Published in: El-Assady, Mennatallah/Hautli-Janisz, Annette/Lyding, Verena (Eds.): Proceedings of the LREC 2018 Workshop. The 3rd workshop on visualization as added value in the development, use and evaluation of language resources (VisLR II), 12 May 2018 – Miyazaki, Japan. Paris, France: ELRA, 2018. Pp. 1-8

Bridging Standards Development and Infrastructure Usage by Means of Concept Graphs: the Liaison of CLARIN and ISO TC37SC4 in Practice

Karolina Suchowolec[1], Piotr Bański[2], Andreas Witt[2,3]
[1]Technische Hochschule Köln, [2]Institut für Deutsche Sprache, [3]Universität zu Köln
[1]Campus Südstadt, Ubierring 48, 50678 Köln, [2]R5, 6-13, 68161 Mannheim, [3]Albertus Magnus Platz, 50923 Köln
karolina.suchowolec@th-koeln.de,{banski, witt}@ids-mannheim.de, andreas.witt@uni-koeln.de

Abstract

The present submission reports on a pilot project conducted at the Institute for the German Language (IDS), aiming at strengthening the connection between ISO TC37SC4 "Language Resource Management" and the CLARIN infrastructure. In terminology management, attempts have recently been made to use graph-theoretical analyses to get a better understanding of the structure of terminology resources. The project described here aims at applying some of these methods to potentially incomplete concept fields produced over years by numerous researchers serving as experts and editors of ISO standards. The main results of the project are twofold. On the one hand, they comprise concept networks dynamically generated from a relational database and browsable by the user. On the other, the project has yielded significant qualitative feedback that will be offered to ISO. We provide the institutional context of this endeavour, its theoretical background, and an overview of data preparation and tools used. Finally, we discuss the results and illustrate some of them.

Keywords: CLARIN, ISO, standardisation, terminology, concept systems, network analysis, terminology visualisation, interactive graph visualisation

1. Introduction

The present submission reports on a pilot project conducted at the Institute for the German Language (IDS), aiming at strengthening the connection between ISO TC37SC4 "Language Resource Management" and CLARIN[1]. The former is a subcommittee of an international standards body focusing on the codification of norms for various aspects of language resource creation, management, and use, whereas the latter is a federated research infrastructure, currently bringing together dozens of centres of various types in almost 20 countries, mainly but not exclusively in Europe.

A simplified approach would paint ISO as the producer of standards and CLARIN as a consumer, but in fact, the relationship holds in both directions: numerous CLARIN researchers are also ISO experts, and CLARIN centres are responsible for the emergence of best practices that often end up becoming codified into standards. This relationship has recently been strengthened also on the political plane: since June 2017, ISO TC37SC4 and CLARIN-ERIC have a formal liaison agreement. The project described here may be seen as a practical expression of the new level of cooperation that both partners have entered.

As a CLARIN centre, IDS offers access to an interactive information system about standards endorsed or simply used by CLARIN centres. In order to make this system even more comprehensive, we plan to add to it a visualisation of the concept network extracted from standards published by ISO TC37 SC4. The present submission describes steps taken towards that goal and discusses the results.

[1] ISO TC37SC4 stands for Subcommittee 4 ("Language resource management") of Technical Committee 37 "Terminology and other language and content resources", cf. https://www.iso.org/committee/297592.html. CLARIN stands for "Common Language Resources and Technology Infrastructure", cf. https://www.clarin.eu/. CLARIN is an ERIC (European Research Infrastructure Consortium).

Section 2 couches the project in a theoretical background, section 3 provides a view of the current ISO offer, section 4 describes the aims of the project and the technical preparations, while section 5 discusses the results. Section 6 focuses on the evaluation of the results, and section 7 outlines further directions for research.

2. Theoretical Background: Definitions and Concept Systems in Terminology Work

ISO Standard 10241-1:2011 ("Terminological entries in standards – Part 1: General requirements and examples of presentation") provides recommendations on how to prepare and write terminological entries in ISO standards. Those recommendations follow the general terminology approach as described in e.g. ISO 704:2009-11 ("Terminology work – Principles and methods"). This approach is concept-oriented and in the following we focus on its two elements: definitions and concept systems.

2.1 Definitions

Definitions are a central means in terminology management for concept description. Definitions used in standards or other normative terminology sources differ from definitions used in other applications.

As stated in ISO 704:2009-11 (p. 22–24), in normative terminological resources, definitions are needed to clearly identify concepts in a domain and distinguish them from one another. Preferable definitions are short and concise, giving just enough information for concept distinction. This means that a definition in a standard does not aim at providing comprehensive or complete understanding, and that distinguishes it from an encyclopaedic definition. Moreover, definitions should be as general as possible and not limited to one standard only. In particular, a definition should, whenever possible, "be appropriate for other standards within closely related subjects." (ISO10241-1:2011 p. 26)

As for structure and wording, the so-called intensional definition is preferred. It begins with a hyperonym of the concepts and continues with delimiting characteristics. A concept is not defined in isolation, but in relation to other concepts in the domain, which might be referenced in the definition body: "*Definitions* shall be systemic in order to enable a terminologist to reconstruct the *concept system*." (ISO 704:2009-11 p. 23) Given the above-mentioned requirements, we can expect that redundancy, ambiguity and the use of underspecified terms in standard definitions should be avoided.

