

Alignment Cubes: Towards Interactive Visual Exploration and Evaluation of Multiple Ontology Alignments

Valentina Ivanova¹, Benjamin Bach², Emmanuel Pietriga³, and Patrick Lambrix¹

¹ Swedish e-Science Research Centre and Linköping University, Sweden

² University of Edinburgh, United Kingdom

³ INRIA, LRI (Univ Paris-Sud & CNRS), Université Paris-Saclay, France

Abstract. Ontology alignment is an area of active research where many algorithms and approaches are being developed. Their performance is usually evaluated by comparing the produced alignments to a reference alignment in terms of precision, recall and F-measure. These measures, however, only provide an overall assessment of the quality of the alignments, but do not reveal differences and commonalities between alignments at a finer-grained level such as, e.g., regions or individual mappings. Furthermore, reference alignments are often unavailable, which makes the comparative exploration of alignments at different levels of granularity even more important. Making such comparisons efficient calls for a “human-in-the-loop” approach, best supported through interactive visual representations of alignments. Our approach extends a recent tool, Matrix Cubes, used for visualizing dense dynamic networks. We first identify use cases for ontology alignment evaluation that can benefit from interactive visualization, and then detail how our *Alignment Cubes* support interactive exploration of multiple ontology alignments. We demonstrate the usefulness of Alignment Cubes by describing visual exploration scenarios, showing how Alignment Cubes support common tasks identified in the use cases.

Keywords: Ontology alignment evaluation, Visual exploration, Multiple alignment comparison

1 Introduction

The need for automatic alignment of ontologies has sparked the development of a growing number of tools and algorithms. A comprehensive literature review can be found in [18]. It has also led to the creation of an annual event, the Ontology Alignment Evaluation Initiative (OAEI)⁴, where alignments computed by the participating tools are compared against *reference alignments* (RA). In most cases the quality of these alignments is measured in terms of precision, recall, and F-measure. Precision is the ratio of correct suggested mappings over all suggested mappings. Recall is the ratio of correct suggested mappings over all correct mappings. F-measure is a harmonic mean between precision and recall. These measures give a good overall assessment of the quality of

⁴ <http://oaei.ontologymatching.org>

alignments in terms of the ratio of found mappings, missed mappings and wrongly suggested mappings. However, they do not allow for comparing alignments of specific parts of ontologies, or for comparing alignments to each other and to the RA at the detailed level of concepts and relations. Without means to compare the tools and algorithms at a detailed level, their strengths and weaknesses cannot be easily revealed and understood.

Furthermore, RAs are often not available, as their development is time and effort consuming and requires domain expertise. As a consequence, the quality of alignments is difficult to measure. In the absence of RA, the evaluation of alignments requires the exploration and comparison of multiple alignments, which involves expert users (analysts) performing tasks at different levels of granularity [1,8,20]: determining regions with similar or different numbers of mappings between the alignments, determining common or rarely found mappings, characterizing mappings as correct or incorrect. However, there is currently little support for performing these tasks in an interactive and flexible manner. Analysts are relying on custom scripts, which can be error-prone and require time to develop and fine-tune.

This work presents the following contributions to the ontology alignment field: (i) we identify several use cases that would benefit from comparative assessment of several alignments at different level of detail; we discuss their shared analytical tasks and identify features that would benefit from visual support. (ii) To address these use cases and tasks, we propose an interactive visual environment for the simultaneous comparative exploration and evaluation of multiple alignments at different levels of granularity. While visualizing even a single alignment is still a challenge due to the size and complexity of the ontologies, we provide a compact way to visualize multiple alignments. Instead of depicting all mappings together in a single representation, which would cause visual clutter and information overload, we provide an interactive visualization that supports multiple complementary views, and overview and detail techniques to explore alignments at different levels of granularity. We interpret an alignment as a bipartite graph (bi-graph), i.e., a network where links exist only between nodes of different sets (ontologies), and draw from the literature in the field of network visualization (e.g., [3,22]). We identify Matrix Cubes [2], a novel technique introduced for the interactive visual exploration of dynamic networks, as a promising visual approach to serve as a foundation for our tool - *Alignment Cubes*. Alignment Cubes significantly extend Matrix Cubes in order to make it applicable to the visualization of multi-level ontology alignment networks in the form of bi-graphs.

