

Expressing High-Level Scientific Claims with Formal Semantics

Cristina-Iulia Bucur
Vrije Universiteit Amsterdam
Amsterdam, The Netherlands
c.i.bucur@vu.nl

Davide Ceolin
Centrum Wiskunde & Informatica
Amsterdam, The Netherlands
davide.ceolin@cwii.nl

Tobias Kuhn
Vrije Universiteit Amsterdam
Amsterdam, The Netherlands
t.kuhn@vu.nl

Jacco van Ossenbruggen
Vrije Universiteit Amsterdam
Amsterdam, The Netherlands
jacco.van.ossenbruggen@vu.nl

ABSTRACT

The use of semantic technologies is gaining significant traction in science communication with a wide array of applications in disciplines including the life sciences, computer science, and the social sciences. Languages like RDF, OWL, and other formalisms based on formal logic are applied to make scientific knowledge accessible not only to human readers but also to automated systems. These approaches have mostly focused on the structure of scientific publications themselves, on the used scientific methods and equipment, or on the structure of the used datasets. The core claims or hypotheses of scientific work have only been covered in a shallow manner, such as by linking mentioned entities to established identifiers. In this research, we therefore want to find out whether we can use existing semantic formalisms to fully express the content of high-level scientific claims using formal semantics in a systematic way. Analyzing the main claims from a sample of scientific articles from all disciplines, we find that their semantics are more complex than what a straight-forward application of formalisms like RDF or OWL account for, but we managed to elicit a clear semantic pattern which we call the “super-pattern”. We show here how the instantiation of the five slots of this super-pattern leads to a strictly defined statement in higher-order logic. We successfully applied this super-pattern to an enlarged sample of scientific claims. We show that knowledge representation experts, when instructed to independently instantiate the super-pattern with given scientific claims, show a high degree of consistency and convergence given the complexity of the task and the subject. These results therefore open the door on the longer run for allowing researchers to express their high-level scientific findings in a manner they can be automatically interpreted. This in turn will allow for automated consistency checking, question answering, aggregation, and much more.

CCS CONCEPTS

• **Computing methodologies** → **Knowledge representation and reasoning.**



This work is licensed under a Creative Commons Attribution International 4.0 License.

K-CAP '21, December 2–3, 2021, Virtual Event, USA.

© 2021 Copyright is held by the owner/author(s).

ACM ISBN 978-1-4503-8457-5/21/12.

<https://doi.org/10.1145/3460210.3493561>

KEYWORDS

Semantic Publishing; Knowledge Representation; Scholarly Communication; Formal Semantics

ACM Reference Format:

Cristina-Iulia Bucur, Tobias Kuhn, Davide Ceolin, and Jacco van Ossenbruggen. 2021. Expressing High-Level Scientific Claims with Formal Semantics. In *Proceedings of the 11th Knowledge Capture Conference (K-CAP '21), December 2–3, 2021, Virtual Event, USA*. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3460210.3493561>

1 INTRODUCTION

Semantic Web technologies have started to be more widely applied in science communication to address, among other things, the accelerating growth of scientific literature [8]. This growth makes it increasingly difficult for individual researchers to follow and be up-to-date with all the current developments in their fields [28]. However, despite semantic technologies being applied in many ways in science, there still is a big gap between the existing (and continuously developing) formal ontologies and the content of scientific articles expressed in natural language that is only human-readable [40]. As such, the content of scientific articles of today is not accessible to algorithms. Methods of semantic interlinking and metadata enrichment of scientific articles exist, especially in life sciences [12, 17, 18, 39], but they are mostly about annotating the existing natural language text in a relatively shallow manner. Other successful applications of semantic technologies focus only on the metadata level, such as provenance representation and versioning of scientific works [17].

Applying semantic technologies to not just annotate or describe (on a metalevel) but actually express scientific knowledge is a much more ambitious and much harder task. While technologies like RDF and OWL grounded in first-order logic are now mature and well-tested, full semantic representation of natural language has remained a task that is too complex for automated approaches. Even the extraction of simpler RDF-based structures from the content of scientific articles typically requires manual curation to achieve sufficient level of quality [3, 12, 20]. In fact, even for knowledge representation experts it is very challenging to fully represent typical high-level scientific findings in formal logic, as no general scheme or procedure exists that is known to apply to findings across specific fields.

In order to make progress on these challenges, we focus here on how general high-level scientific claims can be represented in formal logic. We assume in this work that these representations

are manually created, but it could also lay the basis for future research on how to automatically generate these representations from scientific texts (though we assume this to be a very hard problem, as it has been in the past). Concretely, we introduce below our approach to express the content of high-level scientific claims with formal semantics with a semantic template that we call the “super-pattern”.

In this research we therefore aim to answer the following research questions:

- (1) To what extent can our super-pattern be used to formalize the main claims of scientific articles from different disciplines?
- (2) How reliably can the super-pattern be applied to formalize existing claims by knowledge representation experts?

We envisage that in the longer-term future researchers could themselves express their findings with the super-pattern and thereby make their work directly add to a complex knowledge graph of scientific findings. On that basis, it is easy to imagine how technologies like graph query languages and logic reasoners can be used to find similar research, corroborate scientific claims, spot contradictions, provide aggregations and visualizations, answer questions, and many other kinds of tasks. Among countless other applications, this could save months of literature study for first-year PhD students to get an overview of their fields.

