

A Unified Nanopublication Model for Effective and User-Friendly Access to the Elements of Scientific Publishing

Cristina-Iulia Bucur¹[0000–0002–7114–6459], Tobias Kuhn¹[0000–0002–1267–0234],
and Davide Ceolin³[0000–0002–3357–9130]

¹ Vrije Universiteit Amsterdam, Amsterdam, The Netherlands
`{c.i.bucur,t.kuhn}@vu.nl`

² Centrum Wiskunde & Informatica, Amsterdam, The Netherlands
`davide.ceolin@cw.nl`

Abstract. Scientific publishing is the means by which we communicate and share scientific knowledge, but this process currently often lacks transparency and machine-interpretable representations. Scientific articles are published in long coarse-grained text with complicated structures, and they are optimized for human readers and not for automated means of organization and access. Peer reviewing is the main method of quality assessment, but these peer reviews are nowadays rarely published and their own complicated structure and linking to the respective articles are not accessible. In order to address these problems and to better align scientific publishing with the principles of the Web and Linked Data, we propose here an approach to use nanopublications as a unifying model to represent in a semantic way the elements of publications, their assessments, as well as the involved processes, actors, and provenance in general. To evaluate our approach, we present a dataset of 627 nanopublications representing an interlinked network of the elements of articles (such as individual paragraphs) and their reviews (such as individual review comments). Focusing on the specific scenario of editors performing a meta-review, we introduce seven competency questions and show how they can be executed as SPARQL queries. We then present a prototype of a user interface for that scenario that shows different views on the set of review comments provided for a given manuscript, and we show in a user study that editors find the interface useful to answer their competency questions. In summary, we demonstrate that a unified and semantic publication model based on nanopublications can make scientific communication more effective and user-friendly.

1 Introduction

Scientific publishing is about how we disseminate, share and assess research. Despite the fact that technology has changed how we perform and disseminate research, there is much more potential for scientific publishing to become a more transparent and more efficient process, and to improve on the age-old paradigms

of journals, articles, and peer reviews [3, 28]. With scientific publishing often stuck to formats optimized for print such as PDF, we are not using the advances that are available to us with technologies around the Semantic Web and Linked Data [7, 35].

In this work, following our “Linkflows” vision [5], we aim to address some of these problems by looking at the scientific publishing process at a more finer-grained level and recording formal semantics for the different elements. Instead of treating big bulks of text as such, we propose to represent them as small snippets — e.g. paragraphs — that have formal semantics attached and can be treated as independent publication units. They can link to other such units and therefore form a larger entity — such as a full paper or review — by forming a complex network of links. With that approach, we can ensure that provenance of each snippet of information can be accurately tracked together with its creation time and author, and therefore allow for more flexible and more efficient publishing than the current paradigm. A process like peer-reviewing can then be broken down into small snippets and thereby take the specialization of reviewers and the detailed context of their review comments into account, and these review comments can formally and precisely link to exactly the parts of the paper they address. Each article, paragraph and each review comment thereby forms a single node in a network and is each identified by a dereferenceable URI.

We demonstrate here how we can implement such a system with the existing concept and technology of nanopublications, a Linked Data format for storing small assertions together with their provenance and meta-data. We then show how this approach allows us to build powerful and user-friendly interfaces to aggregate and access larger numbers of such small communication elements. In order to assess the concrete benefits, we zoom in to just one out of the countless scenarios in which we can expect substantial advantages from such fine-grained semantic representations. We chose here the concrete case of a system for editors to assess manuscripts based on a set of review comments, and based on this concrete case we demonstrate and assess our approach.

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

1. Can we use nanopublications as a unifying data model to represent the structure and links of manuscripts and their assessments in a precise, transparent, and provenance-aware manner?
2. Is a fine-grained semantic publishing and reviewing model able to provide us with answers to common competency questions that journal editors face in their work as meta-reviewers?
3. Can we design an intuitive and effective interface based on a fine-grained semantic publishing and reviewing model that supports journal editors in judging the quality of manuscripts based on the received reviews?

We address these research questions with the following contributions:

- A general scheme of how nanopublications can be used to represent and publish different kinds of interlinked publication elements

- A dataset of 627 nanopublications, implementing this scheme to represent exemplary articles and their open reviews
- A set of seven competency questions for the scenario of journal editors meta-reviewing a manuscript, together with SPARQL representations of these questions
- A prototype of a fine-grained semantic analysis interface for the above scenario and dataset, powered by nanopublications
- Results from a user study on the perceived importance of the above competency questions and the perceived usefulness of the above prototype for answering them

The rest of this article is structured as follows. In Section 2 we describe the current state of the art in the field of scientific publishing and the reviewing process in particular. In Section 3 we describe our approach with regard to performing the reviewing process in a fine-grained manner based on nanopublications. In Section 4.1 we describe in detail how we performed the evaluation of our approach, while we report and discuss the results of this evaluation in Section 4.2. Future work and conclusion of the present research are outlined in Section 5.

