



HAL
open science

GO2PUB: Querying PubMed with Semantic Expansion of Gene Ontology Terms

Charles Bettembourg, Christian Diot, Anita Burgun, Olivier Dameron

► **To cite this version:**

Charles Bettembourg, Christian Diot, Anita Burgun, Olivier Dameron. GO2PUB: Querying PubMed with Semantic Expansion of Gene Ontology Terms. Journées Ouvertes en Biologie, Informatique et Mathématiques - JOBIM 2012, Jul 2012, Rennes, France. LIRMM UMR CNRS/UM2 5506, 2012, Journées Ouvertes de Biologie, Informatique et Mathématiques. hal-01210643

HAL Id: hal-01210643

<https://hal.science/hal-01210643>

Submitted on 3 Jun 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

GO2PUB: Querying PubMed with Semantic Expansion of Gene Ontology Terms

Charles BETTEMBOURG^{1,2}, Christian DIOT², Anita BURGUN¹ and Olivier DAMERON¹

¹ UMR936 INSERM, Université de Rennes 1, 2 av. Léon Bernard, F-35043 Rennes, France
{charles.bettembourg, anita.burgun, olivier.dameron}@univ-rennes1.fr

² UMR1348 INRA, Agrocampus Ouest, 65 Route de Saint-Brieuc, F-35042 Rennes, France
{charles.bettembourg, christian.diot}@rennes.inra.fr

Abstract *As PubMed grows, literature searches become more complex and time-consuming. Automated search tools with good precision and recall are thus necessary. We developed GO2PUB to address this demand. Our goal was to use the knowledge from Gene Ontology (GO), its annotations and following the true path rule for semantic expansion. GO2PUB enriches PubMed queries based on selected GO terms and keywords. It processes the result and displays the PMID, title, authors, abstract and bibliographic references of the articles. Gene names, symbols and synonyms that have been generated as extra keywords from the GO terms are also highlighted. Two experts manually assessed the relevance of GO2PUB, GoPubMed and PubMed on three queries about lipid metabolism. Experts agreement was high ($\kappa=0.83$). GO2PUB returned 76 % of the relevant articles, GoPubMed: 45 % and PubMed: 21 %. 39 % of the relevant articles were returned only by GO2PUB, and 18 % only by GoPubMed. We also generated 20 queries based on random GO terms with a granularity similar to those of the first three queries and compared the proportions of GO2PUB and GoPubMed results. These were respectively of 70 % and 38 %. The comparison of GO2PUB, based on semantic expansion, with GoPubMed, based on text-mining techniques, showed that both tools are complementary. GO2PUB is available at <http://go2pub.genouest.org>*

Keywords Gene ontology, Semantic expansion, Query enrichment, PubMed

1 Background

Due to the permanent growth of data, information retrieval becomes an increasingly difficult task. PubMed is the most comprehensive public database of biomedical literature. It comprises more than 21 million entries for biomedical literature¹. There are now 4 million more articles than 5 years ago². A well defined query is important for retrieving as many relevant articles as possible with as few irrelevant ones as possible when exploring or when keeping up with a domain. To reach this goal the development of automatic tools helping the users build such complex queries is needed.

PubMed supports MeSH-based query expansion [1]. Some additional literature search tools have been developed [2] and evaluated [3]. These correspond to three major approaches. The first approach, exemplified by tools like SLIM [4], is based on an intuitive interface to set some filters on PubMed queries for obtaining a better precision than with the basic PubMed querying system. A good proficiency with PubMed advanced search brings similar results. The second approach developed in SEGOPubMed uses a Latent Semantic Analysis (LSA) framework. It is based on a semantic similarity measure between the user query and PubMed abstracts [5]. The authors of SEGOPubMed state that the LSA approach outperforms the other approaches when using well-referenced keywords. Unfortunately, no implementation of SEGOPubMed is currently available. Moreover, this method requires that a corpus of well-referenced keywords be constituted and maintained before the search. Such a corpus is not available either. The third approach is based on query enrichment using controlled vocabularies and ontologies. These knowledge representations in which concepts are described both