2 BACKGROUND

Despite being in the digital era, the medium in which the scientific articles are still published are journals, which are now electronically available, but otherwise still follow the old paper paradigm [30, 31, 40]. On top of this, the growing number of articles that are published every day makes it imperative to make the articles machine-readable as well, as human researchers are reaching their limits in terms of how much they are able to process and read every day [28, 36]. With current developments in the Semantic Web, with technologies like RDF, OWL, and SPARQL [8, 40], it is currently possible to enrich the meaning of a traditional article in the digital publishing environment and facilitate its automatic discovery, to have access in a semantic way within the article and link to other related articles or other related parts of articles [25, 29, 38].

Based on these Semantic Web technologies, there are various approaches that try to automatically extract such semantic information from articles to make them more machine-readable, covering the metadata [18] or the article structure [11], providing summarizations of the main concepts and ideas (human assisted and curated) [20], using annotations to semantically enhance articles [11, 17, 39] or semantically interlink them [12], and describing bibliographic references and citations [35]. On the one hand, we have rich vocabularies and ontologies like the SPAR ontology suite [34] that support the classical publishing process starting from formal descriptions of the structure of articles and article discourses [4], to the traditional publishing workflows [6], to accounting for all roles involved in the publishing process [22], and even to the current imperfect peer-reviewing process [37]. On the other hand, there are Artificial Intelligence tools that can identify plagiarism and predict the impact an article might have [10], to natural language processing techniques that make use of strategies such as entity

extraction to identify concepts and ideas in articles, assisted and curated by humans [20].

Data infrastructures such as Linked Open Data allow for connecting data between scientific publications [41], but this needs additional steps as concepts and relations need to be identified beforehand, and the content of the articles need to be interconnected semantically, usually by using humans in the loop [18, 39]. Moreover, more complex approaches that attempt to automate the process of extracting formal semantics from traditionally-structured scientific articles include techniques like the compositional and iterative semantic enhancement method (CSIE) [32], semantic lenses [5], and modelling the context of sentences from scientific articles in conceptual frameworks [13]. Semantic Web technologies are therefore extensively used but the underlying publishing paradigm has stayed the same [44].

New initiatives that try to change this old paradigm of publishing are proposed, together with new publishing workflows. This paradigm shift entails to move from the textual representation of information to a more data-centric one, similar to the proposal for the next-generation Web [1], where instead of documents, the interest shifts from the syntactic (e.g. HTML) to the semantic level (e.g. RDF, OWL). On the semantic level, we would be able to express the content and not just the structure of what is now in narrative documents. Moreover, formats that are based on HTML, like RASH [33], have been proposed, where scientific articles that include semantic annotations can be represented. Other initiatives involve using semantic representations from the start, written by the actual authors of the research in what is named *genuine semantic publishing* [27], and we have proposed in our earlier work to move from the monolithic form structure of scientific articles to smaller, more granular interconnected parts that each contain single scientific claims or statements from the start [7].

The Open Research Knowledge Graph [16] is another initiative that aims to make research articles machine-readable. With this approach, scientific entities are expressed as a semantically interconnected knowledge graph, populated by methods such as extracting scientific concepts from the abstracts of scientific articles with the help of annotators [2]. In the life sciences, especially in the biomedical fields, there have been many initiatives to create controlled vocabularies that can serve as the foundation to represent scientific knowledge in a structured way in order to capture evidence and scientific findings from research in a computable form [14, 21, 43] with specific markup languages [15] and exchange formats [9]. Some of these approaches target high-level claims, but only with restricted coverage, for example BEL [42] for specific kinds of biological relations.

Nanopublications [19] are a concept and technology of Linked Data containers that can be used to represent and share different kinds of scientific knowledge. They are expressed in a way that is fully formal, thus machine-interpretable and can be regarded as “minimal publications” due to the fact that they contain just three basic elements (represented in RDF): an assertion containing a small unit of information that is its main content (such as a scientific finding), provenance of the assertion (e.g. linking to the scientific methods used to derive the scientific finding in the assertion) and a publication information part about the nanopublication as a whole (e.g. when it was created and by whom). Nanopublications can not

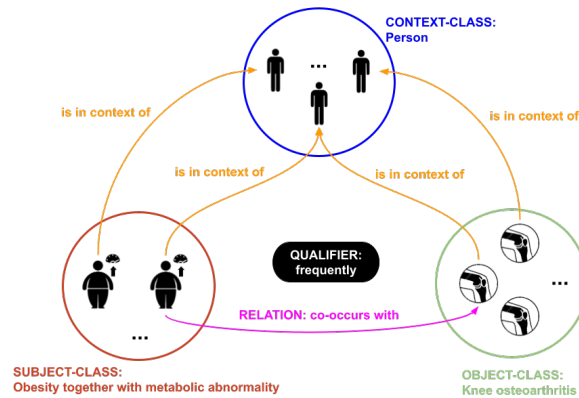


Figure 1: Diagram with an example for the Super-pattern instantiation.

only be an openly accessible structured data container for scientific findings or claims, but they can also be used for different types of meta-level assertions as well, for example for statements or assessments about other nanopublications [23].

3 APPROACH AND METHODS

We present here our approach to formalize high-level scientific claims by proposing what we call a “super-pattern”. This super-pattern is a general template of a logical statement that can be instantiated to represent scientific claims in formal logic. To aid the development of this super-pattern, we created a dataset of 50 claims from scientific publications (Dataset A). After the super-pattern had been finalized, we created a second set of another 25 claims (Dataset B), which did not influence our design decisions of the super-pattern and is therefore not biased towards it.