2.2 Concept Systems

The general terminology approach stresses the importance of arranging concepts into concept systems, which are defined as "set[s] of concepts [...] structured according to the relations among them" (ISO 10241-1:2011 p. 3). More loosely structured sets are called concept fields. Concept systems are crucial for terminology management in many ways. A concept definition should, preferably, reflect the position of a concept within the concept system by referencing other related concepts (e.g. ISO 10241-1:2011, 2011 p. 26). Also, concept systems help standardization bodies in assigning normative status (*preferred*, *admitted*, *deprecated*) to synonymous terms. For terminology users, concept systems offer additional access to the data and help to explore the structure of the given domain.

Although concept systems are generally recognized as a useful tool in terminology management for both authors and users, there has been a lack of appropriate tools to dynamically and interactively generate concept systems from terminology databases (e.g. Drewer, Massion and Pulitano, 2017 p. 21). Hence, one had to resort to static means of visualisation such as regular drawings, which made the creation and maintenance of concept systems rather cumbersome.

Because concept systems play such an essential role in creating a coherent and well-formed terminology resource, a new generation of terminology tools has recently emerged that address the gap of dynamic visualisation (cf. Früh and Deubzer, 2016). In addition, more and more projects use generic dynamic visualisation technology and complement their terminology database with an additional visual access (e.g. *Verweis Viewer*: Chiocchetti, Culy and Ralli, 2015; *EcoLexicon*: Faber, León-Araúz and Reimerink, 2016; Lang, Schwinn and Suchowolec, in press).

Various display modes for concept systems have been proposed (DIN 2331:1980-04, 1980), but from the point of view of the end user, visualisation as a graph is often considered to be the desired one. A graph can be defined as "a set of *nodes*, which are an abstraction of any entities […], and the connecting links between pairs of nodes called *edges* or relationships." (Igual and Seguí, 2017 p. 142) A graph can be undirected or directed, depending on whether the relationship between edges is symmetric or not. There are several characteristics for analysing a graph (or network). Most importantly, the so-called centrality measures compute the importance of a node within a graph. Igual and Seguí (2017 p. 147ff) mention the following measures for graph analysis:

- degree centrality: „number of edges of the node"
- betweenness centrality: "number of times a node is crossed along the shortest path/s between any other pair of nodes"
- closeness centrality: "[quantified] position a node occupies in the network based on a distance calculation"
- eigenvector centrality: "relative score for a node based on its connections"
- PageRank: an algorithm developed for Google to "rate webpages objectively and effectively measure the attention devoted to them. [...] The rank of page P_i is the probability that a surfer on the Internet who starts visiting a random page and follows links, visits the page P_i." (p. 154–156). PageRank can be applied to the nodes of any directed graph.

In terminology management, attempts have recently been made to use graph-theoretical analyses to get a better understanding of the structure of terminology resources (Falke, Lang and Suchowolec, 2017). The pilot project described here aims at applying some of these methods to potentially incomplete concepts (concept fields) produced over years by numerous researchers serving as experts and editors of ISO standards. Apart from supplying proof of concept, the project intends to provide ISO with feedback of various sorts: on what can be done to make ISO definition networks more robust, and on the existing problems that can be eliminated quickly and improve the coherence of the existing networks.

3. Terminological Data in ISO Online Browsing Platform

The present section describes the current state of the ISO Online Browsing Platform[2] – a front-end to the ISO terminological database, where a user can get a glimpse of the content of standards by browsing selected sections (see figure 1).

3.1 Terminological Entry

Terminological entries of the OBP are concept-oriented and in general follow the recommendations of ISO. They can be displayed only as a list (sorted by relevance or by term), in two different views: the basic view and the full view.

The basic view includes one term per concept (no synonyms), its definition, the name of the standard where the concept is defined and the relevant definition paragraph number.

The full view ("full entry") might additionally include all possible terms, (scope) notes, as well as examples. In case the definition is re-used from a different standard, the entry might also contain the reference to the original source of the definition.