1. www.ncbi.nlm.nih.gov/pubmed

2. www.nlm.nih.gov/bsd/licensee/baselinestats.html

by their meaning and their relations to each other are useful for information retrieval [6]. QuExT performs a concept-oriented query expansion to retrieve articles associated with a given list of genes from PubMed and to prioritize them [7]. GoPubMed [8] uses a text extraction algorithm to mine PubMed abstracts with GO terms [9]. However, GO terms are linked by relationships according to strict rules conveying the semantics. Among them, the true path rule requires that every child term inherits the meaning of all of its ancestor terms. GoPubMed does not follow this rule, resulting in a loss of flexibility and power [10].

We hypothesized that genes annotated by GO terms of interest in Gene Ontology Annotation database [11] can be used as additional keywords in gene-oriented PubMed queries. In GO2PUB, we propose a new approach that considers not only the genes annotated with a GO term of interest, but also those annotated by a descendant of this GO term, complying with the true path rule. GO2PUB's user inputs a list of GO terms, one or more species, and a list of keywords. GO2PUB retrieves all the genes annotated by at least one of the GO terms. It generates a PubMed query with the names, symbols and synonyms or aliases of these genes, the species and the keywords and processes PubMed result.

We performed a qualitative relevance study on our domain of expertise using three queries related to lipid metabolism. For each query, we compared PubMed, original GoPubMed and after having manually-generated the semantic expansion of the GO term and GO2PUB results. Two experts manually determined the relevance of all the articles. We computed the precision, relative recall and F-score of GO2PUB and GoPubMed. Then we performed a quantitative study on 20 randomly-generated queries to assess if the results can be generalized.

2 Resources and Methods

2.1 Resources

In the manuscript, we used the March 2011 version of GO for semantic expansion and of Gene Ontology Annotations database (<http://www.ebi.ac.uk/GOA>) for retrieving the genes annotated by the provided GO terms. These tables allowed to build queries about only 7 different species. Since June 2011, GO2PUB uses the Uniprot-GOA table instead of the species-specific tables, allowing to mine literature about over 2000 different species. We represented the overlap of the results using Venn diagrams generated by BioVenn [14].

2.2 Methods

2.2.1 GO2PUB query building. GO2PUB creates an expanded PubMed query. Figure Fig. 1 presents the process.

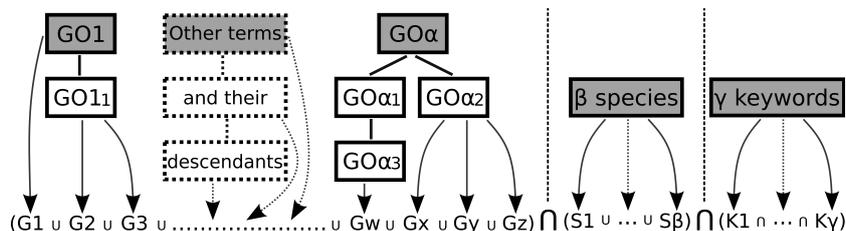


Figure 1. GO2PUB Query composition process using parameters provided by the user. Step 1: initial α GO terms (grey boxes) are enriched by their descendants. Step 2: genes (G_1 to G_z) annotated by the GO terms are retrieved. Step 3: query is composed using names, symbols and synonyms of the genes, β species (S) and γ MeSH or free keywords (K).