3.1 Super-Pattern

The super-pattern is a template with five slots, which are to be filled in when it is instantiated. It has the following structure with five elements with expected types in square brackets:

- **Context class** (“in the context of all ...”): [class identifier]
- **Subject class** (“things of type ...”): [class identifier]
- **Qualifier** (e.g. “mostly”): [qualifier from closed list]
- **Relation type** (“have a relation of type ...”): [relation type from closed list]
- **Object class** (“to things of type ... that are in the same context”): [class identifier]

The context class is optional, whereas all other slots are mandatory. The phrases in quotes above give some hints towards how to interpret it, but we will provide a thorough formal definition below. But let us first have a look at a concrete example.

In Figure 1 we see an example where the context class is named “person”, the subject class is named “obesity together with metabolic abnormality”, the qualifier is “frequently”, the relation is “co-occurs with” and the object class has the name “knee osteoarthritis”. By just filling in the informal phrases above, with a little bit of editing we get the following sentence, hinting at how to interpret it:

Table 1: Super-pattern qualifiers.

qualifier	interpretation	\lesseqgtr	q
(can) always	100%	=	1
(can) generally	at least 90%	\geq	0.9
(can) mostly	at least 50%	\geq	0.5
(can) frequently	at least 10%	\geq	0.1
(can) sometimes	at least 0.1%	\geq	0.001
(can) never	0%	=	0
(can) generally not	at most 10%	\leq	0.1
(can) mostly not	at most 50%	\leq	0.5
(can) frequently not	at most 90%	\leq	0.9
(can) sometimes not	at most 99.9%	\leq	0.999

In the context of all *persons*, things of type *obesity together with metabolic abnormality* frequently have a relation of type *co-occurs with* to things of type *knee osteoarthritis* that are in the same context (i.e. the same *person*).

It is easy to see that this is saying that when people have obesity together with metabolic abnormality then these people frequently also have a thing called knee osteoarthritis.

To arrive at a more formal representation grounded in logic, we can start by looking at the qualifiers like “generally”. Our super-pattern defines the qualifiers shown in Table 1. There are the five base qualifiers *always*, *generally*, *mostly*, *frequently*, and *sometimes*, each with its negative counterpart, such as *never* for *always* and *frequently not* for *frequently*. On top of that, each of these ten qualifiers can be amended by the operator *can*, such as *can generally*, which will be interpreted in a modal logic manner. The table shows how each qualifier is mapped to a quantified interpretation, such as *at most 90%* for *frequently not*. These interpretations are necessarily a bit arbitrary, but to reach our goal of fully formal semantics we had to define a precise value for these otherwise vague terms.

For our formalization to be described below, each of the 20 qualifiers is given a relation predicate \lesseqgtr , a probability value q , and a modality value m . The values for \lesseqgtr and q are shown in Table 1. The modality value m is \diamond for the qualifiers with *can* and the empty symbol ϵ otherwise.

We define the semantics of an instantiated super-pattern by the following logic expression template:

$$P(m(\exists z(o(z) \wedge i(z, x) \wedge r(y, z))) \mid s(y) \wedge c(x) \wedge i(y, x)) \lesseqgtr q$$

Here, c , s , and o stand for the context, subject, and object classes, respectively, of the instantiated super-pattern. Them being classes they map to unary predicates in the expression above. The super-pattern qualifier is expressed with the three parameters \lesseqgtr , q , and m , as explained above. r , finally, represents the relation of the instantiated super-pattern, mapped to a binary predicate. This relation comes from a fixed list of possible relations, consisting of sameness, numerical comparison, causality, and spatio-temporality relations. The full list of relations can be found in the online super-pattern ontology documentation¹ and in Table 3.

The remaining symbols are interpreted as follows. The symbol i stands for the general relation “is in the context of” that maps

¹https://larahack.github.io/linkflows_superpattern/doc/sp/index-en.html

instances in the subject and object classes to instances of the context class.² $P(\dots | \dots)$ represents the conditional probability function of standard statistics, but with the slightly non-standard convention of having logical expressions as elements. They are interpreted as the sets of tuples of all contained free variables (here, x and y) that satisfy the expression. z , finally, is a bound variable in the first of these expressions with the normal first-order semantics.

We can also phrase the $P(\dots | \dots)$ part as a matter of ratios instead of probabilities: It is the ratio of things that satisfy the left-hand side out of all the things that satisfy the right-hand side. This can be defined as such in higher-order logic, but we chose here the notationally simpler expression with conditional probabilities.

To get a better understanding of how the super-pattern works, let us walk through two cases. First, we can have a look at the effect of the *always* qualifier, for which we set \lesseqgtr to $=$, q to 1, and m to ϵ . The above expression then becomes (omitting the empty symbol ϵ , as it has no effect by definition):

$$P(\exists z(o(z) \wedge i(z, x) \wedge r(y, z)) | s(y) \wedge c(x) \wedge i(y, x)) = 1$$

This stating that the probability is 1, this is equivalent to saying that the right-hand side part of the conditional probability expression implies the left-hand side:

$$\exists x \exists y (s(y) \wedge c(x) \wedge i(y, x) \rightarrow \exists z(o(z) \wedge i(z, x) \wedge r(y, z)))$$

We see that this is expressing that for all things y in class s that are in the context of a thing x in class c , there exists another thing z in class o that is in the context of the same thing x such that there is a relation of type r from y to z , which matches how we introduced the super-pattern above.

As a second case, let us pick the more complex case of the qualifier *can generally not*, corresponding to the parameter values $\lesseqgtr = \leq$, $q = 0.1$, and $m = \diamond$. This gives us:

$$P(\diamond(\exists z(o(z) \wedge i(z, x) \wedge r(y, z)) | s(y) \wedge c(x) \wedge i(y, x)) \leq 0.1$$

This is stating that given a thing y of type s that is in the context of a thing x of type c , the probability (or “ratio”) that it is possible that there exists a thing z of type o in the context of the same thing x such that there is a relation of type r from y to z is less or equal 0.1. Or in other words, thing z as specified is possible under the given conditions in at most 10% of the cases.