This article is structured as follows: Section 2 presents the use cases and their shared tasks, and discusses existing approaches for visualizing multiple alignments. Section 3 describes Matrix Cubes, and Section 4 describes the Alignment Cubes we derived from them. We explain how users interact with Alignment Cubes using an example scenario in Section 5, discuss lessons learned and future extensions in Section 6, and conclude in Section 7.

2 Ontology Alignment Evaluation

The interactive exploration and evaluation of several alignments is only rarely considered in the literature. In [1] several analytic tasks have been identified and supported

through multiple connected views, while [15,19,20] focus more on the alignment computation than their presentation. We thus first identify several evaluation use cases and discuss shared activities that could be efficiently supported through interactive visualization in Subsection 2.1. In Subsection 2.2 we connect this discussion with work about cognitive support for ontology mapping [9], requirements and evaluation of user interfaces in ontology alignment systems [7,12], and review capabilities for simultaneous visualization of several alignments in existing systems. In Subsection 2.3, we study the visualization approaches taken in alignment evaluation frameworks.

2.1 Evaluation Use Cases

The following use cases would benefit from users being able to simultaneously explore and evaluate several alignments interactively:

- (UC1) **Selecting, combining and fine tuning alignment algorithms and tools:** OAEI editions, and [8,21] among others, have shown that matchers and tools do not necessarily find the same correct mappings, and may compute different erroneous mappings too. Thus, selecting and combining matchers and tools requires examination of overlapping and divergent mappings to understand the differences in the underlying algorithms. It also includes assessing the impact of parameter changes on single similarity values and parts of the alignments for fine-tuning.
- (UC2) **Matchers development:** Developers alter their algorithms according to some observations. Examining the outcome of such changes over other parts of the alignment and comparing to previous versions at different levels of granularity helps them assess the consequences of these changes.
- (UC3) **Ontology alignment evolution:** Alignments are used for, e.g., data integration, merging ontologies, and database annotation. Changes in the alignments may influence the applications employing them. Understanding how alignments differ will facilitate the assessment of the impact of changes on their client applications [14,23].
- (UC4) **Validating and debugging of ontology alignments and RA:** Analyzing several alignments at the same time may reveal parts with large variations in the number of mappings or similarity values, and help in identifying potential errors and their sources during diagnosis. Developing and debugging RAs is a laborious and error-prone task. Recently, several works have found problems in the OAEI Anatomy track's RA [8]. Manual mapping validation requires a detailed view of each mapping and its context [9]. Understanding consequences of user validations and exploring what-if scenarios are also important [12].
- (UC5) **Collaborative ontology alignment:** Collaborators need to understand the current state of an alignment, where and why their peers have introduced changes in comparison to previous revisions. This is especially needed when collaborative work happens over a long time period, or in distributed teams.

These use cases share common analytic tasks at *different granularity levels*; identifying parts of the alignment covered or not covered by all alignments, agreement or disagreement between matchers, determining incorrect, missed or always found mappings—which could be efficiently supported through interactive visualizations. Without measures to provide well-defined quantitative outputs, and with only high-level goals and

questions, e.g., which threshold to choose and why (UC1), which version of my algorithm is better and why (UC2), what has been changed between two revisions and why (UC3,UC5), all of the use cases above are exploratory in their nature. They aim to find relevant information and answer questions not known in advance.

These are typical scenarios in which visual and interactive representations can be of help. In a visual environment exploration activities are supported by interactively varying parameters and thresholds to estimate the impact of changes, changing visual encodings to highlight different aspects, and reordering elements to facilitate trend and pattern discovery. Interactive exploration at different granularity levels reveals *regions of interest* to guide further exploration and help in identifying *patterns in and among different regions of interest* which may reveal similarities and dissimilarities between the respective alignments. These benefits directly apply to our case, where analysts are faced with the problem of exploring and visualizing multiple alignments.