3.2 Conceptual Relations

Within a terminological entry of a standard, other concepts are explicitly referenced in the definition body – concepts mentioned in a definition of a given concept are highlighted and cross-referenced with a hyperlink and the target paragraph number, as illustrated below:[3]

[2] https://www.iso.org/obp/ui

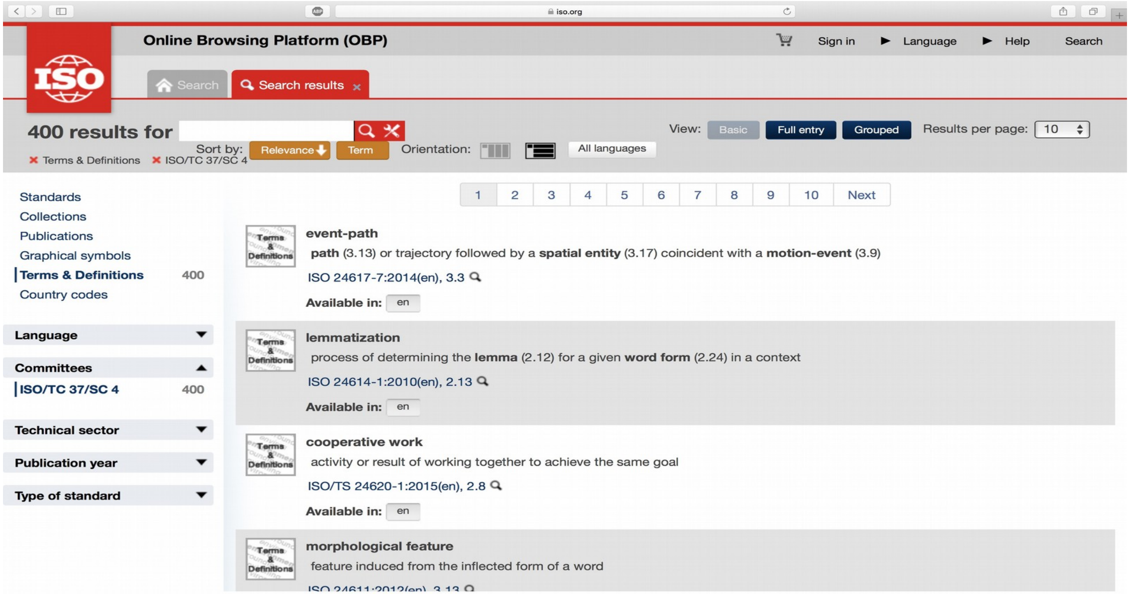

Figure 1: Screenshot of the Terminological Component of the ISO Online Browsing Platform (retrieved 2018-01-11)

> **lemmatization**
> process of determining the **lemma** (2.12) for a given **word form** (2.24) in a context

The current system has some restrictions and limitations that we enumerate in what follows.

- Cross-references are restricted in scope to the current standard only; there is no cross-referencing across standards, even within the same scientific committee.
- The relation type is not specified explicitly and is, therefore, subject to human interpretation. We observe different types of relations: hyperonymy as in [1], associative relation as in [2], and meronymy as in [3]:

  segmentation annotation
  annotation (2.3) [1] that delimits linguistic elements that appear in the **primary data** (2.1) [2]

  sentence
  related group of **word forms** (3.24) [3] containing a predication, usually expressing a complete thought and forming the basic unit of discourse structure

- Some obvious cross references are neither highlighted nor hyperlinked, as in [4] and [5]:

  graph
  set of nodes [4] (vertices) V(G) and a set of edges [5] E(G)

- By analogy to other terminological resources, a list of relations (Chiocchetti, Culy and Ralli, 2015), or even an explicit relation typing (Lang, Schwinn and Suchowolec, in press) could be provided in a dedicated section of an entry. However, as mentioned above, this information is missing.

To sum up, conceptual relations in the terminological entries of the OBP are rather ad-hoc, locally limited to a given standard, and sometimes incomplete. Within the current display options, it is impossible to grasp the conceptual *structure* of a set of standards or even of a single standard.

4. Project Description

The rationale for our project rests on the following two assumptions of standard-related terminology work (see section 2):

- that the conceptual structure of a domain is implied in definitions
- that visualization techniques make conceptual structures more transparent and thus more beneficial.

We aim at making the conceptual structure of the terminological resources in the OBP explicit, by retrieving all concepts mentioned in the definitions regardless of their highlighting and then visualizing the relations as an interactive concept graph. The work is done with two user groups in mind, and thus encompasses two kinds of aims and questions:

- For end users – researchers in HLT and related areas:
  - to add a visual browsing option for the resource,
  - to give an assessment of the importance of concepts within a certain field, also for educational and training reasons,
  - to answer questions such as “Which concepts are mentioned by many definitions and are, therefore, central and need to be studied first?” or “Is there a conceptual overlap between standards?”

[3] All examples were retrieved from the OBP on January 08, 2018. The numbers in square brackets are introduced by the present authors.