The “it is possible that” here corresponds to the “can” in the qualifier and to the logical operator \diamond . It is interpreted according to standard modal logic with possible world semantics with the meaning of “there is an accessible possible world where the following is true”. As a more intuitive variation, we can also say “it can be made true that”. As we will see below, we often use these “can” qualifiers for statements from applied sciences like computer science or medicine, where findings are often stating that a certain kind of treatment or kind of software *can* solve a particular problem, but only if it is actually applied to it. As an extreme example, there might be strong evidence that a novel kind of medical treatment can treat a certain kind of condition, without a single instance of that actual condition having been treated in this way to that point in time. In that situation we need to express the possibility that it

could treat these conditions without necessarily treating them in reality.

For the case when the optional context slot is omitted, we can define a *universal context* class, which contains just one element that by definition has everything in its context. Leaving the context slot empty thereby corresponds to setting it to the universal context. In the expression above, x therefore corresponds to the universal context instance and the $i(\dots, \dots)$ components are always true and can be omitted. We therefore arrive at this simpler expression without context:

$$P(m(\exists z(o(z) \wedge r(y, z)) | s(y)) \lesseqgtr q$$

3.2 Dataset A

In order to see if the content of high-level scientific claims from different disciplines can be expressed with our super-pattern, we created a set of randomly selected scientific articles from Semantic Scholar. All data can be found online³. The general methodology for creating this Dataset A was as follows: (1) selecting a random sample of articles from different disciplines; (2) identifying a high-level scientific claim from each article; (3) applying the super-pattern on this claim with informal classes; and (4) formalizing the classes by finding existing identifiers or defining new ones.

Step 1: select articles. We selected a set of 50 random scientific articles by creating random numbers and then consulting Semantic Scholar⁴ to retrieve the article with that number as Semantic Scholar corpus ID, if it exists. We then checked whether it matched our inclusion criteria of being written in English, published in 2000 or later, having an abstract, and being a research paper. We repeated this process until we collected 50 such matching articles.

Step 2: identify main claim. We check the abstract to find the main scientific claim, finding, or conclusion, e.g. “In a cross-sectional study in Korean women, obesity showed closest association with knee osteoarthritis when accompanied by metabolic abnormality”. When several candidate claims existed, we chose the most high-level one, or if that did not resolve it, we simply chose the one mentioned first. We discarded articles that failed to clearly mention their main claim in the abstract. To arrive at a clear understanding of what exactly we want to formalize, we rephrased the found claims as AIDA sentences [24, 26]. These are single English sentences that are Atomic, Independent, Declarative, and Absolute. These properties make them a good tool to delineate what should go into the formalization to be created. For the claim above, the corresponding AIDA sentence could be “Obesity when accompanied by metabolic abnormality is closely associated with knee osteoarthritis”.

Step 3: apply the super-pattern. Next, we tried to apply the super-pattern by figuring out what classes would have to go into the three class slots, and which qualifier and relation would be most suitable. At this step, we do not yet check whether the needed classes are defined by any of the existing ontologies, but we simply refer to them by their (likely) class name. For the example given above, we would arrive at something similar to what is shown in Figure 1.

²A future version of the super-patterns might also make this relation configurable with more specific relations when the super-pattern is instantiated.

³https://github.com/LaraHack/linkflows_claims_dataset

⁴<https://www.semanticscholar.org/>

Step 4: find identifiers in existing ontologies. Next we tried to find identifiers in existing ontologies for the classes needed for the super-pattern instantiation from the last step. For this, we used (1) Wikidata⁵, the free knowledge base in which new concepts and properties can be added in a structured way, (2) BioOntology⁶, one of the biggest repositories containing biomedical ontologies and (3) Linked Open Vocabularies (or LOV)⁷, a big curated data collection of vocabularies that are reviewed and added continuously. All these chosen sources contain indexed vocabularies and ontologies and as such permit full text searches of concepts and relations. If we found multiple applicable candidate identifiers, we selected the one that seemed most suitable or best documented. For the cases where we could not find a matching identifier, we tried to find identifiers for parts of the required concept. These parts could then be used to construct a (partial) class definition, such as “obesity together with metabolic abnormality”, which can be defined as the intersection of the conditions “obesity” and “metabolic abnormality”. We mint new identifiers for such complex classes and also for the cases where we could not find any existing identifiers at all.

The process described above was itself the result of an iterative process during which the super-pattern was defined in its current form.

3.3 Dataset B

As the super-pattern was developed based on the claims found during the generation of Dataset A, we created an additional dataset that did not influence the super-pattern and can therefore provide a more reliable picture about what the range of scientific claims the super-pattern is able to express. For this Dataset B, we selected an extra set of 25 claims from another set of 25 randomly selected articles from Semantic Scholar, following the same procedure and criteria as for Dataset A above. We therefore arrived at another 25 formalizations in the form of instantiated super-patterns (to the extent it applied; more on this later).

In contrast to Dataset A, Step 3 for Dataset B was performed several times independently by different knowledge representation experts. Only after the independent super-pattern instances were created, the experts discussed and tried to reach an agreement. This will be explained in more detail below.

4 EVALUATION

In this section we will present the design and results of the evaluation for our approach. We performed a descriptive analysis and a vocabulary use analysis for the datasets introduced above. In order to assess how well and how consistently the super-pattern can be applied by knowledge representation experts, we ran a *formalization study* where such experts independently perform super-pattern based formalizations.

