# Evaluating Scientific Hypotheses Using the SPARQL Inferencing Notation

Alison Callahan[1] and Michel Dumontier[1,2,3,*]

[1] Department of Biology, [2] School of Computer Science, [3] Institute of Biochemistry,
Carleton University, Ottawa, Canada
acallaha@connect.carleton.ca,
michel_dumontier@carleton.ca

**Abstract.** Evaluating a hypothesis and its claims against experimental data is an essential scientific activity. However, this task is increasingly challenging given the ever growing volume of publications and data sets. Towards addressing this challenge, we previously developed HyQue, a system for hypothesis formulation and evaluation. HyQue uses domain-specific rulesets to evaluate hypotheses based on well understood scientific principles. However, because scientists may apply differing scientific premises when exploring a hypothesis, flexibility is required in both crafting and executing rulesets to evaluate hypotheses. Here, we report on an extension of HyQue that incorporates rules specified using the SPARQL Inferencing Notation (SPIN). Hypotheses, background knowledge, queries, results and now rulesets are represented and executed using Semantic Web technologies, enabling users to explicitly trace a hypothesis to its evaluation as Linked Data, including the data and rules used by HyQue. We demonstrate the use of HyQue to evaluate hypotheses concerning the yeast galactosegene system.

**Keywords:** hypothesis evaluation, semantic web, linked data, SPARQL.

## 1 Introduction

Developing and evaluating hypotheses in the context of experimental research results is an essential activity for the life scientist, but one which is increasingly difficult to carry out manually given the ever growing volume of publications and data sets[1]. Indeed, biologists perceive that the predominant challenge in research is to "locate, integrate and access" the vast amounts of biological data resulting from small- and large-scale experiments[2]. Life sciences resources for the Semantic Web, such as Bio2RDF[3] and the growing number of bio-ontologies offer the potential to develop systems that consume these resources and computationally reason over the knowledge they contain to infer new facts[4-6]and answer complex questions[7].

With the diversity of research claims that exist in such large resources, there is also the potential for statements to contradict one another. Formally exploring the out-

* Corresponding author.

E. Simperl et al. (Eds.): ESWC 2012, LNCS 7295, pp. 647–658, 2012.

comes of relying on different sets of research claims to assess a hypothesis is necessary to not only confer confidence in the hypothesis evaluation methodology (whether manual or automatic), but also to provide evidence for the likelihood of one interpretation of results compared to another. Previous research efforts that have aimed at formally evaluating scientific data in the context of hypotheses include HYPGENE[8, 9], HinCyc[10], GenePath[11] and Adam the Robot Scientist[12, 13]. Each of these projects use application-specific representations for data and the rules used to assess this data, making their extension to new domains, as well as their comparison and performance evaluation difficult.

Towards addressing the challenge of integrating experimental knowledge with biological hypotheses, we previously developed HyQue[14, 15]. HyQue uses Semantic Web standard languages (RDF/OWL) to represent hypotheses and data, SPARQL queries to retrieve data, and domain-specific rulesets to evaluate hypotheses against this data. While HyQue uses rulesets based on well understood scientific principles[16, 17], finer grained evaluations would require the exclusion or inclusion of additional rules. Problematically, HyQue's domain-specific evaluation rules were hard-coded, which made it implausible for users to construct custom rule sets for hypothesis evaluation.

In this paper, we describe an extension of HyQue that uses evaluation rules specified using the SPARQL Inferencing Notation (SPIN) in place of hardcoded rules. SPIN is a W3C member submission[1] rule language whose scope and expressivity are defined by SPARQL. Thus, SPIN rules are SPARQL queries which can not only be used to assert new facts, but also used to infer OWL class membership for non-hierarchical class membership axioms[2]. Moreover, SPIN rules can be serialized into RDF, and hence can become part of a system that maintains provenance concerning calculations and inferences.

In this new version of HyQue, hypotheses, background knowledge, queries, results and now evaluation rulesets are represented and executed using Semantic Web technologies. Domain specific rules for evaluating experimental data in the context of a hypothesis are now maintained independently of the system rules that are used to calculate overall hypothesis evaluation scores. We demonstrate these features by evaluating hypotheses about the galactose gene system in yeast[16]. HyQue enables users to explicitly trace a hypothesis to its evaluation, including the data and rules used. In addition to making the hypothesis evaluation methodology transparent and reproducible (essential qualities for good e-science), this allows scientists to discover experimental data that support a given hypothesis as well as explore new and potentially uncharacterized links between multiple research outcomes. A unique strength of HyQue is that its design is not dependent upon a specific biological domain, and the assumptions encoded in its hypothesis evaluation rules are changeable and maintained separately from the evaluation system. As our understanding of biological systems evolves and improves through research, the way HyQue evaluates hypotheses, as well as the facts and data it uses, can evolve as well.

