development of text-mining methods that enable circuit reconstruction from scientific documents is an area of active development (*Camon et al., 2005; Hoffmann & Valencia, 2005; Huang et al., 2008; Chen, Liu & Manderick, 2010; Arighi et al., 2011; Kano et al., 2011; Kim et al., 2011; Szklarczyk et al., 2011; Usié et al., 2011; Bossy et al., 2012; Kim et al., 2012; Pyysalo et al., 2012; Arighi et al., 2013; Krallinger et al., 2013*). The performance of those methods for automated identification of the circuits (*Camon et al., 2005*), of their components (genes/proteins), and of the inter-component relationships, has been systematically evaluated over the last few years, for example through the BioNLP (*Kano et al., 2011; Kim et al., 2011; Bossy et al., 2012; Kim et al., 2012; Pyysalo et al., 2012*) and BioCreAtIvE initiatives (*Huang et al., 2008; Chen, Liu & Manderick, 2010; Arighi et al., 2011; Wu et al., 2012; Arighi et al., 2013; Krallinger et al., 2013*).

To briefly summarize, there are three large classes of methods that have been employed for the reconstruction of co-occurrence gene networks: dictionary-based methods, morphology-based methods, and context-based methods (*Vazquez et al., 2011*). Dictionary-based methods rely on matching compiled lists of terms to their appearances in the text of documents (*Yang, Lin & Li, 2008*). Morphology-based methods rely on the morphological structure of specific classes of words to single them out in documents (*Malouf, 2002; Peng & Schuurmans, 2003*). Finally, context-based methods can be divided into Machine Learning or Natural Language Processing techniques: The former identify patterns in the structure of the text that help to recognize the presence of the relevant entities in documents; the latter draw on our knowledge about the grammar and syntax rules of natural languages to recognize those entities. These three general approaches can be combined in order to improve NER (for example see *Arighi et al., 2013; Krallinger et al., 2013* and references therein).

In general, methods participating in evaluations such as BioCreAtIvE or BioNLP are implemented in tools that can be included in web-services and assist curators in the maintenance of large databases of biological knowledge. Examples of this are given in *Arighi et al. (2013), Krallinger et al. (2013)*. In most cases, using these methods and tools requires that one becomes an expert in computer programming.

Experimental scientists that are interested in being users of, without becoming experts in, text-mining methods to directly reconstruct networks of gene co-occurrence for their genes of interest in scientific documents have a much smaller set of available

user-friendly tools. The first that became available was iHOP ([Hoffmann & Valencia, 2005](#)), which was later joined by STRING ([Franceschini et al., 2013](#)). These user friendly and intuitive web applications allow anyone to reconstruct the network of co-occurrences contained in Medline abstracts and/or Pubmed documents. Users of these applications face two important limitations. First, the applications rely on preprocessed versions of the Medline/Pubmed databases, which means that searches are fast but results are always out of date. Second, the coverage of full text documents by the applications is, at best, limited.

Recognizing these limitations, Biblio-MetReS (Bibliometric Metabolic network Reconstruction Server [Usié et al., 2011](#)) was implemented for the same target audience as STRING or iHOP, but relying on two differential features with respect to those applications. The first was that it would search databases and analyze documents on the run, thus providing the users with the most up-to-date results available on the web. The second was that full text documents were also analyzed, as were other databases besides Medline/Pubmed. These two features made Biblio-MetReS significantly slower than STRING and iHOP.

Here, we report an evolution of the application Biblio-MetReS that permits constructing co-occurrence networks for genes, GO processes, Pathways, or any combination of the three types of entities and graphically represent those entities. No other user-friendly application that we are aware of simultaneously allows the type of mixed analysis and graphical representation afforded by Biblio-MetReS. We show in a comparative analysis that the performance of Biblio-MetReS in identifying gene co-occurrence is as least as good as that of other comparable applications (STRING and iHOP). In addition, we also show that the identification of GO processes is on par to that reported in the latest BioCreAtIvE challenge ([Arighi et al., 2013](#)). Finally, we also report the implementation of a new strategy that combines on-the-fly analysis of new documents with preprocessed information from documents that were encountered in previous analyses. This combination simultaneously decreases program run time and maintains 'up-to-dateness' of the results.

## METHODS

### Organism selection

Biblio-MetReS is organism-centric. Users must select their organism of interest from a list of more than 1200 organisms with fully sequenced and annotated genomes before starting any search. They must also decide whether they want to perform GO and/or pathway term co-occurrence analysis. After these decisions are made, the program loads the genes for the organism from the program's central database. If selected, terms from the GO classification, KEGG and/or Panther pathways are also loaded into the application's front end.