2 Background

Before we move on to describe our approach, we give here the relevant background on scientific publishing, semantic papers, and the specific concept and technology of nanopublications.

Scientific publishing is at the core of scientific research, which has moved in the last decades from print to online publishing [36]. It is, however, still mostly following the paradigm from the print age, with narrative articles being published in journals and assessed by peer reviewers, only the printed volumes having been replaced by PDF files that are made accessible via search engines [22]. Considering the ever increasing number of articles and the increasing complexity of research methods, this old paradigm of publishing seems to have reached its limit, and scientists are struggling to stay up to date in their specific fields [21]. Slowly but steadily, these old paradigms are shifting with open access publishing, semantically enriched content, data publication, and machine-readable metadata gaining momentum and importance [33, 37]. Opposition is also growing against the use of impact factor [9, 10, 24] or h-index as metrics for assessment of the participants in this publication process, and it has been shown that these metrics can be tampered with easily [1, 8, 29, 31].

Advances in Semantic Web technologies like RDF, OWL, and SPARQL have allowed for the semantic enhancement of scholarly journal articles when publishing data and metadata [32, 34]. As such, semantic publishing was proposed as a way to make scholarly publications discoverable, interactive, open and reusable for both, humans and machines, and to release them as Open Linked Data [13, 23, 30]. In order to extract formal semantics from already published

papers in an automated manner, sophisticated methods such as the compositional and iterative semantic enhancement method (CSIE) [25], conceptual frameworks for modelling contexts associated with sentences in research articles [2] and semantic lenses were developed [12]. Furthermore, HTML formats like RASH have been proposed to represent scientific papers that include semantic annotations [27], and vocabularies like the SPAR (Semantic Publishing and Referencing) suite of ontologies have been introduced to semantically model all aspects relevant to scientific publishing [26]. These approaches mostly work on already published articles, but it has been argued that scientific findings and their contexts should be expressed in semantic representations from the start by the researchers themselves, in what has been named *genuine semantic publishing* [18].

In our previous work [6], we applied the general principles of the Web and the Semantic Web to promote this kind of genuine semantic publishing [18] by applying it to peer reviews. We proposed a semantic model for reviewing at a finer-grained level called Linkflows and argued that Linked Data principles like dereferenceable URIs using open standards like RDF can be used for publishing small snippets of information, such as an individual review comment, instead of big chunks of text, such as an entire review. These small snippets of text can be represented as nodes in a network and can be linked with one another with semantically-annotated connections, thus forming distributed and semantically annotated networks of contributions. The individual review comments are semantically modeled with respect to what part of the paper they target, whether they are about syntax or content, whether they raise a positive or negative point, and whether they are a suggestion or compulsory, and what their impact on the quality of the paper is. We showed on this model that it is indeed beneficial if we capture these semantics at the source (i.e. the peer reviewer in this case).

Nanopublications [11] are a specific concept and technology based on Linked Data to publish scientific results and their metadata in small publication units. Each nanopublication has an assertion that contains the main content (such as a scientific finding), and comes with provenance about that assertion (e.g. what study was conducted to derive at the assertion; or which documents it was extracted from) and with publication information about the nanopublication as a whole (e.g. by whom and when it was created). All these three parts are represented in RDF and thereby machine-interpretable.

It has been shown how nanopublications can also be used for other kinds of assertions, including meta-statements about other nanopublications [15], and in order to make nanopublications verifiable and immutable, *trusty URIs* [17] can be used as identifiers, which include cryptographic hash values that are calculated on the nanopublication’s content. A decentralized server network has been established based on this, through which anybody can reliably publish and retrieve nanopublications [19]. In order to group nanopublications into larger collections and versions thereof, index nanopublications have been introduced [20]. With these technologies, small interconnected Linked Data snippets can be published in a reliable, decentralized, provenance-aware manner.

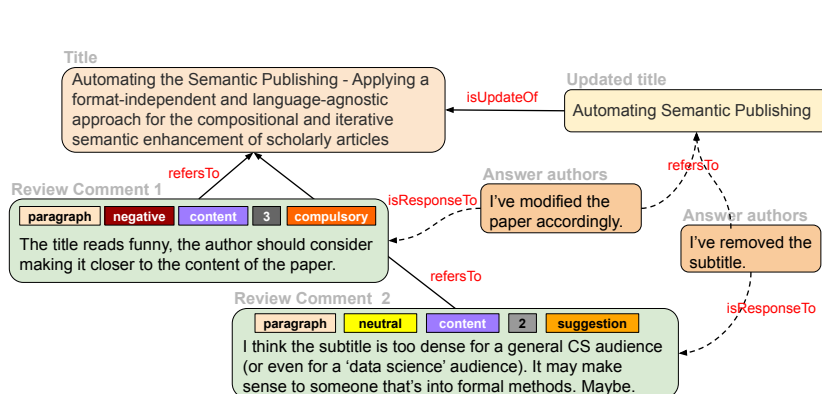


Fig. 1. An example of a nanopublication-style communication interaction.