Our use cases UC1-UC5 involve numerous *compare and contrast tasks* to select a threshold (UC1), an alignment to use (UC2 and UC4), for diagnosis and to identify outliers (UC4), common trends and regions (UC1, UC2), and for identifying changes (UC3, UC5). According to [1], diagnosis is a complex activity which is composed of iterative sequences of exploration and comparison, analyses of clusters of mappings, and comparative evaluation of matchers' performance.

2.2 Existing Approaches

Two recent studies review exploratory features in the user interfaces of ontology alignment systems [7,12]. Only one of the tools [5] (including its recent extensions [1,17]) provides support for visualization and exploration of several alignments together after discussing the need for such in connection to analytic tasks identified in interviews with alignment experts. Navigation and exploration in ontologies and just a single alignment have been highlighted in the context of discovering mappings and verification in [9], through the *inspection* dimension defined in [12], and by the 7 information seeking tasks and visual analytics from [7]. Another desirable functionality identified by alignment experts in [1] is the clustering of mappings according to different statistics in order to analyze each cluster separately. Several other works have also emphasized the importance of visually identifying dense regions in single alignments: for planning manual validation and identifying the most similar areas of the ontologies [9], as part of alignment inspection [12], and grouping to help identifying patterns [7]. Small-world graphs were used to present clusters in a single alignment at different granularity levels [16].

The exploration and comparison task functionality identified by experts in [1], is supported through comparative views by showing juxtaposed matrices computed by a combination of different matchers in the system. An additional matrix highlights variations in computed mappings between the matchers. RAs can be overlaid on top of each matcher's matrix. This approach is similar to one of the views in our approach using Matrix Cubes [2]—small multiples—but in comparison to our work, it forgoes the structure of the ontologies which is shown in another view. Similar to our approach, previous work [6,13] has presented the structure of ontologies as indented trees on the sides of a matrix, but only focused on a single alignment. One of them [6] provides support for reordering and depicts several types of mappings as well as derived and

asserted mappings. Recent work [20] takes another approach to simultaneously visualizing several (externally-generated) alignments, employing linked indented lists and color-coding to depict the edges belonging to different alignments. The authors emphasize that with their tool “[...] users can better compare and analyze alignments (i.e., parts of the ontologies which are covered for most alignments and those which are not, consensus between alignments, etc.)”. Although filtering by threshold and mapping type is supported, the view quickly becomes cluttered as the size of ontologies and the number of alignments involved grow.

2.3 Frameworks for Ontology Alignment Evaluation

A few ontology alignment evaluation frameworks, SEALS⁵, KitAMO [15] and AMC [19], provide rich back-end infrastructures for configuring (to a different extent), executing and storing the results from the execution of alignment components (matchers, filters and combination algorithms). In SEALS, used in OAEI, the alignments computed by the tools are compared to RA and evaluated in terms of precision, recall, F-measure, run time, coherence and number of requests to an oracle. The results are presented in sortable tables, after analysis with custom scripts. In KitAMO [15], probably the earliest system, the results are presented in the form of sortable tables containing either the mappings with their similarity values computed by each component, or aggregated data (number of correct, incorrect and inferred mappings) in comparison to another component. In AMC [19], linked indented lists are used together with sortable tables as well. It additionally introduces a cube view which presents a single alignment where two of the dimensions depict the source and target ontologies, and the third dimension shows similarity values as bars—taller bars representing higher similarity values.

While all three frameworks provide rich back-end infrastructures for configuring and executing alignment algorithms, the tabular views are too limited to adequately address the simultaneous interactive comparison of several alignments and provide visual exploration at different granularity levels. The third framework devotes more attention to the visual presentation of the results, but it only depicts one alignment at a time. In comparison, our work focuses on the user interface and allows users to visually explore multiple alignments together. It can thus be seen as complementary to the back-end functionalities offered by these frameworks.

3 Matrix Cubes and the Cubix Interface

Alignment Cubes are adapted from earlier work that introduced *Matrix Cubes* [2] as an interactive visualization metaphor for the comparison of multiple time steps in dynamic networks. Here, we briefly summarize Matrix Cubes and *Cubix*, the interactive interface prototype used to manipulate and explore Matrix Cubes, and refer readers to [2].