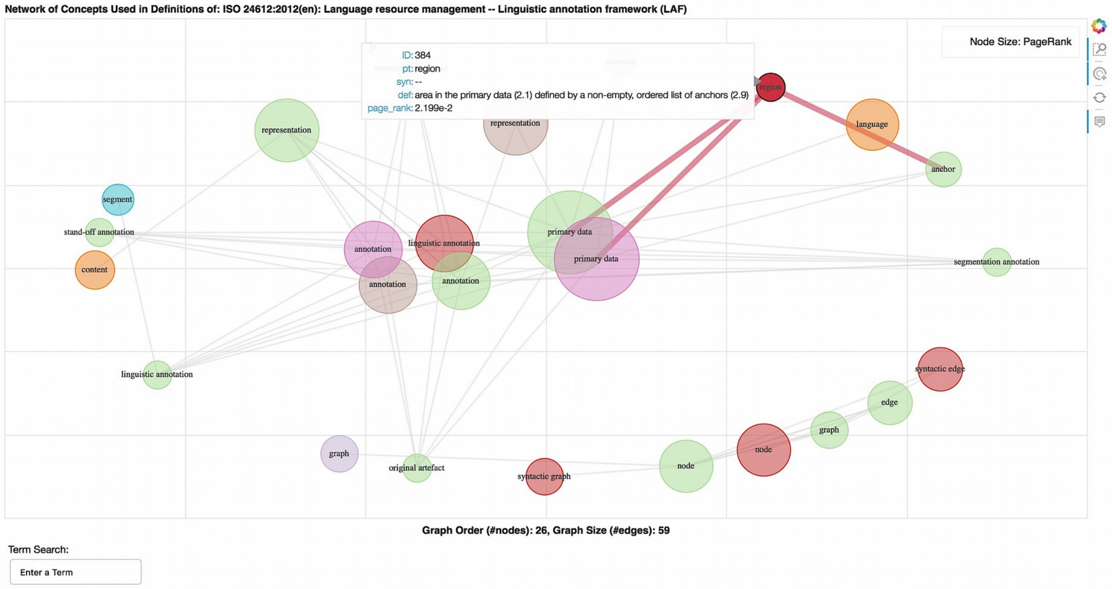

Figure 2: Bokeh Visualisation of ISO 24612:2012 Concepts and their Links to other Standards; hover tool on *region* for concept information and edges of related concepts. In order to give a visual hint of the importance of a node in a graph, we plot the size of the nodes according to an importance measure, here: PageRank.

- For terminology authors, as a diagnostic and evaluation tool providing answers to questions such as:
  - Is the conceptual structure of the given standard well-formed?
  - Is highlighting or hyperlinking missing?
  - Are cross-references correct and complete?
  - Are there entire concepts missing?
  - Are any definitions circular?
  - Are there doublets or overlapping definitions across standards?

The data used in the pilot project come from the ISO Online Browsing Platform and is copyrighted by ISO. Following advice from the CLARIN Legal Helpdesk, the data have been used only internally within the project, with the results planned to be displayed using HTML frames (in this way yielding all control over the data to ISO). The snippets of text mentioned in this paper and in the illustrations therein fulfill the EU copyright law conditions on research exception.

4.1 Data Preparation

Data was downloaded in the form of HTML (which most probably had been exported from TBX, TermBase eXchange, which is the standard representation and exchange format in the terminology community (ISO 30042, 2008)). We have searched for “full entry” in the facet “TC 37 SC 4”, with the option “Terms & Definitions” set. The result was exactly 400 concepts (definition entries) from 20 published standards. The resulting HTML was exploited to harvest Terms, Definitions, Standard, Paragraph and Notes; we assumed that there is always a preferred term (altogether, 405 terms were identified); we used the Python library *lxml* with the ElementTree API (*etree*) to process the HTML tree; the output was a CSV file. We created a relational database for storing the harvested terminology; the database is normalized (3rd normal form) and implemented in MySQL.

4.2 Technical Implementation

This section provides an overview of tools used in the project and an outline of the technical implementation.

4.2.1 Tools

Preprocessing, DB management and visualisation was done in Python 3.6; next to the above mentioned lxml.etree for HTML processing, we used re (for regular expressions), pymysql (for connecting to the MySQL database and reading our terminology database), networkx (for creating basic initial graph and computing measures such as centrality, etc.), and bokeh (v. 0.12.11) and its sub-libraries (for visualising the networkx graph as an interactive HTML file); we used bokeh's hover tool for creating interactive tool tips when moving over the elements of the graph.

4.2.2 Extraction of Links Between the Concepts

The general idea was to check which ISO concepts are mentioned in a given ISO definition: if concept A is mentioned in the definition of concept B, then concept B points to concept A and a directed link (concept B → concept A) is created. All concepts form the nodes and all links form the edges of a graph. Such a graph is directed and records the links between concepts based on concept definitions.