3 Approach

Our general approach is to investigate the benefits of using the philosophy and technology of nanopublications as a unifying publishing unit to establish a new paradigm of scientific communication that is better aligned with the principles of the Web and Linked Data. We investigate how such an approach could allow us to communicate in a more efficient, more precise, and more user-friendly manner.

3.1 Semantic Model and Nanopublications

Our unifying semantic model based on nanopublications uses a number of existing ontologies like SPAR, PROV-O, FAIR* reviews, the Web Annotation Data Model, and our own Linkflows model [6] to break the big bulks of article and review texts into smaller text snippets. An example of a nanopublication-style communication interaction during the reviewing process is illustrated in Figure 1, where the title of a paper is addressed by several review comments that come with semantic classes (e.g. *suggestion*), which are themselves referred to by the authors' answers that link them to the updated version. Each node in this network is represented as a separate nanopublication and all the attributes and relations are formally represented as Linked Data.

As we can see in Figure 1, the properties *refersTo*, *isResponseTo*, *isUpdateOf* play the key role of linking different nodes in this network. *refersTo* is a property that links a review comment to the text snippet in the article it refers to. *isResponseTo* links the answer of the authors to the review comments of the reviewer and also to new versions of the text snippets that these review comments triggered. *isUpdateOf* links a version of the text snippet to another. In our approach, snippets of scientific articles (mostly corresponding to paragraphs) as well as their review comments (corresponding to individual review comments) are semantically represented as nanopublications [11], and thereby

```

@prefix this: <http://purl.org/np/RA28Qw0d22DK7U1Xq94UESAd3Y2jpGmVMCAk6_7aB_ivc> .
@prefix sub: <http://purl.org/np/RA28Qw0d22DK7U1Xq94UESAd3Y2jpGmVMCAk6_7aB_ivc#> .
@prefix dcterms: <http://purl.org/dc/terms/> .
@prefix xsd: <http://www.w3.org/2001/XMLSchema#> .
@prefix prov: <http://www.w3.org/ns/prov#> .
@prefix pav: <http://purl.org/pav/> .
@prefix np: <http://www.nanopub.org/nschema#> .
@prefix linkflows: <https://github.com/LaraHack/linkflows_model/blob/master/Linkflows.ttl#> .

sub:Head {
  this:np:hasAssertion sub:assertion ;
  np:hasProvenance sub:provenance ;
  np:hasPublicationInfo sub:pubinfo ;
  a np:Nanopublication .
}

sub:assertion {
  sub:comment-7 a linkflows:ContentComment, linkflows:NegativeComment, linkflows:ReviewComment, linkflows:SuggestionComment ;
  linkflows:hasCommentText "In section 3.2) you give background information on the Find-Fix-Verify pattern (2nd paragraph). This information (in which scenario it was used first, etc.) is not really necessary for the rest of the paper." ;
  linkflows:hasImpact "2"^^xsd:positiveInteger ;
  linkflows:refersTo <http://purl.org/np/Ra0Fc19T2f1F2j25Q6x8tfc1Z1LYAZrLD5IU6F_yBozGE#paragraph> .
}

sub:provenance {
  sub:assertion prov:hadPrimarySource <http://www.semantic-web-journal.net/content/detecting-linked-data-quality-issues-crowdsourcing-dbpedia-study> ;
  prov:wasAttributedTo <https://orcid.org/0000-0001-7069-9804> .
}

sub:pubinfo {
  this: dcterms:created "2019-11-26T09:05:11+01:00"^^xsd:dateTime ;
  pav:createdBy <https://orcid.org/0000-0002-7114-6459> .
}

```

Fig. 2. Example nanopublication of a review comment.

they each form a node in the network described above. A complete example of such a nanopublication containing a review comment is shown in Figure 2.

Each of the three main parts of a nanopublication — assertion, provenance, and publication info — is represented as an RDF graph. In the example of Figure 2, the assertion graph describes a review comment using the classes and properties of the Linkflows model³. It raises a negative point with an importance of 2 out of 5, and is marked as a suggestion for the authors. Furthermore, we see that this review comment refers to an external element, with a URI ending in **#paragraph**, as the target of this comment. This external element happens to be a paragraph of an article described in another nanopublication, which we can find out by following that trusty URI link.

Moreover, the nanopublication contains information regarding the creator of the assertion and the creator of the nanopublication that contains this assertion. These pieces of information can be found in the *provenance* and *publication info* graphs. As illustrated in Figure 2, the author of the review comment is indicated by his ORCID identifier and the source of the original source of the review comment is indicated by the the URL pointing to a link of the Semantic Web Journal. From the publication info graph, we can see who created the whole nanopublication together with the date and time of its creation.