[1] http://www.w3.org/Submission/2011/SUBM-spin-overview-20110222/
[2] http://www.w3.org/Submission/2011/SUBM-spin-modeling-20110222/

# 2 Methods

## 2.1 Overview

HyQue evaluates hypotheses (and assigns an evaluation score) by executing SPIN rules over the pertinent knowledge extracted from a HyQue Knowledge Base (HKB). A hypothesis is formulated as a logical expression in which elements of the hypothesis correspond to biological entities of interest. HyQue maps the hypothesis, expressed using terminology from the HyQue ontology[3], to the relevant SPIN rules, which execute SPARQL queries to retrieve data from the HKB. Finally, HyQue executes additional SPIN rules over the extracted data to obtain a quantitative measure of hypothesis support. Figure 1 provides a graphical overview of HyQue.

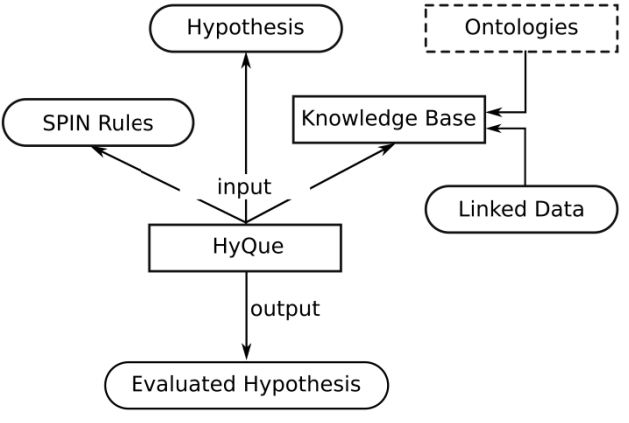


**Fig. 1.** HyQue uses SPIN rules to evaluate a hypothesis over RDF linked data and OWL ontologies. The dashed rectangle represents OWL ontologies. Rounded rectangles are RDF resources.

## 2.2 HyQue Hypothesis Model

A HyQue *hypothesis* may be composed of one or more *propositions* that *specify events* related to each other by AND/OR operators. Events must have an agent (an entity executing an action) and a target (the object of the action), and can optionally have a physical location, a physical operator (*e.g.* 'binding'), a logical operator (*e.g.* 'repression' or 'activation') and a perturbation context (in the case of genes and proteins). HyQue maps these events to SPARQL queries through a SPIN rule, and subsequently executes them over the HyQue Knowledge Base. HyQue currently supports the following kinds of events [14, 15]:

[3] The HyQue ontology, linked data, and SPIN rules are available at the project website: http://hyque.semanticscience.org

1. protein-protein binding
2. protein-nucleic acid binding
3. molecular activation
4. molecular inhibition
5. gene induction
6. gene repression
7. transport

### 2.3 HyQue Knowledge Base (HKB)

A HyQue Knowledge Base (HKB) consists of RDF data, RDFS-based class hierarchies and/or OWL ontologies. For demonstration purposes, our HKB consists of an RDF version of the galactose (GAL) gene network in yeast [17], an extended version of the Bio2RDF compatible yOWL knowledge base [7, 15] and the following bio-ontologies (for the listed entities):

- **Gene Ontology (GO):** cellular components, events (*e.g.* 'nucleus', 'positive regulation of gene expression')
- **Evidence Codes Ontology (ECO):** the type of evidence supporting an event (*e.g.* 'electronic annotation', 'direct assay')
- **Sequence Ontology (SO):** event participants (*e.g.* 'gene')
- **Chemical Entities of Biological Interest (CHEBI) Ontology:** event participants (*e.g.* 'protein', 'galactose')

All Linked Data (encoded using RDF) and ontologies (encoded using OWL) that comprise the HKB are available at the project website.