### Document analysis

To analyze networks of co-occurrence Biblio-MetReS needs users to select at least one gene from their organism of interest and at least one database in which to perform

document analysis. The seed list of genes is then used to identify relevant documents in the selected database(s) of documents. The text in the flagged documents is then analyzed to identify additional genes from the organism of interest, as well as GO and/or Pathway terms. This allows users to identify co-occurrence among GO/Pathway entities and between GO/Pathway entities and gene/protein entities in sentences, paragraphs or documents. For further description of this procedure, please see section 1 of the [Supplementary Materials](#).

Biblio-MetReS uses exact matching of gene names to an internal dictionary of synonyms to identify co-mentions of genes/proteins in the text of scientific records. The gene synonyms are those officially defined by NCBI. For any given gene, all synonyms are searched for in the text of flagged documents. Similarly, Biblio-MetReS uses exact matching of GO terms to an internal dictionary of synonyms defined by the gene ontology consortium (*Gene Ontology, 2013*) to identify co-mentions of GO terms in the text of scientific records. The same exact matching is done to identify mentions to entities from the complete joint sets of KEGG (*Kotera et al., 2012*) and Panther pathways (*Mi et al., 2005*).

Co-occurrence of any two terms is analyzed at three levels. First, all possible pairs of terms are searched for in each document as a whole. Then, each document is divided into paragraphs and the pairs of terms identified in the document are searched for within each paragraph. Finally, each paragraph is divided into single sentences and the pairs of terms identified in that paragraph are searched for within each sentence. Within each level, the distance between each term in the pair is not taken into account.

The database containing the organisms and their gene names, as well as the GO terms, is updated every three months using information compiled automatically from NCBI and GO.

### Calculating the significance of term co-occurrences

To attribute statistical significance to the co-occurrence of a pair of genes, pathway terms, GO terms, gene-pathway terms, gene-GO terms or GO-pathway terms, we calculate several metrics. First, we measure how frequently the different pairs co-occur in sentences, paragraphs and/or documents. We then take the odds ratio of the frequency of occurrences in the first two categories with respect to that of the third. The closer to one these odds ratios are, the more frequent it is that both genes are mentioned only in the same sentences or paragraphs of a document, rather than appearing haphazardly in different sections of the text.

Second, we calculate how much information we gain by having two terms,  $T_i$  and  $T_j$ , co-occur, when compared to the individual occurrences of the terms. To estimate this we use information theory. The individual probability of occurrence of a term is denoted as  $p(T_i)$  and it is formally defined as  $p(T_i) = a/n$ , where  $a$  is the number of documents where Term  $T_i$  appears, and  $n$  is the total number of documents. The joint probability of co-occurrence of two terms,  $p(T_i, T_j)$ , is defined as  $p(T_i, T_j) = b/n$ , where  $b$  is the

number of documents where terms  $T_i$  and  $T_j$  simultaneously appear, and  $n$  is the total number of documents. The mutual information,  $MI(T_i, T_j)$ , is then calculated as follows:

$$MI(T_i, T_j) = p(T_i, T_j) \times \log(p(T_i, T_j)/(p(T_i)p(T_j)))$$

Finally, and in order to attribute some form of statistical significance to the co-occurrence of a pair of terms, we do as follows. Consider a set of  $n$  sentences (paragraphs, documents)  $[1, \dots, n]$ . For a given term  $k$  define

$$y_{i,k} = \begin{cases} 1 & \Leftarrow \text{term } k \text{ occurs } \in \text{ sentence (paragraph, document) } i \\ 0 & \Leftarrow \text{otherwise} \end{cases}$$

Now, for terms  $k_1$  and  $k_2$  define  $\varphi_{k_1,k_2} = y_{i,k_1} \times y_{i,k_2}$  which has value 1 when both terms co-occur and 0 otherwise.

Both these variables have a Bernoulli distribution. If the occurrence of terms  $k_1$  and  $k_2$  is independent, then  $p(\varphi_{k_1,k_2}) = p(y_{k_1})p(y_{k_2})$  would be expected, where  $p(y_{k,\cdot})$  is the relative frequency of occurrence of term  $y_{k,\cdot}$  and  $p(\varphi_{k_1,k_2})$  is the relative frequency of co-occurrence of terms  $k_1$  and  $k_2$  in the total number  $n$  of sentences (paragraphs, documents). Then, a Pearson statistic can be used to test for independence of occurrence between  $k_1$  and  $k_2$  by comparing the observed frequencies,  $n_1 = n \times p(\varphi_{k_1,k_2})$  and  $n_2 = (1 - p(\varphi_{k_1,k_2})) \times n$ , with the expected frequencies under the null hypothesis of independence, which would be  $m_1 = n \times p(y_{k_1}) \times p(y_{k_2})$  and  $m_2 = n \times (1 - p(y_{k_1}) \times p(y_{k_2}))$ . The Pearson statistic is computed as follows