As provenance and immutability of scientific contributions are crucial, we use trusty URIs [16] to enforce these properties. As such, for every nanopublication, in order for it to be published, a unique immutable URI is generated to refer to the node that holds the nanopublication. Any change of this nanopublication results in the generation of a new nanopublication, thus of a new node that is linked to the previous one. Such nanopublications can then be published in the existing decentralized nanopublication network [19].

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

3.2 Use Case with Competency Questions

In the scientific publishing context, editors of journals play a key role, being an important link between content providers for journals (authors), the people who assess the quality of the content (peer reviewers) and the consumers of such content (the readers). While the peer reviewers are the ones that can recommend the acceptance or rejection of an article, it is up to the editors to make the final decision. We will look here into how our approach can benefit the specific scenario of editors assessing a manuscript based on given reviews and having to write a meta-review.

Performing such a meta-review is not a trivial task. As classical reviews are mainly comprised of a large bulks of text in natural language, it is hard to provide a tool with quantitative information about the reviews and their collective implications on the manuscript. As such, an editor needs to spend a lot of time just to read these reviews fully to even get an overview of the nature and range of the raised issues.

In order to apply our approach to this chosen use case, we first define a set of competency questions (CQs), which are natural language questions that are created with the objective to assess the practicality and coverage of an ontology or model [4]. After consulting with publishing experts at IOS Press⁴ and the Netherlands Institute of Sound and Vision⁵ during an informal session, we came up with the following seven quantifiable competency questions from an editor's point of view during meta-reviewing:

- **CQ1:** *What is the number of positive comments and the number of negative comments per reviewer?*
- **CQ2:** *What is the number of positive comments and the number of negative comments per section of the article?*
- **CQ3:** *What is the distribution of the review comments with respect to whether they address the content or the presentation (syntax and style) of the article?*
- **CQ4:** *What is the nature of the review comments with respect to whether they refer to a specific paragraph or a larger structure such as a section or the whole article?*
- **CQ5:** *What are the critical points that were raised by the reviewers in the sense of negative comments with a high impact on the quality of the paper?*
- **CQ6:** *How many points were raised that need to be addressed by the authors, as an estimate for the amount of work needed for a revision?*
- **CQ7:** *How do the review comments cover the different sections and paragraphs of the paper?*

3.3 Dataset

In order to evaluate our approach on the given use case, we need some data first. For this, we selected three papers that were submitted to a journal that

⁴ <https://www.iospress.nl/>

⁵ <https://www.beeldengeluid.nl/en>

has open reviews (Semantic Web Journal). Therefore, we could also access the full text of the reviews these papers received. We then manually modelled all the article, paragraphs, review comments, their interrelations, as well as their larger structures — in the form of sections and full articles and reviews — as individual nanopublications according to our approach. All these elements were thereby semantically modeled, and we could reuse part of our earlier dataset of manually assigned Linkflows categories [6]. Figure 2 above shows an example of a nanopublication that resulted from this manual modeling exercise. We would like to stress here that according to the vision underlying our approach, these semantic representations would in the future be generated as such from the start, and therefore this manual effort is only for evaluation purposes. This should be integrated in the future in smart tools, such that this approach does not come at an additional burden for reviewers but in fact leads to a more efficient way of reviewing.

Apart from nanopublications at the lowest level, such as the one shown in Figure 2, higher-level ones combine them (by simply linking to them) to form larger structures, such as entire sections, papers, and reviews. Section nanopublications, for example, point to their paragraphs and define their order among other metadata. We also created a nanopublication index [20] that refers to this set of manually created nanopublications such that we can retrieve and even reuse parts of this dataset for new versions incrementally. All the nanopublications from our dataset are in an online repository⁶.

3.4 Interface Prototype for Use Case

In order to apply and evaluate our approach on the chosen use case, we developed a prototype of an editor interface that accesses the nanopublications in the dataset presented above to provide a detailed and user-friendly interface to support editors in their meta-reviewing tasks.

This prototype comes with two views: one where the review comments are shown per reviewer in a bar chart broken down into the different dimensions and classes, as shown in Figure 3 and another view that focuses on the distribution of the review comments to the different sections of the article, as shown in Figure 4. The interface for an exemplary article with three reviews can be accessed online⁷. The shown content is aggregated from nanopublications stored in a triple store and displayed by showing color codes for the different Linkflows classes for the individual review comments.

In the reviewer-oriented view (Figure 3), we can see in a more quantitative way the set of review comments and their types represented in different colors, where the checkboxes in the legend can be used to filter the review comments of the given category. To see the content of the review comments that are in a certain dimension, it is sufficient to just click on a bar in the chart.

⁶ https://github.com/LaraHack/linkflows_model_implementation

⁷ <http://linkflows.nanopubs.lod.labs.vu.nl>




Fig. 3. The reviewer-oriented view for the editor study.

The section-oriented view (Figure 4), aggregates all the finer-grained dimensions of the review comments at the level of sections in an article. Again, clicking on one cell in the table, thus selecting one specific dimension of the review comments, will show the content of those review comments underneath the table in the interface.