Dynamic networks are networks whose topology and attributes change over time. The general idea of Matrix Cubes is to represent the different states (time steps) of a dynamic network as adjacency matrices, one matrix showing the state of the network

⁵ <http://seals-project.eu> — Semantic Evaluation At Large Scale




Fig. 1: From dynamic networks to Matrix Cubes and Alignment Cubes. (a) Adjacency matrices stacked to form a (b) 3-dimensional space-time cube (Matrix Cube). (c) Translation of the concept of Matrix Cubes to ontology alignment.

at one given point in time. Each adjacency matrix is organized as follows (fig. 1-a): nodes are the row and column headers, and cells represent the links between nodes. A cell at the intersection of a given row and column will only be filled if there is a link between the two corresponding nodes. Attributes of such links can be encoded visually inside the matrix cells', using variables such as, e.g., color hue or saturation, texture, and glyph size.

All adjacency matrices corresponding to the individual states of the network over time are then stacked, forming a cube (fig. 1-b). The resulting 3D visualization acts mainly as a pivot and metaphor. Users manipulate the cube interactively to derive multiple meaningful 2D projections of the cube or its content, better suited to their visual analysis tasks. Manipulations include slicing the cube along different dimensions, rotating slices or the entire cube, juxtaposing slices to obtain small-multiple views, playing with cells' transparency to enable detailed compare & contrast tasks between slices.

As shown in fig. 1-b, Matrix Cubes hold the network's nodes along two of the cube's three dimensions, the last dimension representing time, i.e., the different states of the network over time. As detailed below, instead of representing time on the cube's third dimension, our Alignment Cubes represent the different alignments being compared thus opening analysis opportunities for other dimensions than time.

An interactive visualization environment for Matrix Cubes, called Cubix, has been developed in Java + OpenGL by the authors of [2]. Our prototype Alignment Cube visualization tool is derived from the Cubix implementation. While Matrix Cubes are the general visualization structure (i.e., the 3D cube consisting of matrices) and Cubix is the interactive interface, the term *Alignment Cubes* refers to both; a specialization of Matrix Cubes *and* the name of our interface.

4 Alignment Cubes

This section details our extensions and adaptations of Matrix Cubes and describes our tool, Alignment Cubes⁶, in connection to the discussion in Section 2.

⁶ <http://www.ida.liu.se/~patla00/publications/ISWC17> provides supplemental material to this submission: all figures from the paper in higher resolution, a screencast of the tool, and a downloadable version of the tool itself.

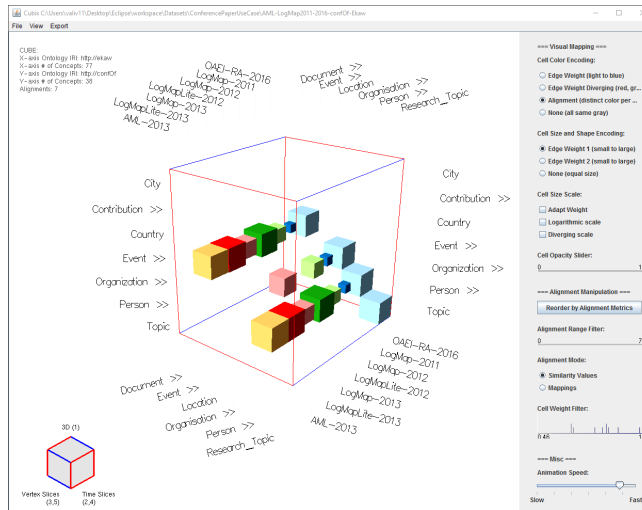


Fig. 2: Default view of Alignment Cubes in *similarities* mode. Rows and columns of the cube represent ontology concepts, individual cells inside the cube represent mapping relationships. Each alignment, corresponding to a slice in the third dimension, is assigned a different color. Widgets (sliders, button groups, etc.) mentioned in the text are referred to directly using their name.

4.1 From Matrix Cubes to Alignment Cubes

Two ontologies and their alignment can be seen as a bipartite network of mappings between individual concepts in the two ontologies. A matrix represents a single alignment between two ontologies. The rows hold the concepts from one ontology, the columns hold the concepts from the other ontology. Cells denote existing mappings between concepts in the respective rows and columns. Stacking several matrices, i.e., several alignments, creates an Alignment Cube.