We used Python NLTK library for linguistic pre-processing (WordNetLemmatizer with the default POS “n”). This step is necessary, because definitions are the

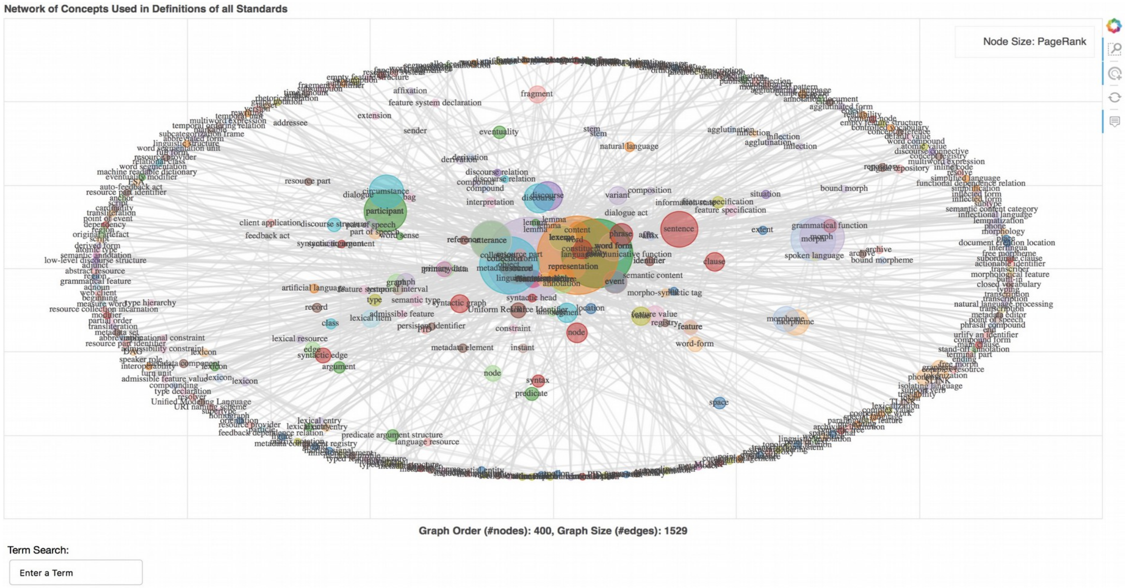

Figure 3: Bokeh Visualisation of the Entire Concept Network of ISO TC37SC4. Node Size: PageRank

only source of information on the conceptual structure in our terminological resource and we test them against a list of all possible terms.

For every concept, we recorded the identified terms. Then, we used the information from the MySQL database to retrieve the corresponding IDs of concepts that use the identified terms. As a result, we obtained a data structure that records links between concept IDs.

In order to maximize automation, in cases where terms are ambiguous, we created links to all possible concepts. In particular, we also linked to concepts where the term in question is assigned as a non-preferred term.

4.2.3 Creation of a Graph Object

We used Python library networkx (v. 1.11) to create a graph object. First, for each concept (ID) from the database we created a node; then we added the edges using the data structure as described in the previous section.

4.2.4 Visualisation of the Graph as an Interactive HTML File

In order to visualise the graph object in an interactive HTML file with the bokeh Python library, we combined the approach of Meier (2016) with bokeh's method *from_networkx*, also customizing the ColumnDataSource by adding columns for concept definitions, preferred terms, synonyms and the source standard, and for the value of the centrality measure (or PageRank) for the size of the nodes. This is done on the fly by consulting the MySQL database for information on concepts and by computing the values of centrality measures with networkx. We have also colour-coded the nodes in such a way that each ISO standard uses a different colour.

We have used the customized columns of the ColumnDataSource in bokeh's hover tool. The hover tool provides interactive tool tips when moving the mouse over an element of the graph. In our case, when pointing to a graph node, the concept ID, definition, preferred term, synonymous terms as well as the source standard are displayed. Moreover, we make use of hover's inspection and selection policy for highlighting nodes and edges of related concepts. Finally, we added a JavaScript callback for term search.

5. Intermediate Results and Open Issues

The outcome of the project thus far is twofold. On the one hand, it comprises concept networks dynamically generated from a relational database. On the other, the project has yielded significant feedback that can be offered to ISO.

5.1.1 Concept Networks

We have been able to dynamically generate different types of interactive graphs, varying with respect to two major parameters: scope of the domain and importance of the node.

Three different scopes of graphs have been tested:

- definitions in *one* particular standard have been searched for concepts defined in the *same* standard
- definitions in *one* particular standard have been searched for concepts defined in *all* standards (see figure 2)
- definitions from *all* standards have been searched for concepts defined in *all* standards (see figure 3).

Recall that referencing concepts across related standards is justified by the ISO recommendations on standard drafting described in section 2.

Once the scope is set, we plot the graph according to the chosen importance measure (centrality measure or PageRank, as described above). At this stage of the

project, the user creates the visualization in the command line. She can either call a function that, first, lists IDs for standards as well as importance measures and, then, prompts for input; or the user can call a different function and pass these IDs directly as arguments.
The resulting graph enables visual and interactive exploration of the ISO concept landscape: by clicking on graph nodes, information on a given concept (definition, used terms, standard, importance value) can be accessed. In addition, the visualisation can be zoomed in and out. Although the underlying networkx graph is directed, this directionality has not been visualized yet.