When data from the triple store is required, the server (implemented in NodeJS with the Express web application framework⁸) sends a request to the Virtuoso triple store where the nanopublications are stored. This request executes a SPARQL query on the stored nanopublications and returns the result to the server that, in turn, passes it further to the client, in the web browser, where the results are postprocessed and visualized. The code for the prototype can be found online⁹.

4 Evaluation

Here we present the evaluation of our approach in the form of a descriptive analysis, the analysis of the SPARQL implementations of our competency questions, and a user study with editors on our prototype interface.

⁸ <https://nodejs.org>, <https://expressjs.com/>

⁹ Interface: https://github.com/LaraHack/linkflows_interfaces

Backend application: https://github.com/LaraHack/linkflows_model_app

Data: https://github.com/LaraHack/linkflows_model_implementation

Article: "Detecting Linked Data Quality Issues via Crowdsourcing: A DBpedia Study"

Authors: Maribel Acosta, Amrapali Zaveri, Elena Simperl, Dimitris Kontokostas, Fabian Flock, Jens Lehmann

Date submitted: 04/15/2015 - 15:45

Section	Title	Aspect	Positivity/Negativity			Impact	Action needed										
			syntax	style	content		negative	neutral	positive	1	2	3	4	5	compulsory	suggestion	no action
1	Introduction				1	1									1		
2	Linked Data Quality Issues				3	3									3		2
3	Crowdsourcing Preliminaries				2	3					1	1	1				1
4	Our Approach: Crowdsourcing Linked Data Quality Assessment	1			7	8					4	3	1		8		
5	Evaluation				1	36	35			2	9	15	10	3	26	6	5
6	Final Discussions				1	1					1				1		
7	Related Work				5	5					1	1	2	1	3		2
8	Conclusions and Future Work																

[Review comments content for 'style' dimension of Section 3:](#)

style	negative	1	compulsory
--------------	-----------------	----------	-------------------

section 3.2 is clearly related to reference [3], so there is no need to include the citation several times

Fig. 4. The section-oriented view for the editor study.

4.1 Evaluation Design

First, we run a small descriptive analysis on the nanopublication dataset that we created. We can quantify the size and interrelation of the represented manuscripts and reviews in new ways, including the number of nanopublications, triples, paragraphs, review comments, and links between them.

As we are dealing with hundreds of individual nanopublications instead of just a hand-full of submission and review files, the performance of downloading them might pose a practical problem. For that reason, we also tested how long it takes to download all 627 nanopublications from the server network, using `nanopub-java` [14] as a command-line tool and giving it only the URI of the index nanopublication. This small download test was performed on a personal computer via a normal home network. For this, we retrieved them all via the library's `get` command and measured the time. We performed this 50 times, in five batches of 10 executions.

Next, we used our dataset to see if we are able to answer the seven competency questions that we defined above, in order to help editors in their meta-reviewing task. With this, we want to find out whether the combination of ontologies and vocabularies we used in our approach is sufficient to cover them, and then whether we can use the SPARQL query language to operationalize them and make them automatically executable on our nanopublication data.

Finally, we perform a user experiment involving editors to find out whether they indeed consider our competency questions important, and how useful they find our interface for getting an answer to these questions. For this study, we created a form that had two parts corresponding to the two parts of the study. We chose an article from our dataset that had a large number of review comments. For the first part, we asked for the importance of the competency questions using a Likert scale (from 1 to 5). For the second part, we provided static screenshots of our tool (the reviewer-oriented or the section-oriented view, depending on the question) together with a link to the live demo and asked about how useful the participants would find such a tool to answer the given competency question. The answers were on the same kind of a Likert scale from 1 to 5. We sent this

Table 1. Descriptive statistics dataset

part of article	number
articles	3
sections	89
paragraphs	279
figures	11
tables	10
formula	8
footnote	2
review comments	213

Table 2. Statistics nanopublications.

	number	average
Nanopublications:	627	
Head triples:	2508	4.00
Assertion triples:	5420	8.64
Provenance triples:	1254	2.00
Publication info triples:	1255	2.00
Total triples:	10 437	16.65

questionnaire (details online¹⁰) to a total of 401 editors of journals that support open reviews, specifically Data Science, the Semantic Web Journal and PeerJ Computer Science.

4.2 Evaluation Results

We can now turn to the results of these three parts of our evaluation. Details about the dataset and how it was generated and further queries and results can be found online¹¹.

Descriptive Analysis. Our representation of the three papers of our dataset together with their reviews leads to a total of 10 437 triples in 627 nanopublications, 279 text snippets and 213 review comments (85 for article 1, 59 for article 2 and 69 for article 3). Each of the three articles had three reviews: first article - 17, 18 and 50 review comments provided by the three reviewers, second article - 16, 21, 22 review comments each and third article - 11, 42, 16 review comments each.