The first part of each query involves one or more GO terms. The user can enter either the name or the identifier of the GO terms. These terms are suggested when the user starts to fill the field. The exact GO term is suggested if the user provides one of its GO synonyms. For example, GO2PUB will search for “lipid biosynthetic process” if the user provides “lipogenesis”. When two or more GO terms are entered, GO2PUB makes the union of them (“OR” connector). Then, the user selects one or several species using a name (common or scientific names and their synonyms are allowed) or a NCBI taxon code. In this case, the user can choose to join them (using “OR”) or intersect them (using “AND”). Logical connectors “AND” and “OR” are set by

default to make the union of species and intersection of keywords, but this can be modified. Next, the user can enter additional MeSH terms to decrease the number of false positives. MeSH terms associated to the articles by PubMed are not all of same importance, some of them are “Major topic” (MAJR) classified. We can qualify each keyword as a simple MeSH term or a Major topic. Again, the user can specify the connector between keywords. At this point, the user has built a simple GO2PUB query. We called this query [BASICq].

The system supports 3 modifications for [BASICq] for studying if minor changes bring additional relevant results. The first modification lets the opportunity to ignore MAJR qualifiers and search all keywords in PubMed [MeSH] tag. As MAJR terms are also MeSH terms, articles associated to them will still be found. We called this query [MeSHq]. The second modification replaces “AND” connectors between keywords by “OR” connectors. However, as it can return many more results with a lot of noise, all keywords in this additional query are tagged with MAJR. Species that are normally searched in MeSH are also tagged with MAJR. We called this query [ORq]. The third modification ignores MeSH and MAJR keywords, and tags species with MAJR. This option must be used carefully because it can yield several hundreds of results if the topic of the search is too large. It is of interest only for very narrow topics if the user does not obtain enough results with the other queries. We called this query [NOKq]. Last, GO2PUB proposes 3 additional options. The first option allows to set a lower limit on the publication year. The second option proposes an exhaustive search of the official synonyms of gene names using Entrez gene³. As the authors sometimes use in their articles synonyms that are absent in the GOA database, this option allows to build more complete PubMed queries in order to obtain more relevant results. The third option toggles the display of the MeSH table associated with each article.

2.2.2 Query rewriting using semantic expansion. After semantic expansion genes annotated with GO terms and those annotated by descendants are selected. Figure Fig. 2 shows that the expansion yields five genes associated with the regulation of fatty acid metabolic process, instead of two if the true path rule is ignored.

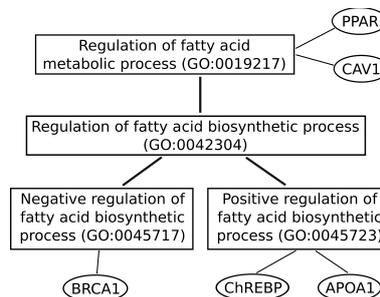


Figure 2. Keyword semantic enrichment. The two genes PPAR and CAV1 are directly annotated by the GO term “Regulation of fatty acid metabolic process” (GO:0019217). However, the true path rule implies that genes annotated by at least one descendant of the original term (BRCA1, ChREBP and APOA1) should also be considered.

GO2PUB retrieves all the genes annotated by each GO term, directly or indirectly through the true path rule. GO2PUB builds a query on the model “(*n* gene names, symbols or synonyms separated by OR) AND (*m* species) AND (*p* MeSH terms)”. The name, symbol and synonyms of each gene compose the first part of the query. They will be searched in title and abstract. Species and keywords make up the second part of the query. Finally, GO2PUB submits to PubMed a query composed of genes annotated directly or indirectly (name OR symbol OR Synonym), at least one species and some MeSH terms and free keywords. GO2PUB compiles the results and displays all citations numbered and sorted by date.

3 Results

3.1 Qualitative study

In order to evaluate GO2PUB relevance and to compare it with GoPubMed, we assessed three queries (Q1, Q2 and Q3) about biological processes related to lipid metabolism and including different GO terms, species

3. <http://www.ncbi.nlm.nih.gov/sites/entrez?db=gene>

and MeSH terms. As GoPubMed only considers the GO term(s) provided by the user and ignores the inheritance rules of GO, we expanded queries manually then submitted them to GoPubMed. Our GoPubMed queries were composed by the GO term(s) and all its descendants separated by “OR”, plus MeSH keywords. This ensured the closest comparison possible. We have also written PubMed queries as close to ours as possible.