$$X^2 = \sum_{i=1}^2 \frac{(n_i - m_i)^2}{m_i}$$

This statistics follows a chi-square distribution with one degree of freedom, i.e.  $\chi_1^2 \sim X^2$ ; hence, the  $p$ -value can be calculated as  $p = \Pr \chi_1^2 > \chi^2$  to assess whether the observed co-occurrence is higher than the one expected by pure chance.

### Precompilation strategy

Biblio-MetReS v2.0 implements a precompilation approach that works in the following way. Any search done will identify a given number of documents in the database(s) selected by the user(s), for an organism of interest. If a given document has not been found in any previous search by any user, in the context of that organism, Biblio-MetReS will analyze it as described in section 1 of [Supplementary Materials](#) and all information contained in that document and relevant for the analysis will be stored in a central database (see section 3 of [Supplementary Materials](#) for detailed information). If a given document has been previously found by any user, its information will be directly accessed from our central database, and the document will not be reanalyzed. This means that newly found documents are mined on the fly by the program to find and count mentions of relevant entities, while mentions in documents that have been previously found are simply looked up in our central database.

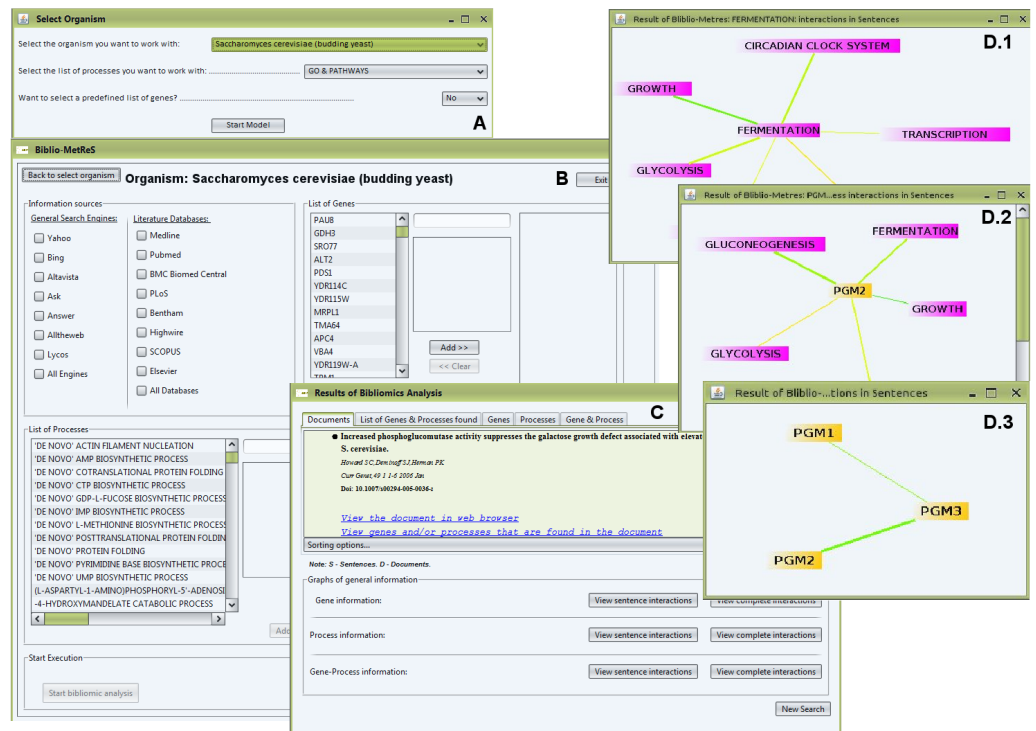
## RESULTS

### Biblio-MetReS and Biblio-MetReS player

Biblio-MetReS v2.0 can be used to identify genes/proteins from more than 1200 different organisms in records stored in a variety of databases. Users download the application and run it locally. A functioning internet connection and a local copy of JAVA are required for the program to work. Upon starting the program, users log into the central Biblio-MetReS database and choose which organism they are interested in and whether they want to search only for co-occurrence of genes/proteins or if they also want to include biological pathways and/or GO biological processes in the analysis. Once this choice is made, the program loads the necessary information from the central database. Taking this approach, instead of including all the data locally in the program installation, permits making an application that is much smaller in size and needs less RAM to function properly. Subsequently, users select the source of documents that they want to analyze, as well as the genes/proteins and/or pathways/GO biological processes that they want to search for. They can also include their own handpicked list of processes to be searched. Once the search is launched, Biblio-MetReS will identify documents in the relevant databases that contain mentions to the relevant search items. After identifying these documents, the application fully analyzes them to identify mentions for any additional gene or Pathway/GO biological process via a dictionary matching approach. The co-occurrence of the different entities is analyzed at the level of the whole document, of individual paragraphs and of individual sentences, and the significance of this co-occurrence is calculated as described in [Usié et al. \(2011\)](#). The information that is relevant for the co-occurrence calculations is stored in the central Biblio-MetReS database. Any subsequent searches that identify the same document will not reanalyze it; instead, these numbers are directly retrieved from that database. Once the analysis is complete, the users can visualize it in graphical and textual form. Links to the documents and sentences where co-occurrences are found are provided. Graphical visualization of the results can be done in different ways. Users can create graphs for the global co-occurrences network and for gene- or pathway/process-centric co-occurrences at the document/paragraph and sentence levels. Significance and Mutual Information of each co-occurrence is also provided in tabular form. The graphical representation of the networks is automatically stored in local xml files. These files can be opened using a small app, Biblio-MetReS Player, which can be downloaded from the Biblio-MetReS website. This permits reviewing previously obtained networks without having to redo the search. All this process is summarized in [Fig. 1](#).