In Table 1 some general statistics of the dataset are presented, while Table 2 shows general statistics about the nanopublications corresponding to the three articles and their reviews. Overall, this demonstrates the working of our approach of representing the elements of scientific communication in a fine-grained semantic manner. Of course, more complex analyses are possible, including network analyses of the complex interaction structure, and the queries for the competency questions that we defined above, to which we come back below.

Our small test on the performance of retrieving all nanopublications from the decentralized nanopublication network showed an average download time of 11.66 seconds overall (with a minimum of 8.39 and a maximum of 13.34 seconds). This operation retrieves each of the 627 nanopublications separately and then combines them in a single output file. The time per nanopublication is thereby just 18.6 milliseconds, which is achieved by executing the request in parallel to several servers in the network at the same time.

¹⁰ https://github.com/LaraHack/linkflows.editor_survey/

¹¹ https://github.com/LaraHack/linkflows.model_implementation

Competency Question Execution. In order to answer the competency questions in Section 3.2, we managed to implement each of them as a concrete SPARQL query. We can’t go into them here in detail due to space limitations, but the complete queries and all the required data and code can be found online¹².

This shows that our model is indeed able to capture the needed aspects for our competency questions, but we still need to find out whether these competency questions are indeed considered important by the editors, and whether the results from the SPARQL queries allow us to satisfy these users’ information needs. These two aspects are covered in our user study.

User Study Results. Out of the total 401 questionnaire requests sent, we received a total of 42 answers (10.5%). The importance of the seven competency questions for editors and the usefulness of the interface presented to answer these competency questions, assessed on a Likert scale from 1 to 5 where 1 is *not important at all* and 5 is *very important* can be seen in Table 3. We marked with * the competency questions that had a significant p -value (< 0.05) and without, the ones that were not significant. We calculate significance with a simple binomial test by splitting the responses into the ones that assign at least medium importance or usefulness (≥ 3) and the ones that assign low importance or usefulness (< 3).

We see the respondents declared high importance to five of the seven competency questions in a significant manner with average values from 3.05 to 4.58 (CQ1, CQ3, CQ4, CQ5 and CQ6), while the remaining two (CQ2 and CQ7) were not considered important in the editors’ view (average values of 2.36 and 2.79, respectively). Apparently, the number of positive and negative comments per section of the article (CQ2) and how the review comments cover the different parts of the article such as sections (CQ7), seem to have mixed reviews from editors, not being considered significantly important. The critical points that were raised by the reviewers (negative comments with a high impact on the paper) seems to be considered the most important competency question for the editors that responded (CQ5) with an average value of 4.58. Also important, in decreasing order, are the distribution of review comments with respect to whether they address the content or the presentation (syntax and style) of the article (CQ3), the number of points raised to be addressed by authors as an estimate for the amount of work needed for a reviewer (CQ6), the number of positive and negative comments per reviewer (CQ1), and the nature of the review comments with respect to whether they refer to a paragraph or a larger structure such as a section or the whole article (CQ4). For CQ2 and CQ7, we can say that editors did find it on average less important which sections of the article the reviews comments addressed. In general, however, we can conclude that most of competency questions are found to be important by most editors. However, we also observe a quite large standard deviation (SD) as seen in Table

¹² https://github.com/LaraHack/linkflows_model_implementation/tree/master/queries

Table 3. Results of the user study with editors.

Question	importance of question						usefulness of interface					
	AVG	MED	SD	count <3 ≥3	Δ count p -value		AVG	MED	SD	count <3 ≥3	Δ count p -value	
CQ1	3.17	3	1.36	15 27	0.044 *		3.48	4	1.17	9 33	1.36e-4 *	
CQ2	2.36	2	1.10	24 18	0.860		3.83	4	1.03	5 37	2.22e-7 *	
CQ3	3.64	4	0.93	5 37	1.36e-4 *		3.40	3.5	1.04	9 33	1.47e-3 *	
CQ4	3.05	3	1.19	14 28	0.022 *		3.26	3	1.20	14 28	0.022 *	
CQ5	4.58	5	0.63	0 42	< e-12 *		3.21	3	1.16	9 33	1.36e-4 *	
CQ6	3.57	4	1.02	6 36	1.41e-6 *		3.43	4	1.06	8 34	3.44e-5 *	
CQ7	2.79	3	1.12	18 24	0.220		3.62	4	1.03	5 37	2.22e-7 *	

3, ranging from 0.93 to 1.36 on our Likert scale that has a maximum distance of 4.0.

Next, we evaluated the usefulness of our prototype interface. Here the Likert scale went from 1 standing for *not useful at all* to 5 standing for *very useful*. As we can see from Table 3, this interface was on average considered useful for all of the seven competency questions, with averages ranging from 3.21 to 3.83. The preference for scores of 3 or larger is clearly significant for all of them. A substantial minority of respondents, however, didn’t find our interface useful leading again to relatively large standard deviation values between 1.06 and 1.19.