5.1.2 Qualitative Feedback to ISO

The pilot project has already produced several kinds of feedback that can be offered to ISO.
First, the evaluation of our approach provides immediate information on missing cross references that should be added to the definitions. Moreover, in order to handle issues reported in section 6.2 below, we have implemented a visualisation of explicit cross references that highlights edges which might require a closer examination: red indicates a mismatch between the term and the paragraph number, green points at the use of an unknown synonym, and blue represents a non-exact match. (see figure 4). Further, the interactive concept networks resulting from our automatic approach can be used by editors of standards to see the textually coded concept relations at a glance, thus helping to identify ill-formed structures.
Finally, processing the content of the OBP has generated several kinds of feedback that can be used for qualitative improvement of ISO terminology. We enumerate the main types of issues below.
Typos/malformedness. In addition to infrequent regular typos in terms and definitions, we have found instances of term merger, e.g. "orientation(al) relation" or trailing punctuation, e.g. "annotate,".
Database Export. We have found instances of HTML tagging that we attribute to faulty database export, e.g. the additional POS tag "noun" in "annotation, noun" is treated as a term, in effect becoming a synonym of "annotation" (similarly with "annotate, verb").
Inconsistency. A quick examination of terms used in definitions revealed some inconsistencies in the treatment of the normative status (i.e. preferred term=PT, admitted term, "bare" term) in the corpus of definitions. Four cases can be identified:

- PT is chosen, but a non-PT is used in a definition
- no normative status is assigned, all terms are treated as equal, "bare" terms (possibly, this is a result of faulty database export); this causes inconsistencies in definitions, because there is no guidance on which term to use as reference
- for the same set of terms, different standards choose different PTs
- it is not always clear whether a term is used as an instance of a particular concept or with a more general meaning (e.g. "content" or "text")
- moreover, multiples of minimally modified definitions can be found across standards (sometimes, the duplication is signalled in the "Source" section, but sometimes this information is missing).

6. Evaluation

We evaluated our approach to automatically retrieving conceptual relations in terms of precision and retrieval ratio by comparing our concept networks of the smallest scope (one standard only) to the explicit cross references in the definitions of a standard as described in section 3.2. We do not use *recall* due to the lack of a gold standard; because the manually annotated cross references are incomplete (see section 3.2), we cannot regard them as a gold standard. The retrieval ratio shows then how well our fully automatic approach performs in comparison to the manual annotation.
First, we generated graph objects for explicit cross references. In order to create graph edges, we exploited HTML markup in combination with the parenthesised information on the paragraph number. Again, we used Python libraries *lxml, re, NLTK* and *networkx.*
Then, for each standard, we compared the edges retrieved by our automatic approach with the edges retrieved from explicit cross references. See Table 1 for the number of edges for each of the 20 standards by TC 37 SC 4.
As for precision, we looked at the number of edges retrieved by the automatic approach that are missing in the explicit cross references. However, because of the above-mentioned incompleteness of the cross references, a human evaluator within the project had to decide whether those edges were valid or invalid candidates. The precision is, then, the number of valid candidates to the total number of edges retrieved by our automatic approach.

$$precision = e\_valid_{auto} / e\_total_{auto}$$

As for retrieval ratio, we look at the number of edges retrieved from explicit cross references that are missing in our automatic approach. The retrieval ratio is, then, the ratio of the number of edges not missing in the automatic approach to the total number of edges from explicit cross references.

$$retrieval\ ratio = e\_not\_missing_{auto} / e\_total_{expl}$$

In the following, we discuss the overall precision and retrieval ratio for all standards in total.

	#nodes	*#edges*			*#nodes*	*#edges*	
		auto	*expl*			*auto*	*expl*
1	20	22	21	**2**	28	38	34
3	22	22	20	**4**	25	17	0
5	15	24	19	**6**	14	22	16
7	24	42	40	**8**	38	63	47
9	5	6	3	**10**	47	89	91
11	32	43	36	**12**	20	21	21
13	4	4	4	**14**	31	69	55
15	4	0	0	**16**	22	49	45
17	23	34	31	**18**	9	7	0
19	9	14	0	**20**	8	2	2
				∑ auto = 588		∑ expl = 485	

Table 1: Number of edges retrieved by the automatic approach ("auto") and from explicit cross references ("expl") in 20 standards of TC 37 SC 4

6.1 Retrieval Ratio

There are 19 edges in the explicit cross references that are not retrieved by our automatic method (i.e. 466 edges are