3.1.1 Relevance criteria We analyzed the results of GO2PUB, GoPubMed and PubMed queries according to the following criteria. We considered that a relevant article has to describe at least one gene product occurring in the chosen metabolism for the selected species, in particular its role, its interactions, and how and when it is activated. For queries Q1, Q2 and Q3, the results obtained by the different tools were mixed for a blind selection by two reviewers. This ensured that they did not know the tool(s) that retrieved the articles. The final list of relevant articles is the union of the two reviewers’ lists.

3.1.2 Relevance measurement For each query and tool, we computed the precision, recall and F-score. Computing the real recall for each query is impossible because it requires to know all the really relevant articles available in Medline. However, it is possible that some of these articles were missed by all of the three tools. Thus, recall was defined as relative to all relevant articles obtained by at least one of the tools. Most of the relevant articles were found in the intersection of the two reviewers’ selections. Indeed, reviewers agreed on 31 relevant and 49 irrelevant articles while selecting separately 3 and 4 articles as relevant. We used Cohen’s kappa coefficient as a statistical measure of inter-rater agreement [12]. The value obtained was 0.83, which corresponds to an almost perfect agreement [13].

3.1.3 Query Q1: Lipogenesis in chicken liver In our first reviewed query in GO2PUB, we used “Lipid biosynthetic process” (GO:0008610) as GO term, “Gallus gallus” as species, “Liver” as Major Topic and “Lipid Metabolism” as MeSH keyword, and we considered the articles published in the last five years. We ran query Q1 on GO2PUB using the [BASICq], [MeSHq] and [ORq]. “Lipid biosynthetic process” has 243 descendants in GO. The mean number of edges to reach the root of the ontology from this term was 3.5. GoPubMed equivalent query was “lipid biosynthetic process”[go] AND Chickens[mesh] AND Liver[majr] AND “Lipid Metabolism”[mesh] AND last5years[time]. This is the “standard” query for GoPubMed. We have also considered the manually-expanded version of this query by adding the descendants of “lipid biosynthetic process” separated by “OR”. It should be noted that 47 of the 243 terms generated by the semantic expansion of “Lipid biosynthetic process” generated a GoPubMed error and had to be ignored. PubMed equivalent query was “Chickens liver lipogenesis”, which PubMed interpreted using MeSH terms combinations. Figure Fig. 3 presents Venn diagrams of Q1 results obtained with PubMed, GoPubMed (after manual expansion) and GO2PUB. Although queries as similar as possible were issued to the three tools, the resulting sets of articles had little overlap (a). Most of relevant articles (b) were yielded by GO2PUB. We notice that most of the articles retrieved by at least two tools (overlaps in a) were relevant (overlaps in b).

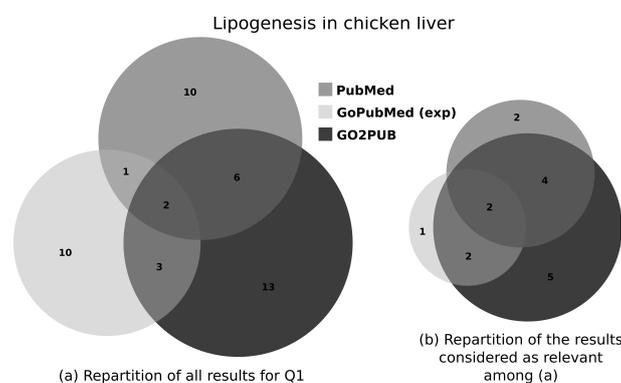


Figure 3. Comparison of the PubMed, GoPubMed and GO2PUB results for Q1. (a) displays the repartition and intersections of these results. (b) displays the repartition and intersections of the results considered as relevant.