The free-text feedback field at the end of the questionnaire, finally, gave us a variety of suggestions for improvement (some of the editors argued that the interface used too many colors, others suggested other ways of grouping the data) but without clear overall tendencies. These responses also did not hint at any competency questions they found to be missing.

5 Discussion and Conclusion

Our results show that we can practically represent the different elements of scientific communication, such as articles and reviews, in a fine-grained and semantic way with nanopublications. We could show that we thereby can automatically answer a wide range of competency questions in the concrete scenario of editors in their meta-reviewing task. We found, however, that some of these were not found to be important, on average, by the editors who participated in our user study. Specifically, the questions about how well the review comments cover the different parts of the paper were not found to be important by a majority of editors. This could indicate that the article structure in terms of its different sections is not a good target for measuring the coverage of reviews. For all the questions, a relatively high variation is observed, which might be hinting at a lack of agreement among editors with respect to how scientific manuscripts should be assessed. This in turn could highlight the importance of more structured and more open reviewing processes. Irrespective of whether the competency questions are important, the majority of editors found our prototype to be useful to

answer them, although again with a large variation. With our approach focusing on interoperability and openness, however, it is not necessary to design a single interface that suits everybody, but we could allow editors to choose from several alternatives in the future.

In summary, we could show that nanopublications might be a suitable format not just for scientific findings but also for their reviewing processes. Their open and semantic nature can moreover allow other participants outside of the assigned editor and invited reviewers to contribute with their suggestions and comments, both before and after publication, while all the provenance needed to understand the context of each contribution is recorded. In this way, publication and reviewing as a whole might become more fluid, more inclusive, and more powerful.

Acknowledgements. This research was partly funded by IOS Press and the Netherlands Institute for Sound and Vision. The authors would like to thank Stephanie Delbeque, Maarten Fröhlich, Erwin Verbruggen, Johan Oomen, and Jacco van Ossenbruggen for providing their insight and expertise.

References

1. Alberts, B.: Impact factor distortions. *Science* **340**, 787–787 (2013). <https://doi.org/10.1126/science.1240319>
2. Angrosh, M., Cranefield, S., Stanger, N.: Contextual information retrieval in research articles: Semantic publishing tools for the research community. *Semantic Web* **5**, 261–293 (2014). <https://doi.org/0.5555/2786113.2786115>
3. Berners-Lee, T., Hendler, J.: Publishing on the semantic web. *Nature* **410**, 1023–1024 (2001). <https://doi.org/10.1038/35074206>
4. Bezerra, C., Freitas, F., Santana, F.: Evaluating ontologies with competency questions. In: 2013 IEEE/WIC/ACM International Joint Conferences on Web Intelligence (WI) and Intelligent Agent Technologies (IAT). vol. 3, pp. 284–285 (2013). <https://doi.org/10.1109/WI-IAT.2013.199>
5. Bucur, C.I.: Linkflows: Enabling a web of linked semantic publishing workflows. In: European Semantic Web Conference. pp. 262–271. Springer (2018). https://doi.org/10.1007/978-3-319-98192-5_45
6. Bucur, C.I., Kuhn, T., Ceolin, D.: Peer reviewing revisited: Assessing research with interlinked semantic comments. In: In K-CAP 2019: Proceedings of the 10th International Conference on Knowledge Capture. pp. 179–187 (2019). <https://doi.org/3360901.3364434>
7. Clark, T.: Next generation scientific publishing and the web of data. *Semantic Web* **5**, 257–259 (2014). <https://doi.org/10.3233/SW-140139>
8. Dong, P., Loh, M., Mondry, A.: The ‘impact factor’ revisited. *Biomed Digital Libraries* **2** (2005). <https://doi.org/10.1186/1742-5581-2-7>
9. Garfield, E.: Journal impact factor: a brief review. *CMAJ* **161**, 979–980 (1999)
10. Garfield, E.: The history and meaning of the journal impact factor. *JAMA* **295**, 90–93 (2006). <https://doi.org/10.1001/jama.295.1.90>
11. Groth, P., Gibson, A., Velterop, J.: The anatomy of a nanopublication. *Information Services & Use* **30**, 51–56 (2010). <https://doi.org/10.3233/ISU-2010-0613>