### 2.4 The HyQue Scoring System

HyQue uses rules to calculate a numerical score for a hypothesis based on the degree of support the hypothesis has from statements in the HKB. HyQue first attempts to identify statements about experimentally verified events in the HKB that have a high degree of matching to a hypothesized event, and then assesses these statements using domain specific rules to assign a score to the hypothesized event. If there is a statement about an experimentally reported GAL gene/protein interaction in the HKB that exactly matches a hypothesized event, then that event will be assigned a maximum score when it is evaluated by HyQue. In contrast, if a hypothesized event describes an interaction between a protein A and a protein B but there is a statement in the HKB asserting that protein A does *not* interact with protein B, then the hypothesis will be assigned a low score based on the negation of the hypothesized event by experimental data. Different HyQue rules add or subtract different numerical values based on whether the relevant experimental data has properties that provide support for a hypothesized event. For instance, if an event is hypothesized to occur in a specific cellular compartment *e.g.* nucleus, but the HKB only contains a statement that such an event takes place in a different cellular component *e.g.* cytoplasm, then a rule could be formulated

such that the hypothesis, while not directly supported by experimental evidence, will be penalized less than if the event had been asserted to not take place at all.

Based on such scoring rules, each event type has a maximum possible score. When a hypothesized event is evaluated by HyQue, it is assigned a normalized score calculated by the sum of the output of the relevant rule(s) divided by the maximum possible score. In this way, if an event has full experimental support, it will have an overall score of 1, while if only some properties of the hypothesized event are supported by statements in the HKB is will have a score between 0 and 1.

Overall proposition and hypothesis scores are calculated by additional rules based on the operators that relate events. If a proposition specifies 'event A' OR 'event B' OR 'event C' then the maximum event score will be assigned as the proposition score, while if the 'AND' operator was used, the mean event score will be assigned as the proposition score. Using the mean reflects the relative contribution of each event score while still maintaining a normalized value between 0 and 1. Similar rules are used to calculate an overall hypothesis score based on proposition scores.

HyQue uses SPIN to execute rules that reflect this scoring system.

### 2.5 HyQue SPIN Rules

HyQue uses two types of rules to evaluate hypotheses: *domain specific rules* that depend on the subject of the hypothesis (in this case, gene regulation) and *system rules* that define how to combine the output of domain specific rules in order to determine an overall hypothesis evaluation score. These rules are defined separately using SPIN and can be changed independently of each other.

HyQue system rules describe how to calculate event, proposition and overall hypothesis scores based on the structure and content of the hypothesis. For example, the following rule (modified with single quoted labels for illustrative purposes) generates four statements that assert the relationship between a HyQue hypothesis (any instance of the class hyque:HYPOTHESIS_0000000) and its evaluation.

```
CONSTRUCT {
   ?this 'has attribute' ?hypothesisEval .
   ?hypothesisEval a 'evaluation'.
   ?hypothesisEval 'obtained from' ?propositionEval .
   ?hypothesisEval 'has value ?hypothesisEvalScore .
} WHERE {
   ?this 'has component part' ?proposition .
   ?proposition 'has attribute' ?propositionEval .
   BIND(:calculateHypothesisScore(?this) AS ?hypothesisEvalScore) .
      BIND(IRI(fn:concat(afn:namespace(?this), afn:localname(?this),"_",
      "evaluation")) AS ?hypothesisEval) .
}
```

This SPIN rule states that a HyQue hypothesis (hyque:HYPOTHESIS_0000000) will be related to a new attribute of type 'evaluation' (hyque:HYPOTHESIS_0000005) by the 'has attribute' (hyque:HYPOTHESIS_0000008) object property. The numeric value of this evaluation is specified using the 'has value' (hyque:HYPOTHESIS_0000013) datatype property. Since the evaluation of the hypothesis comes from evaluating the propositional parts, these are related with the 'is obtained from' (hyque:HYPOTHESIS_0000007) object

property. The SPARQL variable '?this' has a special meaning for SPIN rules, and refers to any instance of the class the rule is linked to. SPIN rules are linked to classes in the HyQue ontology using the spin:rule predicate.