### Comparative performance and benchmarking of new types of entities

The benchmarking of the program and its improvements with respect to version 1.0 was carried out using four organisms of interest: *Saccharomyces cerevisiae*, *Homo sapiens*, *Escherichia coli* and *Drosophila melanogaster*. For each organism we used as a search seed a set of genes belonging to Glycolysis, Lysine metabolism and RNA processing pathways. The genes were chosen to reproduce the experiments reported in [Usié et al.](#)



**Figure 1** User workflow for Biblio-MetReS.

(2011). Details are shown in Table S1. We benchmarked three different aspects of Biblio-MetReS. First, we benchmarked the comparative identification of genes between Biblio-MetReS, iHOP, and STRING, given that these three applications have similar target audiences. Second, we benchmarked the ability of Biblio-MetReS to identify biological processes/pathways. Finally, we benchmarked the improvements in Biblio-MetReS run time made by implementing the combined pre-processing/live analysis strategy.

Benchmarking the comparative identification of genes between Biblio-MetReS, iHOP, and STRING was done in the following way. We used the genes and organisms described in Table S1 to interrogate independently Biblio-MetReS, iHOP, and STRING. For each of the three applications, the complete set of results for each gene from the same pathway were pooled together for analysis. The results from STRING were further filtered to eliminate all genes and interactions that were not literature based. Table S2 compares the results for the three applications. In summary, Biblio-MetReS find the largest number of genes, followed by STRING, and iHOP. The number of genes found by STRING and Biblio-MetReS are of the same order of magnitude, while iHOP finds between one and two orders of magnitude less genes. This result derives from the fact that iHOP analyzes only Medline abstracts, while Biblio-MetReS and STRING analyze the full text of Pubmed publications, in addition to the Medline abstracts. This is confirmed by the fact that, when Biblio-MetReS is run only to analyze Medline abstracts, it finds a similar number of

genes as iHOP (data not shown). As was observed in [Table S3](#), the genes found by each of the three applications for the same experiments only partially overlap and this is explained by the different datasets analyzed by each of the programs and by partially different dictionaries of gene synonyms ([Usié et al., 2011](#)).

Neither STRING nor iHOP permit identifying GO terms and their associations to genes. Therefore we cannot perform experiments that are similar to the comparative benchmarking experiments described above. In light of this, benchmarking of Biblio-MetReS' ability to identify biological processes/pathways was done in the following way. To perform the GO identification benchmark we used the test and development sets of the BioCreAtIvE IV GO task corpus [3]. We used Biblio-MetReS dictionary matching approach to identify GO terms in the non-annotated documents and then analyzed the corresponding annotated documents. We found that Biblio-MetReS identified 100% of the annotated GO terms in both sets (2963 terms in the training set and 2243 terms in the development set). Biblio-MetReS also identifies 2259 additional GO terms in the development set and 2119 additional GO terms in the training set. For the purpose of our testing these terms must be considered false positive. Taking this into account, the precision in the development set is 50%, while in the training set it increases to 58%. The F-score performance of Biblio-MetReS is 33% in the development set and 37% in the training set, which is on par with the best approaches presented in the latest BioCreAtIvE IV challenge ([Mao et al., 2013](#)).