4.1 Descriptive Analysis

We performed a descriptive analysis on the combined datasets used in this research, dataset A and dataset B, one with 50 claims and the other with 25 other claims respectively.

⁵<https://www.wikidata.org>

⁶<https://biportal.bioontology.org/>

⁷<https://lov.linkeddata.es/>

Table 2: Usage of qualifiers in the dataset.

qualifier	sometimes	frequently	mostly	generally	always
positive	4	3	3	29	4
negative			1	5	
can positive	1			17	1
can negative					

91% of the claims (68 out of the total 75 claims) could be expressed with the super-pattern. The remaining ones were rather *too simple* and not too complex to be expressed, as they could be formalized with a simple *subject–predicate–object* triple in RDF style. The super-pattern was therefore successful in covering all the non-trivial scientific claims encountered. 76% of the claims (57 out of the total 75 claims) used the optional context slot of the super-pattern, whereas the remaining ones 15% (11 claims) used the simpler version of the pattern without context.

Next we can look at the distribution of qualifiers and relations for both datasets, which is shown in Tables 2 and 3. As we can see from Table 2, the most used qualifier is “generally” in almost 39% of cases (29 claims), together with its modal counterpart, “can generally” in 23% of cases (17 claims). The modal negative of qualifiers was never used, while using the negative for qualifiers seems to be less common, in just 6% of cases (6 claims), while the most used qualifiers are positive with 57% (43 claims) and modal positive with 25% (19 claims). In terms of the relations used in the datasets, Table 3 shows that relations that express causal relations are the most common with 57% (43 claims), then the equivalency relation “is same as” is the next most used with 16% (12 claims), then in a smaller ratio, the relations marking numerical comparisons (the “compares to” relations) and the relations about spatio-temporal relationships (the “has spatio-temporal relationship with”) are used in about 6-7% of cases (7 and 6 claims, respectively).

Overall, we see a broad but far from uniform distribution. Certain qualifiers and relation types were much more often used than others, which is in general not very surprising. That most claims are of a positive and causal nature is also consistent with expectations.

For Dataset A we can do an analysis of the classes used from the existing ontologies. We only restrict ourselves here to the top-level classes that could directly be used in the slots of the super-pattern, and we exclude here the classes and relations from existing ontologies that can be used to build a complex class definition.

In terms of the vocabulary usage for the classes for the subject, context, and object slots of the super-pattern, we noticed that most of these classes were not present in existing vocabularies, but had to be newly minted. In only 18 out of the 128 classes (14%), we could directly use an existing class identifier. Most of the time (16 cases), these identifiers came from Wikidata. Therefore we can say that formalizations that use the super-pattern mostly depend on defining new classes, which can typically only be partially defined from existing concepts and relations.

4.2 Design of Formalization Study

In order to evaluate the proposed method and our super-pattern, we designed a formalization study where several knowledge representation experts independently apply the super-pattern. The goal

Table 3: Usage of relations in the dataset.

relation	count	total per group
is same as	12	12
compares to	0	7
has similar value as	0	
has same value as	0	
has different value from	0	
has smaller value than	1	
has larger value than	6	
has causal relationship with	0	43
affects	2	
contributes to	11	
enables	6	
inhibits	2	
prevents	1	
increases	4	
decreases	3	
requires	4	
causes	5	
is necessary and sufficient for	2	
is caused by	3	
has spatio-temporal relationship with	0	6
includes	5	
is included in	0	
co-occurs with	1	
is followed by	0	
follows	0	

of this study was to find out how reliably a super-pattern could be created from a given scientific claim. The extent to which users who are not knowledge representation experts would be able to create such formalizations is beyond the scope of this work.

We have therefore designed a three-stage formalization study where the four co-authors of this article participated as knowledge representation experts. In the first stage, the experts independently instantiate the super-pattern for the given claims. In the second stage, each expert is asked to review all four formalizations (which are anonymized and randomly shuffled) from the first stage and select the best ones. In the last stage, all experts meet, discuss their choices, and are given the opportunity to adjust their selection. All data can be found online⁸.

Stage 1. The first stage started with an introduction session where the participating knowledge representation experts learned about the details of the super-pattern and could discuss any questions. The formalizations of Dataset A were used as examples for all the experts to have a common and consistent understanding of the super-pattern. After this joint introduction session, all experts worked independently to apply the super-pattern on the 25 claims of Dataset B. For each of these claims, the participants also had to rate (on a scale from 1 to 5) how confident they were in their formalization. For the class slots, the experts were only asked to provide class names, but not existing class identifiers.

Stage 2. In this stage, the participants were asked to independently inspect and review all four formalizations (i.e. the own one as

⁸https://github.com/LaraHack/linkflows_formalization_study

well as the ones from the other participants). These formalizations were randomly shuffled and it was not visible who created which of them. The participants were asked to select the best formalization (in their opinion) and to indicate which formalizations in their view contained clear formalization mistakes. If several equally good formalizations were present, participants were allowed to select more than one “best” formalization. After this stage, we can assess how much the participants agreed on the best formalization, and how much they agreed on the presence of formalization mistakes. As the task was complex and multiple valid solutions are possible, we can expect a substantial level of disagreement at this stage.

Stage 3. The goal of this last part of the formalization study was to see the extent to which disagreement was resolved when the experts were allowed to directly discuss the cases. In several discussion sessions, the experts went through all 25 claims and their candidate formalizations. The experts could explain why they chose a particular formalization as the best one and what kind of mistakes they identified. During this discussion, the experts could update their previous decisions, e.g. by choosing a different formalization as the best one, or by recognizing a formalization mistake they previously missed. After this stage, we can assess the degree of agreement when the experts have the chance to consult the others’ opinion and have the chance to react on that. This agreement can however also include the effect of social dynamics, such as participants trying to achieve (or prevent) consensus for social rather than conceptual reasons.