This hypothesis rule uses another rule, `calculateHypothesisScore`, to calculate the hypothesis score, and the output of executing this rule is bound to the variable `?hypothesisEvalScore`. Note that the hypothesis rule is constrained to a HyQue hypothesis that 'has component part' (hyque:HYPOTHESIS_0000010) some 'proposition' (hyque:HYPOTHESIS_0000001) that 'has attribute' a proposition evaluation. In this way HyQue rules are chained together – when one rule is executed, all the rules it depends on are executed until no new statements are created. In this case, because a hypothesis evaluation score requires a proposition evaluation score, when the hypothesis evaluation rule is executed, the HyQue SPIN rule for calculating a proposition score is executed as well. Each proposition evaluation is asserted to be 'obtained from' the event evaluations corresponding to the event(s) specified by (hyque:HYPOTHESIS_0000012) the proposition. Each event evaluation is also asserted to be 'obtained from' the scores determined for each event property (the agent, target, location *etc.*) and the statements in the HKB the scores are based on.

Domain specific rules for HyQue pertain to the domain of interest. An example of a domain specific rule is `calculateActivateEventScore` corresponding to the following SPARQL query:

```
SELECT ?activateEventScore
WHERE {
   BIND (:calculateActivateAgentTypeScore(?arg1)
     AS ?agentTypeScore) .
   BIND (:calculateActivateTargetTypeScore(?arg1)
     AS ?targetTypeScore) .
   BIND (:calculateActivateLogicalOperatorScore(?arg1)
     AS ?logicalOperatorScore) .
   BIND (:penalizeNegation(?arg1) AS ?negationScore) .
   BIND (3 AS ?maxScore) .
   BIND (((((?agentTypeScore + ?targetTypeScore) +
        ?logicalOperatorScore) + ?negationScore) /
        ?maxScore) AS ?activateEventScore) .
}
```

In this rule, a numeric score (`?activateEventScore`) is calculated from the sum of a set of outputs from other sub-rules divided by the maximum score possible (in this case, 3). This rule uses a special variable `?arg1`, which corresponds to any entities linked using the SPIN sp:arg1 predicate. This special variable is selected by specifying a spin:constraint on the rule, which states that any variable passed to the rule when it is called can be referred to within the rule to by '?arg1'. For example, if the rule were called by including `calculateActivateEventScore(?data)` in a SPARQL query WHERE statement, `?data` will be the variable referenced by `?arg1` in the rule definition.

The sub-rule `calculateActivateLogicalOperatorScore` determines a score for the type of logical operator specified in a HyQue hypothesis based on domain specific knowledge about the GAL gene network. This rule corresponds to the following SPARQL query:

```
SELECT ?score
WHERE {
   ?arg1 'has logical operator' ?logical_operator .
   BIND (IF((?logical_operator = 'positive regulation of molecular
        function'), 1, -1) AS ?score) .
}
```

Thus, if the logical operator specified in a hypothesis event is of type 'positive regulation of molecular function' (GO:0044093) the rule will return 1, and otherwise the rule will return -1.The `calculateActivateEventScore` rule is composed of several sub-rules of this format. HyQue uses similar rules for each of the seven event types listed in section 2.2 to evaluate hypotheses.

SPIN rules were composed using the free edition of TopBraid Composer 3.5. HyQue executes SPIN rules using the open source SPIN API 1.2.0 and Jena 2.6.4.

### 2.6 Executing HyQue SPIN Rules over the HKB

To execute the HyQue SPIN rules over an input hypothesis using data from the HKB, a Java program was written with the open source SPIN API (version 1.2.0) and the Jena API (version 2.6.4). Users can submit a hypothesis to the program via a servlet available at http://hyque.semantiscience.org. The servlet returns the RDF-based hypothesis evaluation.

## 3 Results

HyQue currently uses a total of 63 SPIN rules to evaluate hypotheses. 18 of these are system rules, and the remaining 45 are domain specific rules that calculate evaluation scores based on well understood principles of the GAL gene network in yeast as described in section 2.5. These rules have been used to evaluate 5 representative hypotheses about the GAL domain, one of which is presented in detail in section 3.1.