The example in fig. 2 shows two of the ontologies from the OAEI Conference track, *ekaw* (columns, 77 concepts) and *confOf* (rows, 38 concepts), as well as seven alignments (laid out along the depth dimension), i.e., the RA for 2016 and alignments from AML and the LogMap-family of systems from 2011 to 2013. Each alignment is color-coded to make it easy to visually differentiate the mappings, by grouping the cells that belong to each of them using a pre-attentive variable. Position is not sufficient to clearly identify which alignment a cell belongs to, as the cube can be manipulated in many ways (rotation, slicing, etc.).

Concepts in ontologies often form a taxonomic hierarchy, but as Matrix Cubes do not support hierarchical networks, we extended the original framework. With Alignment Cubes, we represent the ontologies as indented lists with collapsible rows and columns. Fig. 2 depicts the first level in both ontologies. Concepts that feature sub-concepts display the >> symbol after the concept label (e.g., *Event*, present in both ontologies, features sub-concepts). Clicking on a concept label expands the corresponding row or column. Expanded concepts then show the > symbol, and sub-concepts are

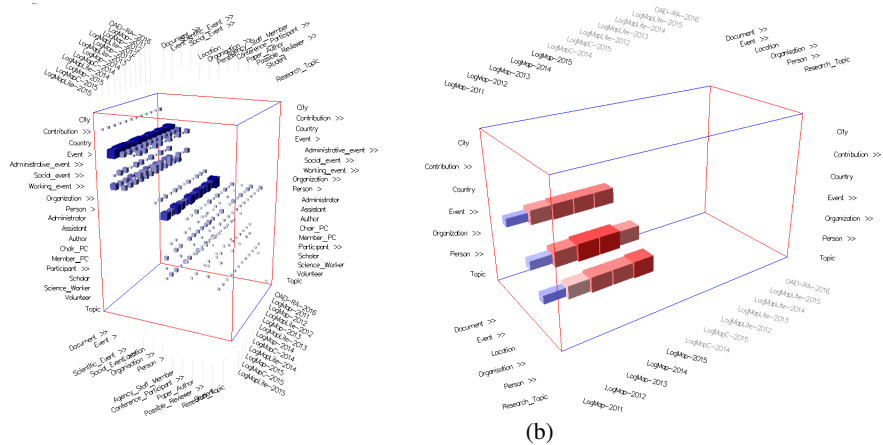


Fig. 3: Views of Alignment Cubes with two different cell weight color encodings: (a) after expanding the *Events* and *Persons* concepts in both ontologies and (b) after filtering out several alignments.

indented according to their level in the hierarchy (fig. 3-a). Concepts that have multiple parents appear under each parent, i.e., potentially multiple times in the hierarchy.

4.2 Granularity Levels

As discussed in Section 2, we aim to support views at different levels of granularity—from an overall view to regions based on the is-a hierarchy, and down to single mappings. To do so we introduce *alignment modes*. In *similarities* mode (fig. 2), a filled cell represents an existing mapping between a pair of concepts. In *mappings* mode (fig. 3-a and 4), a filled cell indicates that there is at least one existing mapping between a pair of concepts or their descendants. The cell weight represents either the similarity value (in the former case), or the number of mappings (in the latter case). Each mode is focused on performing one of two tasks: to compare similarity values for a pair of concepts, and to identify regions in the alignments with few or many mappings. The latter task provides a starting point for exploration and highlights regions of interest where many or few mappings have been calculated. When a concept is expanded in *mappings* mode, a cell is shown for both the concept itself and its sub-concepts. This forms regions in the cube (as in fig. 3-a) where smaller cells indicate mappings deeper in the hierarchy.

4.3 Interactive Visual Exploration

The Alignment Cubes user interface provides a variety of interactions for visual exploration, shown in fig. 2: changing alignment modes (see above), cell color and size encodings, switching between individual views, adapting cell transparency, brushing and linking, as well as alignment slice reordering. The cubelet widget (bottom left corner in fig. 2) allows for a quick navigation between a set of predefined views by clicking

or dragging the mouse on its faces. For example, cells can be filtered out by specifying minimum or maximum value thresholds using a range slider. This also allows to simulate different thresholds and explore what-if questions and cases. Entire alignments can be hidden. To support pattern discovery, the order of alignments (slices) in the cube can be changed to facilitate comparison. A specific order can be calculated based on measures of precision, recall and F-measure between the ontologies, or based on alphanumeric label sorting (labels representing matcher, tool or alignment name). After sorting by label, matchers from one family are displayed together next to each other.