4.3 Formalization Study Results

For each of the 25 claims from Dataset B we have four different possible formalizations, therefore 100 formalizations in total. After Stage 2 (before the joint discussion), 38% of the individual formalizations were marked as containing a mistake by at least one of the experts. This dropped to just 6% after the discussion, showing agreement being increased by the discussion. Given the complexity of the task, these mistake ratios seem reasonably low. Overall, just 2% of the individual formalizations had both, at least one mistake mark as well as at least one best mark (this value didn’t change after discussion).

Table 4 summarizes the most important results from the Stages 2 and 3 of the formalization study. In order to analyze the level of agreement between the participants before and after the discussion meeting in Stage 3, we devised an agreement score of four levels A to D. Based on the participants’ best and mistake marks given, each of the 25 claims is assigned a level from A (most agreement) to D (least agreement).

In Level A, we include all scientific claims with a formalization that everybody agreed it was the best (or one of several “best”). In Level B, there is no full agreement but a majority agreement of three out of four participants on the best formalization, with no mistake mark from the remaining participant. Level C includes all scientific claims that are not in levels A or B, but for which there exists at least one formalization for which all participants agree that it contains no mistakes. Level D, finally, applies to all claims that are not in any of the categories above.

We see in Table 4 that full agreement on the best formalization was rare after Stage 2, with only 8% (2 claims of 25). The slightly less

Table 4: Levels of agreement in the formalization study.

Agreement level: before and after discussion	before (Stage 2)		after (Stage 3)	
	abs.	rel.	abs.	rel.
Level A: Full agreement on best formalization	2	0.08	21	0.84
Level B: Majority agreement on best formalization	12	0.48	3	0.12
Level C: Full agreement on absence of formalization mistakes	9	0.36	1	0.04
Level D: No agreement of the types above	2	0.08	0	0.00

restrictive majority agreement of Level B, however, was achieved in 48% cases, excluding the perfect agreement cases above. Therefore, in a majority of 56% of cases full or majority agreement on the best formalization was achieved, which seems quite remarkable given that this happened without discussion. Agreement on a formalization without mistakes (Level C) was achieved in another 36% of the cases, leaving only 8% to the lowest level D of no agreement at all. Again, considering the complexity and non-deterministic nature of the task, these seem very favorable outcomes that show that the super-pattern can be applied by knowledge representation experts in a reliable and consistent manner.

Looking at the agreement levels after the discussion meeting (Stage 3), we can observe that for most of the formalizations of the scientific claims, full agreement (84%) or majority agreement (12%) is reached. In all cases, there is at least a formalization on which all participants agree that it does not contain any mistakes. As expected, such a discussion session drives up agreement values substantially, which underlines the importance of checking the independent assessments of Stage 2 first. The results show that knowledge representation experts can agree on the best formalization in the vast majority of cases.

The individual pair-wise agreements between the four experts are visualized in Figure 2. To calculate these agreements, we use a value of 1 when the two participants fully agree on the status of an individual formalization (either “best”, “mistake”, or no mark), a value of 0.5 when they are different but without conflict (i.e. not “best” and “mistake” at the same time), and a value of 0 in the case of conflicting “best” and “mistake” marks. We then take the average of these values of the 100 data points per expert pair (4 formalization candidates for each of 25 claims).

We see that the average values before the discussion (Stage 2) are all above the 0.5 threshold (0.57 being the lowest value), and move quite close to full agreement after the discussion in Stage 3 (all values being over 0.82). Also, these pair-wise disagreement values are quite uniform across the pairs, and disagreement was therefore not driven by just one of the experts. Before as well as after discussion, the uniformity of the disagreement can be seen as an indication that it was rather due to the difficulty of the task and not to fundamental misunderstandings by some of the experts.

These pair-wise results on the individual formalization candidates therefore confirm the more general picture of Table 4. In the future, following the discussion in Stage 3 of the study, minor changes were proposed to the Super-Pattern ontology. As such, a future version might contain small changes to the definitions of some relations where, for example, the definition for the “includes” relation, will be “spatio-temporally or conceptually includes” instead of just “spatio-temporally includes”. In summary, we can

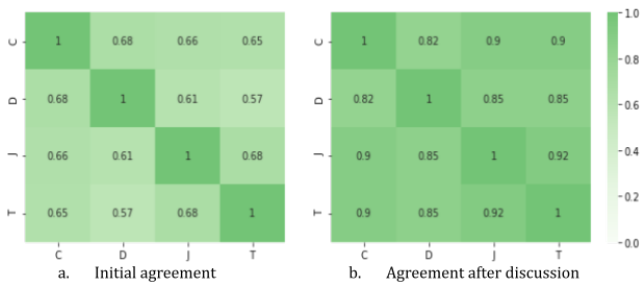


Figure 2: Agreement between participants on the formalizations before and after discussion.

conclude that writing a formalization that expresses the content of a high-level scientific claim using the super-pattern is an achievable task with a reasonably high success rate.

5 DISCUSSION AND CONCLUSION

While Semantic technologies have been extensively applied to a broad range of applications in various science disciplines, their use is quite limited with regards to science communication. The rapid increase in the number of scientific articles and the outlook to be able to communicate, link, and validate scientific discoveries in an automated manner should make it a priority to allow for the content of scientific articles to become readable not only by humans, but by machines as well. However, the use of these semantic technologies have mainly targeted the structure of the scientific publications themselves and not the content of the core scientific discoveries of these articles. In this research, we managed to show how we can create a framework in which we can represent these scientific claims with formal semantics.