### 3.1 Evaluating a Hypothesis about GAL Gene Induction and Protein Inhibition

The following is a natural language description of a hypothesis about the GAL gene network that has been evaluated by HyQue. Individual events are indicated by the letter 'e', followed by a number to uniquely identify them. Events are related by the AND operator in this hypothesis, while the two sets of events (typed as propositions in the HyQue hypothesis ontology) are related by the OR operator.

| | |
|---|---|
| (Gal4p induces the expression of GAL1 | *e1* |
| *AND* Gal3p induces the expression of GAL2 | *e2* |
| *AND* Gal4p induces the expression of GAL7) | *e3* |
| *OR* | |
| (Gal4p induces the expression of GAL7 | *e4* |
| *AND* Gal80p induces the expression of GAL7 | *e5* |
| *AND* Gal80p does not inhibit the activity of Gal4p when GAL3 is over-expressed) | *e6* |

Two domain specific SPIN rules were executed to evaluate this hypothesis: `calculateInduceEventScore` for *e1-e5* and `calculateInhibitEventScore` for *e6,* in conjunction with system rules to calculate overall proposition and hypothesis scores based on the event scores.

By identifying and evaluating statements in the HKB that experimentally support *e1*, the `calculateInduceEventScore` rule assigns *e1* a score of 4 out of a maximum score of 5 (see Table 1). This corresponds to a normalized score of 0.8. Similarly, events 2-5 also receive a score of 0.8. The `calculateInhibitEventScore` rule assigns event 6 a score of 1 based on comparable scoring rules. Therefore, the proposition specifying *e4, e5* and *e6* receives a higher score (0.87 – the mean of the individual event scores) than the proposition specifying *e1, e2* and *e3* (with a mean score of 0.8). Because the two propositions were related by the OR operator, the hypothesis is assigned an overall score that is the maximum of the two proposition scores, in this case, a value of 0.87.

**Table 1.** SPIN rules executed to evaluate a hypothetical GAL gene induction event, their outcomes, and contribution to an overall hypothesis score assigned by HyQue

| SPIN Rule | Rule output | Score |
|---|---|---|
| penalizeNegation | Event is not negated | 0 |
| calculateInduceAgentTypeScore | Actor is a‘protein’ (CHEBI:36080) | +1 |
| calculateInduceTargetTypeScore | Target is a‘gene’ (SO:0000236) | +1 |
| calculateInduceLogical OperatorScore | Logical operator is ‘induce’ (GO:0010628) | +1 |
| calculateInduceAgentFunction Score | Actor does not have ‘transcription factor activity’ (GO:0003700) | 0 |
| calculateInduceLocationScore | Location is ‘nucleus’ (GO:0005634) | +1 |

The complete HyQue evaluations of this hypothesis as well as that of four additional hypotheses are available as RDF at the project website.

### 3.2 Changing a Domain Specific Rule Affects Hypothesis Evaluation

The `calculateInhibitEventScore` used to evaluate event 6 in section 3.1 in its current form does not take into account the physical location of the event participants. In other words, the score does not depend on data describing where the event participants are known (or not) to be located in the cell. However, some experimental evidence suggests that physical location in the context of an inhibition event plays an important role. Specifically, the inhibition of Gal4p activity by Gal80p is known to take place in the nucleus, yet this inhibition is interrupted when Gal80p is bound by Gal3p, which is typically found in the cytoplasm[18].

The effect of changing the `calculateInhibitEventScore`rule to require that all event participants be located in the nucleus to achieve a maximum score (a reasonable assumption given published findings[19]) on the hypothesis in section 3.1 would be that the score for *e6* is reduced. This is because adding an additional

sub-rule (let us call it `calculateInhibitEventParticipantLocationScore`) would increase the maximum score, while experimental data in the HKB is not available to satisfy the conditions of this new sub-rule – there is not experimental data available about the location of the Gal4p or Gal80p proteins in the cell. More specifically, let us say that the maximum score possible for `calculateInhibitEventScore` with the new sub-rule is now 4, and that event 6 is therefore assigned a score of 0.75 (3/4) based on the output of this rule. This changes the overall hypothesis score in that the first proposition (specifying events 1-3) now has a higher mean score (0.8, *versus* 0.78 for the second proposition as calculated using the new rule), and thus this is assigned as the overall hypothesis score.

This example demonstrates how using a different domain specific rule affects an overall hypothesis evaluation, and how the effect can be traced to both the rule(s) used and the data the rules are executed over.