4.4 Compare and Contrast

As with Matrix Cubes, Alignment Cubes provide several views onto the data, resulting from manipulations of the 3D cube. The individual views are: (a) 3D view, (b) 2D projection on 2 of the orthogonal cube faces, (c) side-by-side layout (small-multiples view) of the cube's slices (along 2 of the orthogonal dimensions). The 3D cube provides an overview of the number of alignments, number, size, and distribution of cells (mappings). It helps identify regions of interest and thus drive the initial exploration phase, and can possibly yield some high-level insights (fig. 3-a). It allows for interactive rotation and zoom but suffers from the typical drawbacks of 3D visualization, including occlusion and perspective distortion. Projection views allow for a clutter-free aggregated view on all alignments by orthogonally overlapping cells (fig. 6). Side-by-side views provide the most detailed view onto the data by entirely decomposing the view and showing each alignment in detail (fig. 5). Individual views, together with the ability to vary cell size, color, and translucency, allow for flexible multi-perspective exploration of the entire data set.

Each of the two projections (alignment topology and concepts network) is paired with its respective small-multiples view. Both projections/small-multiple pairs allow for investigating the behavior of matchers—the alignment topology pair focuses on the similarities and differences between the alignments as a whole, while the concepts network pair allows for analyzing the behavior of matchers for a particular concept.

5 Use Cases Support

This section revisits the tasks and use cases identified in Section 2 and demonstrates how our tool supports interactive visual exploration and comparative evaluation of multiple alignments in Subsection 5.1. We further show that Alignment Cubes satisfy requirements for ontology alignment evaluation systems in Subsection 5.2.

5.1 Comparing Alignments & Systems and Support for Comparative Evaluation

We demonstrate how the tasks discussed in Subsection 2.1 can be performed with our tool by conducting a walk through scenario. As the evaluation of interactive visualization tools is challenging [4], methods such as case studies and usage scenario, are often employed [4,10] as they are more likely to provide insightful observations than traditional controlled experiments. In this scenario we aim to answer analytical questions

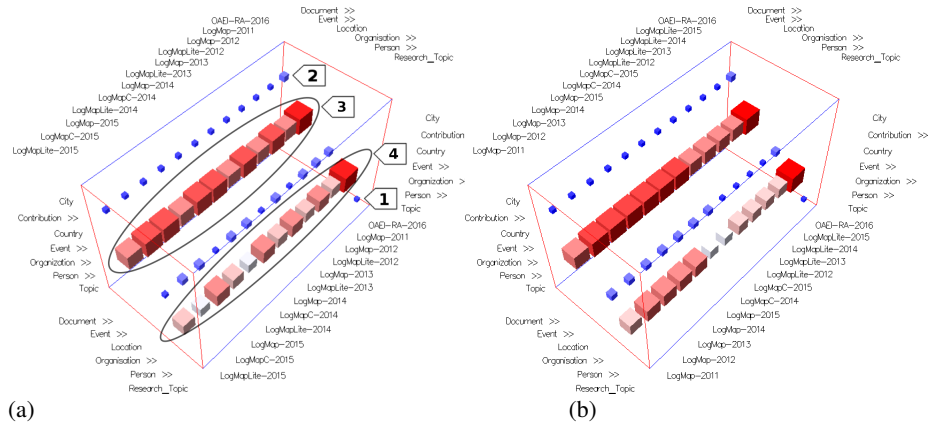


Fig. 4: 3D view in mappings mode: (a) initial view and (b) after reordering by name.

including the following: Are the same regions covered in all alignments? Do matchers agree or disagree? Are there consistently stable and changing regions? Do similarity values differ? Are there missing and wrong mappings? Could we obtain other insights? These analytical questions are shared by all use cases. We focus on observations which would be problematic to obtain without a visual representation of the alignments, as was the case in, e.g., [8].