Table Table 1 presents the precision, relative recall and F-score for each tool. GO2PUB had better precision and relative recall than GoPubMed and Pubmed. There was no difference between “standard” and “expanded” GoPubMed results.

	PubMed	GoPubMed (std)	GoPubMed (exp)	GO2PUB
(a) Number of results	19	16	16	24
(b) Relevant among (a)	8	5	5	13
Precision	0.421	0.313	0.313	0.542
Relative Recall	0.5	0.313	0.313	0.813
F-score	0.457	0.313	0.313	0.650

Table 1. Precision, relative recall and F-score values using PubMed, GoPubMed without (std) or with (exp) manual expansion and GO2PUB search tools for query Q1 about lipogenesis in chicken liver. Values are calculated from (a) and (b) lines using a total number of relevant results of 16.

3.1.4 Query Q2: Lipid transport in human blood In query Q2, we used “Lipid transport” (GO:0006869) as GO term, “Homo sapiens” as species, “Blood” as Major Topic and “Lipid Metabolism” as MeSH keyword, and we considered the articles published in the last five years. “Lipid transport” has 109 descendants in the GO graph. The mean number of edges to reach the root of the ontology from this term was 4.3. We ran equivalent queries on GoPubMed (“standard” and “expanded” versions) and PubMed. 46 of the 109 terms generated by the semantic expansion of “Lipid Transport” generated a GoPubMed error and had to be ignored. PubMed equivalent query yielded 45 articles, none of which was retrieved by GO2PUB nor GoPubMed while there was an overlap between GO2PUB and GoPubMed results. This discrepancy is probably due to the absence of [GO] search field tag in PubMed. Thus, comparing PubMed with GoPubMed or GO2PUB appeared meaningless. Consequently, PubMed results were not analyzed for this query.

Figure Fig. 4 presents the results obtained by GoPubMed (after manual expansion) and GO2PUB for Q2. As observed for query Q1, the majority of the results were tool-specific (a). Three of the four common articles were relevant (b). GO2PUB yielded the half of GoPubMed relevant results while having an important specific relevant results set.

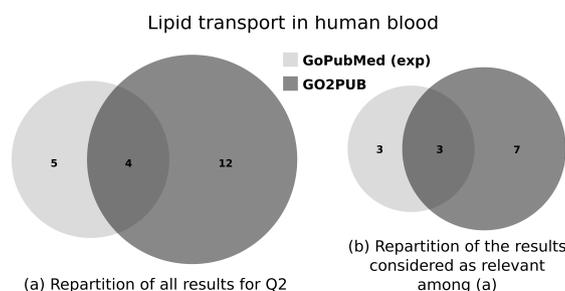


Figure 4. Comparison of Q2 PubMed, GoPubMed and GO2PUB results. (a) repartition and intersections of these results. (b) repartition and intersections of relevant results.

Table Table 2 gives Q2 precision, relative recall and F-score of GoPubMed and GO2PUB. GO2PUB has a slightly lower precision than GoPubMed (standard and after manual expansion) and better relative recall and F-score. For GoPubMed, there was no difference between “standard” and “expanded” results.

3.1.5 Query Q3: Regulation of lipase activity in human cell membrane In a third query, we used “Regulation of lipase activity” (GO:0060191) as GO term, “Homo sapiens” as species and “Cell Membrane” and “Lipid Metabolism” as MeSH keywords, and we considered the articles published in the last ten years. “Regulation of lipase activity” has 35 descendants in the GO graph. The mean number of edges to reach the root of the ontology from this term was 5.25. We ran equivalent queries on GoPubMed (“standard” and “expanded” versions) and PubMed. 16 of the 35 terms generated by the semantic expansion of “Regulation of lipase activity” generated a GoPubMed error and had to be ignored. PubMed yielded 23 articles for a query