12. Iorio, A.D., Peroni, S., Vitali, F., Zingoni, J.: Semantic lenses to bring digital and semantic publishing together. In: Proceedings of the 4th International Conference on Linked Science@ISWC. vol. 128, pp. 12–23 (2014)
13. Jacob, B., Ortiz, J.: Data.world: A platform for global-scale semantic publishing. In: Proceedings of the ISWC posters & demonstrations and industry tracks co-located with 16th International Semantic Web Conference (ISWC) (2017)
14. Kuhn, T.: nanopub-java: A java library for nanopublications. In: Proceedings of the 5th Workshop on Linked Science (LISC 2015). vol. 1572, pp. 19–25 (2015)
15. Kuhn, T., Barbano, P.E., Nagy, M.L., Krauthammer, M.: Broadening the scope of nanopublications. In: Extended Semantic Web Conference. pp. 487–5017 (2013). https://doi.org/10.1007/978-3-642-38288-8_33
16. Kuhn, T., Dumontier, M.: Trusty uris: Verifiable, immutable, and permanent digital artifacts for linked data. In: Proceedings of the 11th Extended Semantic Web Conference (ESWC). vol. 8465, pp. 395–410 (2014). https://doi.org/10.1007/978-3-319-07443-6_27
17. Kuhn, T., Dumontier, M.: Making digital artifacts on the web verifiable and reliable. *IEEE Transactions on Knowledge and Data Engineering* **27**, 2390–2400 (2015). <https://doi.org/10.1109/TKDE.2015.2419657>
18. Kuhn, T., Dumontier, M.: Genuine semantic publishing. *Data Science* **1**, 139–154 (2017). <https://doi.org/10.3233/DS-170010>
19. Kuhn, T., et al.: Decentralized provenance-aware publishing with nanopublications. *PeerJ Computer Science* **2** (2016). <https://doi.org/10.7717/peerj-cs.78>
20. Kuhn, T., et al.: Reliable granular references to changing linked data. In: ISWC Lecture Notes in Computer Science. vol. 10587 (2017). https://doi.org/10.1007/978-3-319-68288-4_26
21. Landhuis, E.: Scientific literature: Information overload. *Nature* **535**, 457–458 (2016). <https://doi.org/10.1038/nj7612-457a>
22. Lippi, G., et al.: Scientific publishing in the “predatory” era. *Clinical Chemistry and Laboratory Medicine (CCLM)* **56**(5), 683–684 (2018). <https://doi.org/10.1515/cclm-2017-1079>
23. Mirri, S., et al.: Towards accessible graphs in html-based scientific articles. In: 4th IEEE Annual Consumer Communications & Networking Conference (CCNC). pp. 1067–1072 (2017). <https://doi.org/10.1109/CCNC.2017.7983287>
24. Opthof, T.: Sense and nonsense about the impact factor. *Cardiovascular Research* **33**, 1–7 (1997). [https://doi.org/10.1016/S0008-6363\(96\)00215-5](https://doi.org/10.1016/S0008-6363(96)00215-5)
25. Peroni, S.: Automating semantic publishing. *Data Science* **1**, 155–173 (2017). <https://doi.org/10.3233/DS-170012>
26. Peroni, S., Shotton, D.: The spar ontologies. In: Proceedings of the 17th International Semantic Web Conference. vol. 128, pp. 119–136 (2018). https://doi.org/10.1007/978-3-030-00668-6_8
27. Peroni, S., et al.: Research articles in simplified html: a web-first format for html-based scholarly articles. *PeerJ Computer Science* **3** (2017). <https://doi.org/10.7717/peerj-cs.132>
28. Priem, J.: Beyond the paper. *Nature* **495**, 437–440 (2013). <https://doi.org/10.1038/495437a>
29. Saha, S., Saint, S., Christakis, D.A.: Impact factor: a valid measure of journal quality? *JMLA* **91**, 42–46 (2003)
30. Sateli, B., Witte, R.: From papers to triples: An open source workflow for semantic publishing experiments. In: *Semantics, Analytics, Visualization. Enhancing Scholarly Data*. pp. 39–44. Springer International Publishing, Cham (2016). https://doi.org/10.1007/978-3-319-53637-8_5

31. Seglen, P.O.: Why the impact factor of journals should not be used for evaluating research. *BMJ* **314**, 498–502 (1997). <https://doi.org/10.1136/bmj.314.7079.497>
32. Shotton, D.: Semantic publishing: the coming revolution in scientific journal publishing. *Learn. Publ.* **22**, 85–94 (2009). <https://doi.org/10.1087/2009202>
33. Shotton, D.: The five stars of online journal articles - a framework for article evaluation. *D-Lib Magazine* **18**, 457–458 (2012). <https://doi.org/10.1045/january2012-shotton>
34. Shotton, D., Portwin, K., Klyne, G., Miles, A.: Adventures in semantic publishing: Exemplar semantic enhancements of a research article. *PLoS computational biology* **5** (2009). <https://doi.org/10.1371/journal.pcbi.1000361>
35. Sikos, L.F.: *Knowledge Representation with Semantic Web Standards*, pp. 11–49. Springer International Publishing, Cham (2017). https://doi.org/10.1007/978-3-319-54066-5_2
36. Stern, B.M., O’Shea, E.K.: A proposal for the future of scientific publishing in the life sciences. *PLoS Biol* **17**(2), 683–684 (2019). <https://doi.org/10.1371/journal.pbio.3000116>
37. Wang, P., Rath, M., Deike, M., Qiang, W.: Open peer review: An innovation in scientific publishing. In: *ICConference 2016 Proceedings* (2016). <https://doi.org/10.9776/16315>