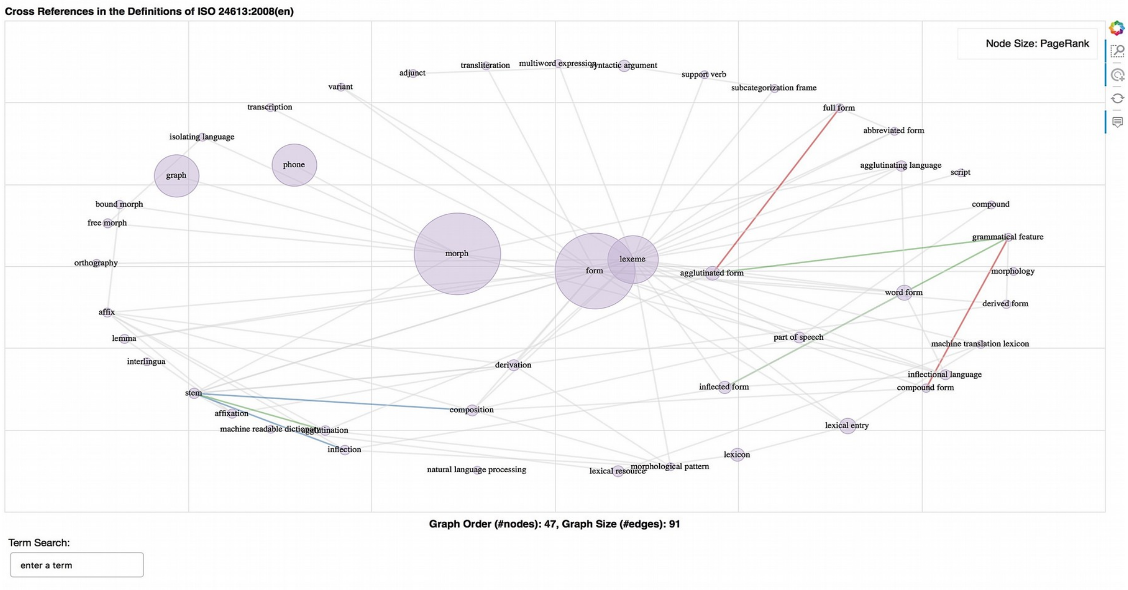

Figure 4: Bokeh Visualisation of ISO 24613:2008 as a Diagnostic Tool. Edges for revision are colour-marked: red: mismatch between the term and the paragraph number, green: unknown synonym used, blue: non-exact match.

"not missing"). 16 edges got missed because of the wording in the definitions: using unknown synonyms (e.g. "time unit" instead of "temporal unit" or "temporal interval"), referencing non-exact matches (e.g. "agglutinative" instead of "agglutination") or elliptical constructions ("compound or derived form" instead of "compound form", "derived form"). In two cases, there is a mismatch between the term used in the definition and the paragraph number ("URI (3.1.1)" instead of "URI (3.2.2)". Finally, one case involves the use of an inflected verbal concept ("annotated"), which was missed by our lemmatizer that was set for detecting nouns only. The resulting retrieval ratio is 0.96.

6.2 Precision

There are 122 edges in our automatic method that are not retrieved from explicit cross references. 86 of these edges were considered valid candidates by a human evaluator.[4] Out of the remaining 36 edges, 24 need a closer evaluation by a domain expert and 12 are invalid candidates, mostly because of the above-mentioned lemmatizer setting. Considering these 36 edges only, our precision reaches 0.94.

7. Future Work

We judge the results of the pilot project as very promising for our intended user groups, while naturally noting some issues that deserve further attention.

Already at this point, the project supplies new ways of both browsing the concept network (for the user) and of increasing its coherence (for the maintainer).

The dynamically generated interactive graphs offer additional visual access to the OBP, complementing its traditional list view. They can be used to explore the textually coded concept structure, also allowing to spot central concepts at a glance. We believe that this facilitates a better understanding of the structure of the domain, although empirical evidence is still needed.

Terminology authors benefit from this kind of visual access as well, because it helps to spot issues in definition texts that are rather difficult to find using the list display.

Already at this stage of the project, it is clear that even the simplest graph-theoretical characteristics such as graph size and graph order provide a valuable contribution to terminology management and help to improve the quality of the resource. In our future work, we therefore intend to exploit other graph-theoretical characteristics for the target user groups. In the longer perspective and with the scientific community in mind, we intend to perform exploratory network analysis in the sense of Igual and Seguí (2017) in order to gain more insight into the nature of the data at hand.

In further future work, we are going to implement more evaluation metrics for our fully automatised approach with larger scopes. At this point, we anticipate the need to address the general issue of how to deal with term ambiguity and underspecification, which, on one hand, is a notorious problem in terminology management, but might be, on the other hand, limited by the specific features of definitions in standards. The latter might be remedied with the help of user feedback.

[4] Note that this high number is partially due to the fact that some standards do not indicate any cross references at all, see standard number 4, 18 and 19 in Table 1.