Taking a random sample of scientific claims extracted from scientific articles from all disciplines, we succeeded to create formalizations of these scientific claims by using our super-pattern approach. It appears that this approach works for scientific claims from a variety of disciplines, confirming the cross-disciplinary nature and applicability of the super-pattern. Most of the classes used in the super-pattern formalizations had to be minted and could not be fully defined from existing vocabularies, but we argue that we managed to solve the more difficult part of the problem, as expressing the high-level logical structure of scientific claims is a more complex task than defining new classes. We could also demonstrate that knowledge representation experts can create such formalizations using the super-pattern in a fairly consistent and reliable manner, despite the inherent complexity and difficulty of this task.

In future work, we plan to investigate how our formal super-pattern representations can be used in practice to let researchers themselves, who are not necessarily knowledge representation experts, publish their own scientific claims. Furthermore, it will be interesting to investigate how advanced reasoning can be applied, considering that super-pattern formalizations use higher-order logic with its theoretical and practical problems around efficiency and decidability. In any case, any kind of reasoning with partial results or incomplete heuristics would constitute a huge advance compared to what we can currently do in terms of reasoning on scientific knowledge. We can then imagine countless novel applications, such as overarching aggregations, finding supporting or conflicting claims, answering high-level questions, making visualization of scientific knowledge, and much more. In summary, it would allow us to harness computers in a more effective way to increase our understanding of scientific discoveries and thereby amplify their impact.

ACKNOWLEDGMENTS

The authors thank IOS Press and the Netherlands Institute for Sound and Vision, who partly funded this research.