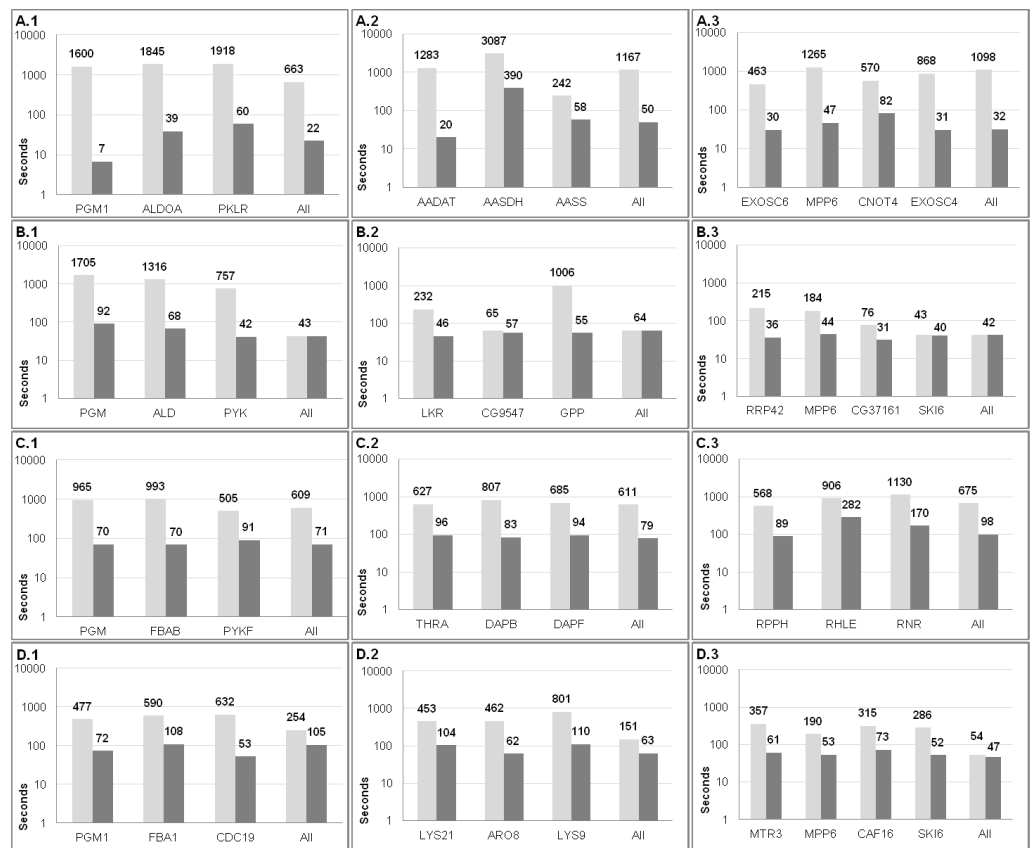
Benchmarking run time was done in two ways. First, we search only the Pubmed database. Second, we search by selecting all literature databases available in Biblio-MetReS. In both tests we used all the seed genes from [Table S1](#). Each seed is used by Biblio-MetReS as a query search. This query search is launched twice. When the first search is done there are no preprocessed documents in Biblio-MetReS' database. The information in documents is analyzed on-the-fly and stored. Then the searches are repeated, now with the documents stored in Biblio-MetReS' database. This allows us to estimate the percentage of run-time saved by preprocessing the documents. The results are shown in [Fig. 2](#) and in [Table S2](#) of the supplementary materials. On average we get decreases in run time of more than 90%.

## DISCUSSION

Here we present the new version of Biblio-MetReS, a user friendly tool for the identification of gene/protein co-occurrence networks in scientific documents. The major changes with respect to version 1.0 have to do with the search and analysis process of the documents, which can now be up to 95% faster than in the previous version. In addition, the tool now also searches for co-occurrences of biological processes and pathways, to help users to more easily establish the biological circuits in which their genes of interest may be involved in.