Furthermore, we are going to deal with the negative impact of duplicate definitions on the readability of graph visualisations. At present, duplicate concepts are plotted as separate nodes, see figure 2, containing among others a cluster of 3 *annotation* nodes. We attribute this both to the source data (*ad hoc* modifications of definition bodies) and to the current visualisation solution. Further improvements will, therefore, deal with the link clutter of the larger-scope graphs as well as with a suitable directionality visualization for the edges (cf. Holten et al., 2011).

Next steps will also address the need to extract/add some hierarchical structure to the rather flat concept networks. We plan to look at pattern-based approaches such as Arnold and Rahm (2014), Sierra et al. (2008) or Storrer and Wellinghoff (2006), and to explore the options offered by graph-theoretical methods.

Respecting the ISO copyright restrictions, we intend to offer framed access to concept graphs to researchers from the CLARIN service offered by IDS Mannheim.

8. Acknowledgements

We would like to thank the anonymous reviewers for their detailed comments and valuable suggestions.

9. Bibliographical References

Arnold, P. and Rahm, E. (2014). Extracting Semantic Concept Relations from Wikipedia. In proceedings of WIMS '14, Thessaloniki, Greece, no pages, preprint: https://dbs.uni-leipzig.de/file/wims2014-final.pdf.

Chiocchetti, C., Culy, C. and Ralli, N. (2015): Visualising conceptual relations in the domain of law: Verweis Viewer. In Roche, C. (Ed.): TOTh 2013. Actes de la septième conférence TOTh – Chambéry – 6 & 7 juin 2013. Annecy: Institut Porphyre, Savoir et Connaissance, pages 135–154.

DIN 2331:1980-04 (1980). Begriffssysteme und ihre Darstellung. April 1980. Berlin: Beuth.

Drewer P., Massion F. and Pulitano, D. (2017). Was haben Wissensmodellierung, Wissensstruktur, künstliche Intelligenz und Terminologie miteinander zu tun? Deutsches Institut für Terminologie e.V., pages 1–28 http://dttev.org/images/img/abbildungen/DITeV_org_Terminologie_und_KI_2017_03_22_v2.pdf.

Faber, P., León-Araúz P. and Reimerink, A. (2016). EcoLexicon: New Features and Challenges. In Proceedings of GLOBALEX 2016 – Lexicographic Resources for Human Language Technology Workshop, pages 73–80, Portorož, Slovenia.

Falke, S., Lang, Ch. and Suchowolec, K. (2017). Entwicklung eines graphentheoretischen Terminologie-Analysetools. Talk given at ars grammatica: Grammatische Terminologie – Inhalte und Methoden, June 2017, IDS Mannheim, Germany.

Früh B. and Deubzer F. (2016). Von der Terminologieverwaltung zur Wissensorganisation. *Edition* 16(1):27–32.

Holten, D., Isenberg P., van Wijk, J. J. and Fekete, J.-D. (2011). An extended evaluation of the readability of tapered, animated, and textured directed-edge representations in node-link graphs. In Di Battista, G., Fekete, J.-D. and Huamin, Q. (Eds.): IEEE Pacific Visualization Symposium 2011. Proceedings, pages 195–202.

Igual, L. and Seguí, S. (2017). Network Analysis. In L. Igual & S. Seguí (Eds.), *Introduction to Data Science. A Python Approach to Concepts, Techniques and Applications*. Cham: Springer, pages 141–164.

ISO 704:2009-11 (2009). Terminology work – Principles and methods. November 2009.

ISO 10241-1:2011 (2011). Terminological entries in standards – Part 1: General requirements and examples of presentation. April 2011.

ISO 24612:2012 (2012). Language resource management – Linguistic annotation framework (LAF). June 2012.

ISO 30042:2008 (2008). Systems to manage terminology, knowledge and content – TermBase eXchange TBX. First edition

Lang, Ch., Schwinn, H. and Suchowolec, K. (in press). Grammatische Terminologie am IDS – Ein terminologisches Online-Wörterbuch als ein vernetztes Begriffssystem. In *Sprachreport* 34(1):16–26.

Meier, B. (2016). NetworkX Visualization Powered by Bokeh. Talk given at EuroPython 2016. video: https://av.tib.eu/media/21112; jupyter notebook: https://ep2016.europython.eu/media/conference/slides/networkx-visualization-powered-by-bokeh.ipynb, Bilbao, Spain.

Sierra, G., Alarcón, R., Aguilar, C. and Bach, C. (2008). Definitional verbal patterns for semantic relation extraction. *Terminology* 14(1):74–98.

Storrer, A. and Wellinghoff, S. (2006). Automated detection and annotation of term definitions in German text corpora. In Proceedings of LREC 2006, pages 2373–2376 http://www.lrec-conf.org/proceedings/lrec2006/pdf/128_pdf.pdf.