REFERENCES

- [1] Tim Berners-Lee and James A. Hendler. 2001. Publishing on the semantic web. *Nature* 410 (2001), 1023–1024.
- [2] Arthur Brack et al. 2020. Domain-Independent Extraction of Scientific Concepts from Research Articles. In *Advances in Information Retrieval*. https://doi.org/10.1007/978-3-030-45439-5_17
- [3] Adrien Coulet et al. 2020. Integration and publication of heterogeneous text-mined relationships on the Semantic Web. *J Biomed Semant* 2 (2011). <https://doi.org/10.1186/2041-1480-2-S2-S10>
- [4] Alexandru Constantin et al. 2016. The Document Components Ontology (DoCO). *Semantic Web* 7, 2 (February 2016), 167–181. <https://doi.org/10.3233/SW-150177>
- [5] Angelo Di Iorio et al. 2014. Semantic Lenses to Bring Digital and Semantic Publishing Together. In *ISWC'14*, Vol. 128, 12–23.
- [6] Aldo Gangemi et al. 2014. The Publishing Workflow Ontology (PWO). *Semantic Web* 8 (2014). <https://doi.org/10.3233/SW-160230>
- [7] Cristina-Iulia Bucur et al. 2020. A Unified Nanopublication Model for Effective and User-Friendly Access to the Elements of Scientific Publishing. *EKAW2020* (2020). https://doi.org/10.1007/978-3-030-61244-3_7
- [8] David Shotton et al. 2009. Adventures in semantic publishing: Exemplar semantic enhancements of a research article. *PLoS computational biology* 5 (2009), Issue 4.
- [9] Emek Demir et al. 2010. The BioPAX community standard for pathway data sharing. *Nat. Biotechnol.* 28 (2010). <https://doi.org/10.1038/nbt.1666>
- [10] Habeeb Ibrahim Abdul Razack et al. 2021. Artificial intelligence-assisted tools for redefining the communication landscape of the scholarly world. *Science Editing* 8 (2021). <https://doi.org/10.6087/kcse.244>
- [11] Jaana Taakis et al. 2015. Crowdsourced semantic annotation of scientific publications and tabular data in PDF. *SEMANTICS'15* (2015).
- [12] L. Garcia-Castro et al. 2013. Connections across Scientific Publications based on Semantic Annotations. *SEPublica* (2013). <https://doi.org/10.5167/UZH-82214>
- [13] M.A. Angrosh et al. 2014. Contextual information retrieval in research articles: Semantic publishing tools for the research community. *Semantic Web* 5 (2014), 261–293, Issue 4. <https://doi.org/10.5555/2786113.2786115>
- [14] Marcus C. Chibucos et al. 2014. Standardized description of scientific evidence using the Evidence Ontology (ECO). *Database : the journal of biological databases and curation* (2014). <https://doi.org/10.1093/database/bau075>
- [15] M. Hucka et al. 2003. The systems biology markup language (SBML): a medium for representation and exchange of biochemical network models. *Bioinformatics* 19 (2003). <https://doi.org/10.1093/bioinformatics/btg015>
- [16] Mohamad Yaser Jaradeh et al. 2019. Open Research Knowledge Graph: Next Generation Infrastructure for Semantic Scholarly Knowledge. In *KCAP'19*. <https://doi.org/10.1145/3360901.3364435>
- [17] Paolo Ciccarese et al. 2011. An open annotation ontology for science on web 3.0. *J Biomed Semant* 2 (2011). <https://doi.org/10.1186/2041-1480-2-S2-S4>
- [18] Paolo Ciccarese et al. 2012. Open semantic annotation of scientific publications using DOME0. *Journal of Biomedical Semantics* 3 (2012). <https://doi.org/10.1186/2041-1480-3-S1-S1>
- [19] Paul Groth et al. 2010. The anatomy of a nanopublication. *Information Services and Use* 30, 1-2 (2010). <https://doi.org/10.3233/ISU-2010-0613>
- [20] Pedro Sernadela et al. 2015. A Semantic Layer for Unifying and Exploring Biomedical Document Curation Results. *WBIO'2015* (2015). https://doi.org/10.1007/978-3-319-16483-0_2
- [21] Sumit Madan et al. 2019. The extraction of complex relationships and their conversion to biological expression language (BEL) overview of the BioCreative VI (2017) BEL track. *Database : the journal of biological databases and curation* (2019). <https://doi.org/10.1093/database/baz084>
- [22] Silvio Peroni et al. 2012. Scholarly publishing and Linked Data: describing roles, statuses, temporal and contextual extents. *i-Semantics 2012* (2012). <https://doi.org/10.1145/2362499.2362502>
- [23] Tobias Kuhn et al. 2013. Broadening the scope of nanopublications. In *Extended Semantic Web Conference*, 487–5017. https://doi.org/10.1007/978-3-642-38288-8_33
- [24] Tobias Kuhn et al. 2013. Broadening the Scope of Nanopublications (*ESWC'13*, Vol. 7882). *ESWC: European Semantic Web Conference*, Springer, 487–501. https://doi.org/10.1007/978-3-642-38288-8_33
- [25] Bryon Jacob and Jonathan Ortiz. 2017. Data.world: A Platform for Global-Scale Semantic Publishing. In *ISWC'17*.
- [26] Tobias Kuhn. 2018. Using the AIDA Language to Formally Organize Scientific Claims. *CNL 2018* (2018). <https://doi.org/10.3233/978-1-61499-904-1-52>
- [27] Tobias Kuhn and Michel Dumontier. 2017. Genuine semantic publishing. *Data Science* 1 (2017), 139–154, Issue 1/2. <https://doi.org/10.3233/DS-170010>
- [28] Esther Landhuis. 2016. Scientific literature: Information overload. *Nature* 535 (2016), 457–458. <https://doi.org/10.1038/nj7612-457a>
- [29] Silvia Mirri et al. 2017. Towards accessible graphs in HTML-based scientific articles. In *4th IEEE Annual Consumer Communications & Networking Conference (CCNC)*, 1067–1072. <https://doi.org/10.1109/CCNC.2017.7983287>
- [30] Silvio Peroni. 2012. *Semantic Publishing: issues, solutions and new trends in scholarly publishing within the Semantic Web era*. PhD Thesis. Universita di Bologna. <http://speroni.web.cs.unibo.it/publications/peroni-2012-semantic-publishing-issues.pdf>
- [31] Silvio Peroni. 2014. The Digital Publishing Revolution. In *Semantic Web Technologies and Legal Scholarly Publishing*, Pompeu Casanovas and Giovanni Sartor (Eds.). Law, Governance and Technology Series, Vol. 15. Springer International Publishing, Switzerland, Chapter 2, 7–43. https://doi.org/10.1007/978-3-319-04777-5_2
- [32] Silvio Peroni. 2017. Automating semantic publishing. *Data Science* 1 (2017), 155–173, Issue 1/2. <https://doi.org/10.3233/DS-170012>
- [33] Silvio Peroni et al. 2017. Research Articles in Simplified HTML: a Web-first format for HTML-based scholarly articles. *PeerJ Computer Science* 3 (2017), Issue e132. <https://doi.org/10.7717/peerj-cs.132>
- [34] Silvio Peroni and David Schotton. 2008. The SPAR Ontologies. *ISWC 2018 Proceedings* (2008). https://doi.org/10.1007/978-3-030-00668-6_8
- [35] Silvio Peroni and David Shotton. 2012. FaBiO and CiTO: ontologies for describing bibliographic resources and citations. *Web Semantics: Science, Services and Agents on the World Wide Web* 17 (December 2012), 33–43. <https://doi.org/10.1016/j.websem.2012.08.001>
- [36] Peter Gordon Roetzel. 2019. Information overload in the information age: a review of the literature from business administration, business psychology, and related disciplines with a bibliometric approach and framework development. *Business Research volume* 12 (2019). <https://doi.org/10.1007/s40685-018-0069-z>
- [37] Idafen Santana-Perez and Maria Poveda-Villalon. 2018. FAIR* Reviews Ontology (FR). Retrieved September 14, 2021 from <http://purl.org/spar/fr>
- [38] Bahar Sateli and René Witte. 2016. From Papers to Triples: An Open Source Workflow for Semantic Publishing Experiments. In *Semantics, Analytics, Visualization. Enhancing Scholarly Data*. Springer, 39–44. https://doi.org/10.1007/978-3-319-53637-8_5
- [39] Pedro Sernadela and Jose Luis Oliveira. 2017. A semantic-based workflow for biomedical literature annotation. *Database (Oxford)* (2017). <https://doi.org/10.1093/database/bax088>
- [40] David Shotton. 2009. Semantic publishing: the coming revolution in scientific journal publishing. *Learn. Publ.* 22 (2009), 85–94, Issue 2. <https://doi.org/10.1087/2009202>
- [41] Leslie F. Sikos. 2017. *Knowledge Representation with Semantic Web Standards*. Springer International Publishing, Cham, 11–49. https://doi.org/10.1007/978-3-319-54066-5_2
- [42] Ted Slater. 2014. Recent advances in modeling languages for pathway maps and computable biological networks. *Drug Discovery Today* 19 (2014). <https://doi.org/10.1016/j.drudis.2013.12.011>
- [43] Ted Slater and Diane H. Song. 2012. Saved by the BEL: ringing in a common language for the life sciences. *Fall* (2012).
- [44] Bodo M. Stern and Erin K. O'Shea. 2019. A proposal for the future of scientific publishing in the life sciences. *PLoS Biol* 17, 2 (2019), 683–684. <https://doi.org/10.1371/journal.pbio.3000116>