Appendix: Content Summaries of Selected Best Papers for the IMIA Yearbook 2021, Section Knowledge Representation and Management

Le DH

UFO: A tool for unifying biomedical ontology-based semantic similarity calculation, enrichment analysis and visualization

PLoS One 2020;15(7):e0235670

In this article, Le presents a unified tool to support semantic similarity-based research called UFO. Ontology-based similarity has become a routine approach in many applications (ontology curation, enrichment, decision support, gene association studies…). Similarity measurements in ontologies can be performed using many metrics and methods, and current solutions are scattered over different tools and platforms.

UFO is implemented as an app for the Cytoscape platform (open-source software for visualizing complex networks and integrating data in molecular and systems biology, genomics, and proteomics) and supports the OBO format. This article (and supplementary material) describes all UFO features and refers to case studies of relevant implementations: human disease phenotype similarity based on Human Phenotype Ontology (HPO), prediction of disease-associated genes and protein complexes based on gene and protein complex similarity networks using Gene ontology, prediction of disease-associated genes and long non-coding RNAs based on disease similarity network using HPO and Disease ontology, and enrichment analysis with HPO.

The main detailed functions are similarity calculation, enrichment analysis, and visualization. The similarity matrices can be calculated between terms (with 11 metrics, node and/or edge-based), between annotated entities (pairwise or groupwise) and between two sets of entities. Statistical tests (binomial or Fischer's exact) can be applied to a set of entities to search for additional salient terms to enrich the set. Graph visualization facilitates the understanding of the relationships among selected terms, their ancestors (e.g., shared ancestors) and descendants, or among entities (similarity networks).

This tool brings to the KRM community a unified solution for semantic similarity research. Now, only OBO format ontologies are supported (other ontology formats will need conversion) and the tool is tailored for a molecular and systems biology platform. However, this solution supports any application domain.

Robinson PN, Ravanmehr V, Jacobsen JOB, Danis D, Zhang XA, Carmody LC, Gargano MA, Thaxton CL, Core UNCB, Karlebach G, Reese J, Holtgrewe M, Kohler S, McMurry JA, Haendel MA, Smedley D

Interpretable clinical genomics with a likelihood ratio paradigm

Am J Hum Genet 2020;107(3):403-17

In this paper, Robinson *et al.* address the phenotype-driven prioritization of variants with a metric providing robust estimates of the strength of the predictions of candidate genes or diseases, beyond the usual placement in a ranked list.

They present a novel algorithm, the LIkelihood Ratio Interpretation of Clinical AbnormaLities (LIRICAL), that calculates the likelihood ratio (LHR) of each observed or excluded phenotypic abnormality. For each candidate diagnosis, LIRICAL calculates the extent to which each phenotypic abnormality (and if available genotype) is consistent with the diagnosis. Phenotypic abnormalities are represented by Human Phenotype Ontology (HPO) terms and the LHR calculations are derived from the subsumption hierarchies in the HPO. In the methods section, the algorithm is entirely described.

This work illustrates how the structure of a knowledge representation (HPO) can contribute to a bioinformatics workflow. The performances of LIRICAL are demonstrated to be state-of-the-art, on simulated data from 384 published case reports and data from 116 solved cases from the 100,000 Genomes Project.

LIRICAL is available for academic use for free, and source code can be downloaded on a GitHub repository.

Slater LT, Gkoutos GV, Hoehndorf R

Towards semantic interoperability: finding and repairing hidden contradictions in biomedical ontologies

BMC Med Inform Decis Mak 2020;20(Suppl 10):311

In this paper, the authors present a method to identify hidden unsatisfiabilities in an ontology that arise when combined with other ontologies. They identified a large set of inconsistencies across a broad range of biomedical ontologies.

One way to combine ontologies is to use the MIREOT (Minimum Information to Reference an External Ontology Term) guidelines which were originally developed to support inclusion of classes from biomedical ontologies. MIREOT has become a standard for term re-use and inclusion throughout the biomedical ontology community. The authors advocate that, while this method allows ontologies to reuse classes in a scalable and efficient manner, the inclusion of external classes without the context of the external ontology's axioms means that contradictions may arise. These contradictions cannot be detected using an automated reasoner that evaluates only the target ontology.

The authors use automated reasoning to determine whether unsatisfiable classes are present. In addition, they designed a novel algorithm that suggest justifications for contradictions across large and complex ontologies. Their experiments identify contradictions that lead to unsatisfiable classes in the OBO ontologies and highlight the axioms that can be removed to solve most cases of unsatisfiability.

Such a work is important since researchers often import pieces of ontologies without considering the associated axioms. It is also important because it shows the challenge of maintaining a coherent group of ontologies in a large repository like the OBO ontologies.

Zheng F, Shi J, Yang Y, Zheng WJ, Cui L

A transformation-based method for

auditing the IS-A hierarchy of biomedical terminologies in the Unified Medical Language System

J Am Med Inform Assoc 2020;27(10):1568-75

In this paper, Zheng *et al.* use the rich knowledge provided by the Unified Medical Language System (UMLS) for auditing and improving the quality of its source terminologies. Given a concept name in the UMLS, they first identify its base and secondary noun chunks. For each identified noun chunk, they generate replacement candidates that are more general than the noun chunk. Then, they replace the noun chunks with their replacement candidates to generate new potential concept names that may serve as supertypes of the original concept. If a newly generated name is an existing concept name in the same source terminology as the original concept, then a potentially missing IS-A relation between the original and the new concept is identified.

This method gives very good results during the tests: a total of 39,359 potentially missing IS-A relations were detected in 13 source terminologies. Domain experts evaluated a random sample of 200 potentially missing IS-A relations identified in SNOMED CT and 100 in Gene ontology. A total of 173 of 200, and 63 of 100 potentially missing IS-A relations were confirmed by domain experts, indicating that the method achieved a precision of 86.5% and 63% for SNOMED CT and for Gene ontology, respectively.

This method is very suitable for large ontologies that are assemblages of thoughts and teams over time and for which it is difficult to audit the whole resource. It would be interesting to test this method on smaller ontologies.