## 4 Discussion

Using SPIN rules to evaluate HyQue hypotheses has several advantages. While HyQue "version 1.0" used SPARQL queries to obtain relevant statements from the HKB, the scoring rules used to evaluate those statements were hard-coded in system code. HyQue's SPIN evaluation rules can be represented as RDF, which allows the potential for users to query for HyQue rules that meet specific conditions, as well as potentially link to and aggregate those rules. In addition, users can create their own SPIN rules to meet specific evaluation criteria and augment existing HyQue rules to include them. In this way, different scientists may use the same data to evaluate the same hypotheses and arrive at unique evaluations depending on the domain principles encoded by the SPIN rules they use, as demonstrated in section 3.2. Encoding evaluation criteria as SPIN rules also ensures that the source of an evaluation can be explicitly stated, both in terms of the rules executed and the data the rules were executed over. This is crucial for formalizing the outcomes of scientific reasoning such that research conclusions can be confidently stated.

Separating HyQue system rules from the GAL domain specific rules highlights the two aspects of the HyQue scoring system. Specifically, HyQue currently encodes certain assumptions about how events in hypotheses may be related to one another, and how these relations are used to determine an overall hypothesis score, as well as domain specific assumptions about how to evaluate data in the context of knowledge about the GAL gene network. However, because assumptions about hypothesis structure are encapsulated by HyQue system rules, they may be changed or augmented without affecting the GAL domain specific rules, and *vice versa*. HyQue system rules can be extended over time to facilitate the evaluation of hypotheses that have fundamentally different structures than those currently presented as demonstrations. We envision a future iteration of HyQue where users can submit unique system and domain specific rules to use for evaluating hypotheses and in this way further research in their field by exploring novel interpretations of experimental data and hypotheses.

Similarly, it may be possible in future for HyQue users to select from multiple sets of evaluation rules and to compare the hypothesis evaluations that result.

Crafting SPIN rules requires knowledge of SPARQL, which, while being used in a number of life-science related projects[3, 5, 20-22], may present a barrier to some users. Similarly, representing hypotheses as RDF to submit to HyQue is not a trivial activity. To address the latter, we have developed an online form based system for specifying hypothesis details and converting them to RDF, available at the project website.

The Rule Interchange Format (RIF)[4] is the W3C standard for representing and exchanging rules between rule systems. SPIN, a W3C member submission, has been identified as an effort complimentary to RIF[23] and because there is some discussion of RIF and RDF compatibility[5], SPIN and RIF may become compatible if the RIF working group remains active[6]. HyQue provides a relevant use case and motivation for enabling such compatibility. Given that SPIN rules may be represented as RDF and executed over any RDF store using SPARQL (both W3C standards), however, and that the motivation of SPIN is specifically to execute SPARQL as rules, in the context of HyQue compatibility with RIF is not of immediate concern.

## 5 Conclusions

We present an extended version of HyQue that uses SPIN rules to evaluate hypotheses encoded as RDF, and makes the evaluation, including the data it is based upon, also available as RDF. In this way, users are able to explicitly trace a path from hypothesis to evaluation and the supporting experimental data, and *vice versa*. We have demonstrated how HyQue evaluates a specific hypothesis about the GAL gene network in yeast with an explanation of the scoring rules used and their outcomes. Evaluations of additional hypotheses, as well as HKB data and HyQue SPIN rules are available at http://hyque.semanticscience.org.

## References

1. Neylon, C., Wu, S.: Article-Level Metrics and the Evolution of Scientific Impact. PloS Biology 7(11), e1000242 (2009)
2. Howe, D., Costanzo, M., Fey, P., Gojobori, T., Hannick, L., Hide, W., Hill, D.P., Kania, R., Schaeffer, M., St Pierre, S., et al.: Big data: The future of biocuration. Nature 455(7209), 47–50 (2008)
3. Belleau, F., Nolin, M.-A., Tourigny, N., Rigault, P., Morissette, J.: Bio2RDF: towards a mashup to build bioinformatics knowledge systems. Journal of Biomedical Informatics 41(5), 706–716 (2008)

[4] http://www.w3.org/TR/2010/NOTE-rif-overview-20100622/
[5] http://www.w3.org/TR/2010/REC-rif-rdf-owl-20100622/
[6] http://www.w3.org/Submission/2011/02/Comment/