	GoPubMed (std)	GoPubMed (exp)	GO2PUB
(a) Number of results	9	9	16
(b) Relevant among (a)	6	6	10
Precision	0.667	0.667	0.625
Relative Recall	0.462	0.462	0.769
F-score	0.546	0.546	0.690

Table 2. Values of precision, relative recall and F-score using GoPubMed without (std) or with (exp) manual expansion and GO2PUB search tools for query Q2 about lipid transport in human blood. Values are calculated from (a) and (b) lines using a total number of relevant results of 13.

using “regulation”, (“lipase” AND “activity”), “human” and (“cell” AND “membrane”) as keywords. Again, as none of these articles were found by GoPubMed or GO2PUB, PubMed results were not analyzed.

As shown in Fig. 5 A, a larger set of results was returned by GO2PUB compared to GoPubMed (24 and 8, respectively), but most of these results are irrelevant (Fig. 5 B).

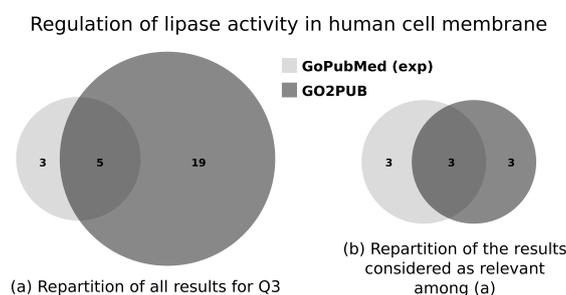


Figure 5. Comparison of the PubMed, GoPubMed and GO2PUB results for Q3. (a) repartition and intersections of results. (b) repartition and intersections of relevant results. GoPubMed data result from manual semantic expansion.

Table Table 3 gives Q3 precision, relative recall and F-score of GoPubMed and GO2PUB. GO2PUB has a relative recall equivalent to GoPubMed’s and lower precision and F-score. For GoPubMed, the “manually-expanded” results have a higher relative recall and F-score and a lower precision than the “standard” ones.

	GoPubMed (std)	GoPubMed (exp)	GO2PUB
(a) Number of results	6	8	24
(b) Relevant among (a)	5	6	6
Precision	0.833	0.75	0.25
Relative Recall	0.556	0.667	0.667
F-score	0.667	0.706	0.364

Table 3. Values of precision, relative recall and F-score using PubMed, GoPubMed without (GoPM std) or with (GoPM exp) manual expansion and GO2PUB search tools for query Q3 about regulation of lipase activity in human cell membrane. Values are calculated from (a) and (b) lines using a total number of relevant results of 9.

3.2 Quantitative study

In order to assess if previous results can be generalized, we performed a quantitative study on 20 randomly-generated queries. We built queries following the pattern: “a random GO term + a species (mouse) + a publication date limit (2011) + a keyword (the GO term name)”. To be coherent with the qualitative study, we randomly selected 20 GO terms among all Biological Process terms having a granularity similar to those of the three GO terms used in the qualitative study. We assumed that the granularity of a term depends on the mean length of its path to the root, and its number of descendants. Each GO term of the quantitative study had a mean path length to the root between 3.5 and 5.25 edges and had between 35 and 244 descendants. As we could not add a MeSH keyword in relation with the random GO term of each query, we simply added the name of this GO term. This keyword was added in the free field for GO2PUB and without [go] tag for GoPubMed. We submitted these queries to GO2PUB and to GoPubMed.

GO2PUB yielded 46 articles on average (min 6, max 189) compared to 21.33 on the qualitative study. They represented 70 % of the total number of articles (77 % in the qualitative study). GoPubMed yielded 25.1 articles on average (min 2, max 88) compared to 11.0 on the qualitative study. They represented 38 % of the total number of articles (40 % in the qualitative study). There were 5.75 articles on average (min 0, max 59) in the common set. They represented 9 % of the total (17 % in the qualitative study). The profile of these results is close to the qualitative study one.