Dataset: We use the same two ontologies from the example in Section 4, but this time with a different set of alignments. We downloaded the alignments for the LogMap-family of systems between 2011 and 2015, and the RA for 2016. During this period, LogMap contributed three different versions: LogMapLite 2012–2015, LogMapC 2014–2015, and LogMap 2011–2015. Matrices for the alignment (fig. 5 and 6) appear sparse since the respective alignments are the final alignments for the OAEI evaluation campaign, and thresholds have already been applied to remove mappings with a low similarity value. Note that in the process of matcher development, selection and fine tuning (see UC1 and UC2), the matrices would often be denser since more potential mappings would be considered. Alignment Cubes would scale to such higher-density matrices, as the technique was designed for dense dynamic networks in the first place [2].

Exploration: Our exploration starts with the initial 3D cube view in *mappings* mode (fig. 4-a). The colors and sizes of cells show that most of the mappings (large red cells) are concentrated around the two sets of concepts between *Person* (④) and *Event* (③). In contrast, few mappings have been found between *Organization* and *confOf:Contribution-ekaw:Document*. To see the subconcepts of *Person* and *Event*, we expand these two concepts in both ontologies (fig. 3-a). After expanding these concepts, as the visualization is in *mappings* mode, the cells that appear smaller and lighter (because of this particular color scheme) indicate under which subconcepts the single mappings are located. We now collapse these concepts and return to the view in fig. 4-a. We see that all matchers have computed similar numbers for the *Event* concept pair, as the size of cells is more or less constant along the mapping dimension (cube depth).

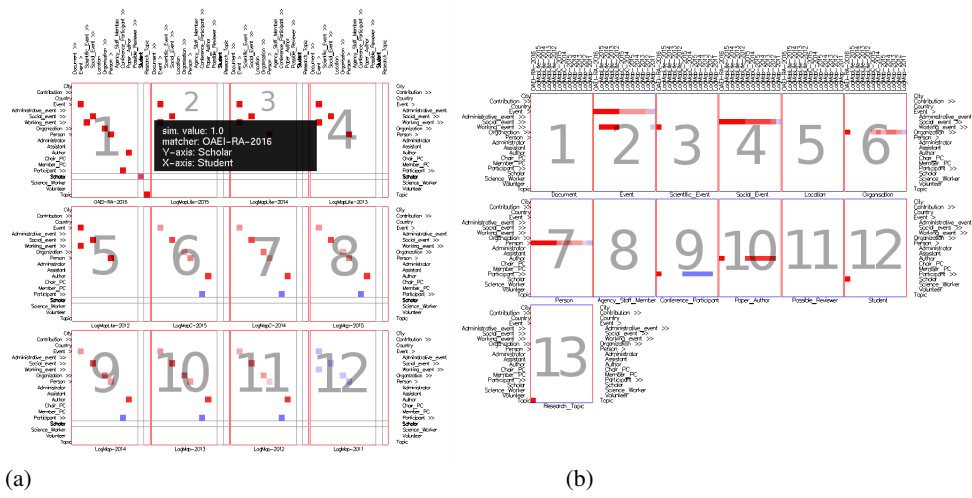


Fig. 5: Small-multiples views: (a) alignment topology (b) concept networks. Gray numbers have been added manually to the figure.

However, somewhat larger differences in cell size can be observed along the depth dimension (across alignments) for the *Person* concept pair, which indicates a lower level of agreement across matching algorithms in this case. Hovering the respective cells with the mouse reveals that numbers vary between 4 and 8. Apparently, only RA (the last, backmost slice) contains a mapping between *confOf:Topic* & *ekaw:Research Topic*, meaning that we found a missing mapping: the single cell not aligned with others (fig. 4-a ①).

For further exploration, we investigate different alignment orderings (by interactively reordering alignment slices). Reorder-by-name (fig. 4-b) clearly separates the three versions of the LogMap-family systems and allows to compare the performance for each version during the years. LogMap performs consistently better in terms of number of mappings found between 2012 and 2015 than in 2011, as visually indicated by the larger and more saturated (red) cells. To see if the same is true for similarity values, we switch to the *similarities* mode: if both modes show a similar picture, the system likely has used the same combination of alignment algorithms. To focus only on the LogMap alignments, we filter out all other alignments in fig. 3-b, leaving the cube half-empty. The similarity values computed by LogMap vary between the years (fig. 3-b), but not for its lightweight version, LogMapLite (not shown due to the filtering).