4. Tari, L., Anwar, S., Liang, S., Cai, J., Baral, C.: Discovering drug-drug interactions: a text-mining and reasoning approach based on properties of drug metabolism. Bioinformatics 26 (18), i547–i553
5. Nolin, M.A., Dumontier, M., Belleau, F., Corbeil, J.: Building an HIV data mashup using Bio2RDF. Briefings in Bioinformatics (2011)
6. Blonde, W., Mironov, V., Venkatesan, A., Antezana, E., De Baets, B., Kuiper, M.: Reasoning with bio-ontologies: using relational closure rules to enable practical querying. Bioinformatics 27(11), 1562–1568 (2011)
7. Villanueva-Rosales, N., Dumontier, M.: yOWL: an ontology-driven knowledge base for yeast biologists. Journal of Biomedical Informatics 41(5), 11 (2008)
8. Karp, P.D.: Artificial intelligence methods for theory representation and hypothesis formation. Comput. Appl. Biosci. 7(3), 301–308 (1991)
9. Karp, P.: Design Methods for Scientific Hypothesis Formation and Their Application to Molecular Biology. Machine Learning 12(1-3), 89–116 (1993)
10. Karp, P.D., Ouzounis, C., Paley, S.: HinCyc: a knowledge base of the complete genome and metabolic pathways of H. influenzae. In: Proc. Int. Conf. Intell. Syst. Mol. Biol., vol. 4, pp. 116–124 (1996)
11. Zupan, B., Bratko, I., Demsar, J., Juvan, P., Curk, T., Borstnik, U., Beck, J.R., Halter, J., Kuspa, A., Shaulsky, G.: GenePath: a system for inference of genetic networks and proposal of genetic experiments. Artif. Intell. Med. 29(1-2), 107–130 (2003)
12. King, R.D., Rowland, J., Oliver, S.G., Young, M., Aubrey, W., Byrne, E., Liakata, M., Markham, M., Pir, P., Soldatova, L., et al.: The automation of science. Science 324(5923), 85–89 (2009)
13. Soldatova, L., King, R.D.: Representation of research hypotheses. In: Bio-Ontologies 2010: Semantic Applications in Life Sciences, Boston, MA (2010)
14. Callahan, A., Dumontier, M., Shah, N.: HyQue: Evaluating hypotheses using Semantic Web technologies. In: Bio-Ontologies: Semantic Applications in the Life Sciences, Boston, MA (2010)
15. Callahan, A., Dumontier, M., Shah, N.H.: HyQue: evaluating hypotheses using Semantic Web technologies. J. Biomed. Semantics 2(suppl. 2), S3 (2011)
16. Ideker, T., Thorsson, V., Ranish, J.A., Christmas, R., Buhler, J., Eng, J.K., Bumgarner, R., Goodlett, D.R., Aebersold, R., Hood, L.: Integrated genomic and proteomic analyses of a systematically perturbed metabolic network. Science 292(5518), 929–934 (2001)
17. Racunas, S.A., Shah, N.H., Albert, I., Fedoroff, N.V.: HyBrow: a prototype system for computer-aided hypothesis evaluation. Bioinformatics 20(suppl. 1), i257–i264 (2004)
18. Bhat, P.J., Murthy, T.V.: Transcriptional control of the GAL/MEL regulon of yeast Saccharomyces cerevisiae: mechanism of galactose-mediated signal transduction. Mol. Microbiol. 40(5), 1059–1066 (2001)
19. Peng, G., Hopper, J.E.: Evidence for Gal3p's cytoplasmic location and Gal80p's dual cytoplasmic-nuclear location implicates new mechanisms for controlling Gal4p activity in Saccharomyces cerevisiae. Mol. Cell Biol. 20(14), 5140–5148 (2000)
20. Kobayashi, N., Ishii, M., Takahashi, S., Mochizuki, Y., Matsushima, A., Toyoda, T.: Semantic-JSON: a lightweight web service interface for Semantic Web contents integrating multiple life science databases. Nucleic Acids Res. 39(Web Server issue), W533–W540 (2011)

21. Chen, B., Dong, X., Jiao, D., Wang, H., Zhu, Q., Ding, Y., Wild, D.J.: Chem2Bio2RDF: a semantic framework for linking and data mining chemogenomic and systems chemical biology data. BMC Bioinformatics 11, 255 (2010)
22. Antezana, E., Kuiper, M., Mironov, V.: Biological knowledge management: the emerging role of the Semantic Web technologies. Brief Bioinform. 10(4), 392–407 (2009)
23. Polikoff, I.: Comparing SPIN with RIF (July 05, 2011), http://topquadrantblog.blogspot.com/2011/06/comparing-spin-with-rif.html (accessed December 7, 2011)