4 Discussion

Our goal was to develop a tool that uses the knowledge from the Gene Ontology (GO) and its annotations for generating semantically-expanded gene-related PubMed queries. Indeed, there is no [GO] tag for a search in PubMed. We compared the results of four different queries sent to PubMed, to our tool GO2PUB and to GoPubMed, another tool that allows to query PubMed with GO terms.

The qualitative study showed that both GO2PUB and GoPubMed retrieved relevant articles ignored by PubMed. For Q1, 26 of the 35 articles (8 of 14 relevant) returned by either GO2PUB or GoPubMed were ignored by PubMed. Conversely, 9 of the 19 articles (6 out of 8 relevant) returned by PubMed were also returned by either GO2PUB or GoPubMed. For Q2 and Q3, the set of articles returned by PubMed was disjoint from both GO2PUB and GoPubMed results. We have not considered PubMed results for these queries because the scope of this article is to assess the relevance of Gene Ontology-based query enrichment using the true path rule, and not to assess PubMed.

Overall, GO2PUB had better performances than GoPubMed and PubMed, but both GoPubMed and GO2PUB yielded relevant articles ignored by the other. The differences varied among the queries. For Q1 and Q2, GO2PUB yielded most of the relevant articles and had therefore the highest relative recall value while its precision was slightly lower than that of GoPubMed. Consequently, GO2PUB had the best F-score. For Q3, GO2PUB yielded as many relevant articles as GoPubMed but had a higher proportion of noise. GO2PUB had a slightly better relative recall than GoPubMed, but its precision was much lower. Consequently, GoPubMed had the best F-score. We can also notice that for Q3, the query expansion on GoPubMed improved its performances with a better relative recall and F-score at the cost of a small loss of precision.

GO2PUB performs a semantic expansion of the GO terms of interest complying with the GO true path rule before retrieving the corresponding genes for enriching the query. During the development of GO2PUB, we have also run queries without this expansion. We obtained empty or very small sets of results.

Using the true path rule is useful. The more descendants a GO term has, the more relevant results GO2PUB yields. GO2PUB performances decreased from Q1 to Q3. For Q1, “lipid biosynthetic process” has 243 descendants and annotates 646 genes for human and 145 genes for chicken. For Q2, “lipid transport” has 109 descendants and annotates 253 genes for human and 63 genes for chicken. For Q3, “regulation of lipase activity” has 35 descendants and annotates 168 genes for human and 18 genes for chicken. The more descendants a GO term has, the more genes it is likely to annotate, which increases the gene-related enrichment importance. Moreover, Q1 concerned chicken, which is less annotated than human. On less annotated species, the annotations focus on the major genes. This explains why GO2PUB yields a high proportion of relevant articles.

GoPubMed does not follow the true path rule. We manually expanded GoPubMed queries and compared it to GO2PUB. The added value of semantic expansion was null for Q1 and Q2, and important for Q3 (+33%). So query expansion is a built-in functionality in GO2PUB, and would be a valuable extension for GoPubMed. In GoPubMed results, a “missing term” error occurred for 19% of the expanded set of GO terms for Q1, 42% for Q2 and 44% for Q3. We assume that the benefits of query expansion on GoPubMed might be higher by considering the articles related to these currently omitted GO terms.

Most of the results obtained by both tools are relevant (4 out of 4 for Q1, 3 out of 4 for Q2 and 3 out of 5 for Q3). So, the intersection of GoPubMed and GO2PUB results decreases noise. As each tool yields relevant articles that are ignored by the other, the union of their results also decreases silence.

In order to consolidate the qualitative study, we submitted 20 randomly-generated queries to GO2PUB and GoPubMed. These queries were more general than Q1, Q2 and Q3 because of the use of non-MeSH keywords. Each query contained a random GO term of a granularity similar to that of the terms used in the qualitative study. The proportion of articles returned by GO2PUB was 70 %, the one of GoPubMed was 38 % and the proportion of articles returned by both was 9 %. We assume that the difference of proportions between the qualitative and the quantitative studies can be attributed to Q1, Q2 and Q3 being more specific because of the use of MeSH keywords.