The methods used by the application to identify genes and proteins, as well as biological processes and pathways, in the documents are dictionary-based. These methods perform





**Figure 2** Effect of preprocessing documents on Biblio-MetReS' run time. In brief, genes from three KEGG-defined pathways are used for this test. Panels A.x show experimental results for glycolysis genes. Panels B.x show experimental results for Lysine metabolism genes. Panels C.x show experimental results for RNA processing genes. Three organisms are used in this benchmark. Panels Y.1 show results for *Homo sapiens*, panels Y.2 show results for *Drosophila melanogaster*, panels Y.3 show results for *Escherichia coli*, and panels Y.4 show results for *Saccharomyces cerevisiae*. These pathways and organisms were chosen to remain consistent with the tests performed in *Usié et al. (2011)*. Searches were done selecting all the databases in the application. Graphs can be interpreted as follows. Light gray bars indicate the run time for Biblio-MetReS when the corresponding gene is searched for the first time. In this case the program has to do a full document analysis on the fly and no information has been preprocessed. Darker gray bars indicate the run time for Biblio-MetReS when the search for the corresponding gene is repeated, and preprocessed information is already present in Biblio-MetReS' central database. The column 'All' indicates the run-time for searching all genes in the graph simultaneously, after individual searches for each gene had already been done and results preprocessed and stored.

on par with iHOP and STRING for gene and protein identification and with the best BioCreAtIvE methods for biological process identification (see [Supplementary Materials](#)).

Taken together, the new application further facilitates the identification of functional relationships between proteins and aids in identifying the biological processes and circuits in which those proteins may be involved. Although GO term search has been implemented in several literature search tools (see for example *Domis & Schroeder, 2005*;

*Plake et al., 2009*, among others), no other user-friendly tool permits simultaneous graphical reconstruction of networks of co-occurrence between genes, GO terms, and Pathway terms.

As is demonstrated by the BioCreAtIvE challenge (*Mao et al., 2013*), the problem of identifying entities in scientific texts is far from solved. Although Biblio-MetReS aims at giving non-expert users the possibility of performing such identification and use that identification to extract biological knowledge, there is much room for improvement. We are implementing an offline system to automatically search, analyze, and store information about gene/protein and pathway/biological processes co-occurrences in the documents. This will contribute to decrease the dependence of Biblio-MetReS on the users and their searches to preprocess information and make searches faster.

## ADDITIONAL INFORMATION AND DECLARATIONS

### Funding

RA was partially supported by the Ministerio de Ciencia e Innovación (MICINN, Spain through grant BFU2010-17704). FS was partially funded by the MICINN, with grants TIN2011-28689-C02-02. The authors are members of the research groups 2009SGR809 and 2009SGR145, funded by the “Generalitat de Catalunya”. AU is funded by a Generalitat de Catalunya (AGAUR) PhD fellowship. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

### Grant Disclosures

The following grant information was disclosed by the authors:  
Ministerio de Ciencia e Innovación (MICINN, Spain) grants BFU2010-17704, TIN2011-28689-C02-02.  
Generalitat de Catalunya research group grants 2009SGR809, 2009SGR145.

### Competing Interests

The authors declare that they have no competing interests.

### Author Contributions

- Anabel Usie and Rui Alves conceived and designed the experiments, performed the experiments, analyzed the data, contributed reagents/materials/analysis tools, wrote the paper.
- Hiren Karathia and Ivan Teixidó contributed reagents/materials/analysis tools.
- Francesc Solsona conceived and designed the experiments, analyzed the data, contributed reagents/materials/analysis tools, wrote the paper.

### Supplemental Information

Supplemental information for this article can be found online at <http://dx.doi.org/10.7717/peerj.276>.

## REFERENCES

- Alves R, Sorribas A. 2007. In silico pathway reconstruction: Iron-sulfur cluster biogenesis in *Saccharomyces cerevisiae*. *BMC Systems Biology* 1:10 DOI 10.1186/1752-0509-1-10.
- Arighi CN, Chohen KB, Hirschman L, Lu Z, Tudor CO, Wiegiers T, Wilbur WJ, Wu CH, eds. 2013. *Proceedings of the fourth biocreative challenge evaluation workshop* Vol 1. Washington: Biocreative Challenge1–302.
- Arighi CN, Lu Z, Krallinger M, Cohen KB, Wilbur WJ, Valencia A, Hirshman L, Wu CH. 2011. Overview of the BioCreative III workshop. *BMC Bioinformatics* 12 (Suppl 8):S1 DOI 10.1186/1471-2105-12-S8-S1.
- Bossy R, Jourde J, Manine A-P, Veber O, Alphonse E, van de Guchte M, Bessires P, Nellec C. 2012. BioNLP shared task—the bacteria track. *BMC Bioinformatics* 13 (Suppl 11):S3 DOI 10.1186/1471-2105-13-S11-S3.
- Camon EB, Barrell DG, Dimmer EC, Lee V, Magrane M, Maslen J, Binns D, Apweiler R. 2005. An evaluation of GO annotation retrieval for BioCreAtIvE and GOA. *BMC Bioinformatics* 6 (Suppl 1):S17 DOI 10.1186/1471-2105-6-S1-S17.
- Chen Y, Liu F, Manderick B. 2010. BioLMiner and the BioCreative II.5 challenge. *BMC Bioinformatics* 11(Suppl 5):P6 DOI 10.1186/1471-2105-11-S5-P6.
- Doms A, Schroeder M. 2005. GoPubMed: exploring PubMed with the Gene Ontology. *Nucleic Acids Research* 33 (Suppl 2):W783–W786 DOI 10.1093/nar/gki470.
- Franceschini A, Szklarczyk D, Frankild S, Kuhn M, Simonovic M, Roth A, Lin J, Minquez P, Bork P, von Mering C, Jensen LJ. 2013. STRING v9.1: protein–protein interaction networks, with increased coverage and integration. *Nucleic Acids Research* 41(Database issue):D808–D815 DOI 10.1093/nar/gks1094.
- Gene Ontology Consortium, Blake JA, Dolan M, Drabkin H, Hill DP, Li N, Sitnikov D, Bridges S, Burgess S, Buza T, McCarthy F, Peddinti D, Pillai L, Carbon S, Dietze H, Ire A, Lewis SE, Mungall CJ, Gaudet P, Chrisholm RL, Fey P, Kibbe WA, Basu S, Siegele DA, McIntosh BK, Renfro DP, Zweifel AE, Hu JC, Brown NH, Tweedie S, Alam-Faruque Y, Apweiler R, Auchinchloss A, Axelsen K, Bely B, Blatter M, Bonilla C, Bouguerleret L, Boutet E, Breuza L, Bridge A, Chan WM, Chavali G, Coudert E, Dimmer E, Estreicher A, Famiglietti L, Feuermann M, Gos A, Gruaz-Gumowski N, Hieta R, Hinz C, Hulo C, Huntley R, James J, Jungo F, Keller G, Laiho K, Legge D, Lemercier P, Lieberherr D, Magrane M, Martin MJ, Masson P, Mutowo-Muellenet P, O'Donovan C, Pedruzzi I, Pichler K, Poggioli D, Porras Millán P, Poux S, Rivoire C, Roechert B, Sawford T, Schneider M, Stutz A, Sundaram S, Tognolli M, Xenarios I, Foulgar R, Lomax J, Roncaglia P, Khodiyar VK, Lovering RC, Talmud PJ, Chibucos M, Giglio MG, Chang H, Hunter S, McAnulla C, Mitchell A, Sangrador A, Stephan R, Harris MA, Oliver SG, Rutherford K, Wood V, Bahler J, Lock A, Kersey PJ, McDowall DM, Staines DM, Dwinell M, Shimoyama M, Laulederkind S, Hayman T, Wang S, Petri V, Lowry T, D'Eustachio P, Matthews L, Balakrishnan R, Binkley G, Cherry JM, Costanzo MC, Dwight SS, Engel SR, Fisk DG, Hitz BC, Hong EL, Karra K, Miyasato SR, Nash RS, Park J, Skrzypek MS, Weng S, Wong ED, Berardini TZ, Huala E, Mi H, Thomas PD, Chan J, Kishore R, Sternberg P, Van Auken K, Howe D, Westerfield M. 2013. Gene Ontology annotations and resources. *Nucleic Acids Research* 41(Database issue):D530–D535 DOI 10.1093/nar/gks1050.
- Hoffmann R, Valencia A. 2005. Implementing the iHOP concept for navigation of biomedical literature. *Bioinformatics* 21:ii252–ii258 DOI 10.1093/bioinformatics/bti1142.