5 Conclusion

GO2PUB brings relevant results ignored by GoPubMed (9 GO2PUB' specific results for Q1, 7 for Q2 and 3 for Q3) even doing a manual query expansion for GoPubMed. Conversely GoPubMed text-mining approach finds relevant articles ignored by GO2PUB (1 GoPubMed' specific result for Q1, 3 for Q2 and 3 for Q3). This demonstrates GO2PUB relevance and its complementarity with GoPubMed for our domain of interest. The quantitative analysis shows that a similar profile of results is obtained using random queries, especially when using keywords for narrowing the queries. This suggests that the results of the qualitative study can be generalized.

Acknowledgements

Biogenouest for computer support and hosting. CB was supported by a fellowship from the French ministry of research.

References

- [1] Z. Lu, W. Kim, W.J. Wilbur, Evaluation of Query Expansion Using MeSH in PubMed. *Inf Retr Boston*, 12(1):69-80, 2009
- [2] Z. Lu, PubMed and beyond: a survey of web tools for searching biomedical literature. *Database (Oxford)*, :baq036, 2011
- [3] A. Bajpai, S. Davuluri, H. Haridas, G. Kasliwal, H. Deepti, K.S. Sreelakshmi, D.S. Chandrashekar, P. Bora, M. Farouk, N. Chitturi, V. Samudiyata, K.P. ArunNehru, K. Acharya, In search of the right literature search engine(s). *Nature Precedings*, doi:10.1038/npre.2011.2101.3, 2011
- [4] M. Muin, P. Fontelo, F. Liu, M. Ackerman, SLIM: an alternative Web interface for MEDLINE/PubMed searches - a preliminary study. *BMC Med Inform Decis Mak.*, 5:37, 2005
- [5] B.C. Vanteru, J.S. Shaik, M. Yeasin, Semantically linking and browsing PubMed abstracts with gene ontology. *BMC Genomics*, 9(Suppl 1):S10, 2008
- [6] J.B. Bard, S.Y. Rhee, Ontologies in biology: design, applications and future challenges. *Nat. Rev. Genet.*, 5:213-222, 2004
- [7] S. Matos, J.P. Arrais, J. Maia-Rodrigues, J.L. Oliveira, Concept-based query expansion for retrieving gene related publications from MEDLINE. *BMC Bioinformatics*, 11:212,2010
- [8] A. Doms, M. Schroeder, GoPubMed: exploring PubMed with the Gene Ontology. *Nucleic Acids Res.*, 33:W783-W786, 2005
- [9] M.A. Harris, Consortium GO, The Gene Ontology (GO) database and informatics resource. *Nucleic Acids Res.*, 1:D258-D261, 2004
- [10] S.Y. Rhee, V. Wood, K. Dolinski, S. Draghici, Use and misuse of the gene ontology annotations. *Nat. Rev. Genet.*, 9:509-515, 2008
- [11] E. Camon, M. Magrane, D. Barrell, V. Lee, E. Dimmer, J. Maslen, D. Binns, N. Harte, R. Lopez, The Gene Ontology Annotation (GOA) Database: sharing knowledge in Uniprot with Gene Ontology. *Nucleic Acids Res.*, 32:D262-D266, 2004
- [12] J. Cohen, A coefficient of agreement for nominal scales. *Educ Psychol Meas.*, 20(1):37-46, 1960
- [13] J.R. Landis, G.G. Koch, The measurement of observer agreement for categorical data. *Biometrics*, 33(1):159-74, 1977
- [14] T. Hulsen, J. de Vlieg, W. Alkema, BioVenn - a web application for the comparison and visualization of biological lists using area-proportional Venn diagrams. *BMC Genomics*, 9(1):488, 2008