- Huang M, Ding S, Wang H, Zhu Z. 2008. Mining physical protein–protein interactions from the literature. *Genome Biology* 9(Suppl 2):S12 DOI 10.1186/gb-2008-9-s2-s12.
- Kano Y, Bjorne J, Ginter F, Salakoski T, Buyko E, Hahn U, Cohen KB, Verspoor K, Roeder C, Hunter LE, Kikicoglu H, Bergler S, Van eghem S, Van Parys T, Van de Peer Y, Miwa M, Ananiadou S, Neves M, Pascual-Montano A, zgr A, Radev DR, Riedel S, Saetre R, Chun H, Kim J-D, Pyysalo S, Ohta T, Tsujii J. 2011. U-Compare bio-event meta-service: compatible BioNLP event extraction services. *BMC Bioinformatics* 12:481 DOI 10.1186/1471-2105-12-481.
- Kim JD, Nguyen N, Wang Y, Tsujii J, Takaqi T, Yonezawa A. 2012. The genia event and protein coreference tasks of the BioNLP shared task 2011. *BMC Bioinformatics* 13(Suppl 11):S1 DOI 10.1186/1471-2105-13-S11-S1.
- Kim J-D, Pyysalo S, Ohta T, Bossy R, Nguyen N, Tsujii J. 2011. *Overview of BioNLP shared task*. Portland: Association for Computational Linguistics.
- Kotera M, Hirakawa M, Tokitatsu T, Goto S, Kanehisa M. 2012. The KEGG databases and tools facilitating omics analysis: latest developments involving human diseases and pharmaceuticals. *Methods in Molecular Biology* 802:19–39.
- Krallinger M, Leitner F, Rabal O, Vazquez M, Oyarzabal J, Valencia A. 2013. *Proceedings of the fourth biocreative challenge evaluation workshop*. Washington: BioCreative Challenge.
- Mao Y, Van Auken K, Li D, Arighi CN, Lu Z. 2013. The gene ontology task at biocreative IV. *Proceedings of the Fourth Biocreative Challenge Evaluation Workshop* vol. 1. 119–127.
- Markowitz F, Spang R. 2007. Inferring cellular networks—a review. *BMC Bioinformatics* 8(Suppl 6):S5 DOI 10.1186/1471-2105-8-S6-S5.
- Malouf R. 2002. A comparison of algorithms for maximum entropy parameter estimation. *Proceedings of the sixth conference on natural language learning*. 49–55.
- Mi H, Lazareva-Ulitsky B, Loo R, Kejariwal A, Vandergriff J, Rabkin S, Guo N, Muruganujan A, Doremiex O, Campbell MJ, Kitano H, Thomas PD. 2005. The PANTHER database of protein families, subfamilies, functions and pathways. *Nucleic Acids Research* 33(Database issue):D284–D288 DOI 10.1093/nar/gki078.
- NCBI Fact Sheet. Accessed 01/07/2013, Available at: <http://www.nlm.nih.gov/pubs/factsheets/medline.html>.
- Peng FC, D Schuurmans. 2003. Combining naive Bayes and n-gram language models for text classification. *Advances in Information Retrieval* 2633:335–350.
- Plake C, Royer L, Winnenburger R, Hakenberg J, Schroeder M. 2009. GoGene: gene annotation in the fast lane. *Nucleic Acids Research* 37(Web Server issue):W300–W304 DOI 10.1093/nar/gkp429.
- Pyysalo S, Ohta T, Rak R, Sullivan S, Mao C, Wang C, Sobral B, Tsujii J, Ananiadou S. 2012. Overview of the ID, EPI and REL tasks of BioNLP Shared Task 2011. *BMC Bioinformatics* 13(Suppl 11):S2 DOI 10.1186/1471-2105-13-S11-S2.
- Szklarczyk D, Franceschini A, Kuhn M, Simonovic M, Roth A, Minquez P, Stark M, Muller J, Bork P, Jensen LJ, von Mering C. 2011. The STRING database in 2011: functional interaction networks of proteins, globally integrated and scored. *Nucleic Acids Research* 39(Database issue):D561–D568 DOI 10.1093/nar/gkq973.
- Usi  A, Karathia H, Teixeira I, Valls J, Faus X, Alves R, Solsona S. 2011. Biblio-MetReS: a bibliometric network reconstruction application and server. *BMC Bioinformatics* 12:387 DOI 10.1186/1471-2105-12-387.

- Vazquez M, Krallinger M, Leitner F, Valencia A. 2011.** Text mining for drugs and chemical compounds: methods, tools and applications. *Molecular Informatics* **30(6-7)**:506–519  
[DOI 10.1002/minf.201100005](https://doi.org/10.1002/minf.201100005).
- Wu CH, Arighi CN, Cohen KB, Hirshman L, Krallinger M, Lu Z, Mattingly C, Valencia A, Wiegers TC, Wilbur WJ. 2012.** BioCreative-2012 virtual issue. *Database* **2012**:bas049  
[DOI 10.1093/database/bas049](https://doi.org/10.1093/database/bas049).
- Yang Z, H Lin, Li Y. 2008.** Exploiting the performance of dictionary-based bio-entity name recognition in biomedical literature. *Computational Biology and Chemistry* **32**:287–291  
[DOI 10.1016/j.compbiolchem.2008.03.008](https://doi.org/10.1016/j.compbiolchem.2008.03.008).