Focusing on the regions with many mappings, we expand the respective concepts in both ontologies and switch to the small-multiples view that shows alignment topology (fig. 5-a). In both modes the patterns for each of the versions are clearly noticeable. The lightweight LogMapLite in all four years (matrices ②–⑤) consistently finds fewer mappings than the versions of LogMap between 2012 and 2015 (matrices ⑧–⑪). LogMap 2011’s pattern (matrix ⑫) is closer to the LogMapLite alignments than to the rest of the LogMaps. We may conclude that LogMapLite reuses LogMap 2011

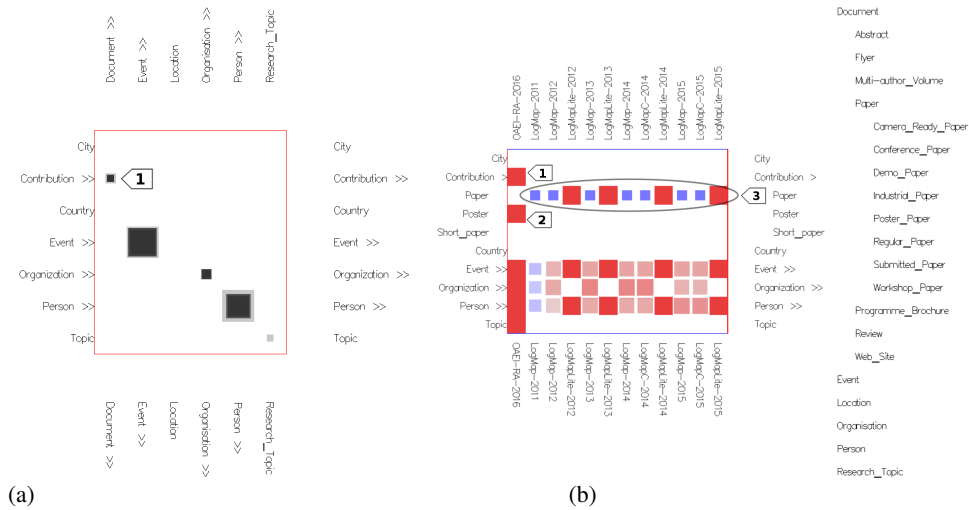


Fig. 6: Aggregated projections: (a) alignment topology and (b) concept networks.

algorithms and settings, and that there were likely significant changes in algorithms or their settings between LogMap 2011 and the following LogMaps. Further evidence about this can be found in the small-multiples view that shows concept-networks (fig. 5-b). Looking at the Event matrix (2), all LogMapLites and LogMap 2011 have computed an *incorrect* mapping between *confOf:Working Event* & *ekaw:Event*, and behave similarly in the *ekaw:Conference Participant* and *ekaw:Paper Author* matrices (9),(10). Additionally, when observing alignment matrices, we can easily see that the following two mappings are missing from all LogMap-family alignments: *confOf:Scholar* & *ekaw:Student* in the alignment topology small-multiples view (on fig. 5-a the labeled cell on matrix 1 does not exist in the other matrices) and *confOf:Working Event* & *ekaw:Scientific Event* in the concept-networks small-multiples view (matrix 3 on fig. 5-b). We can also observe on fig. 5-b two of the missing mappings noticed before, *confOf:Scholar* & *ekaw:Student* (matrix 2), and *confOf:Topic* & *ekaw:Research Topic* (matrix 3).

Going back to the initial 3D cube view in *mappings* mode (fig. 4-a), we make another observation: the cell in the top-left corner seems larger than the other cells aligned with it (fig. 4-a 2). We switch to the alignment topology view, change the cell color encoding to a uniform color (gray) and decrease transparency using the cell-opacity slider. In this aggregated view (fig. 6-a), cells that appear nested within each other (cells still have different sizes) indicate that the different matchers have computed different values for the same pair of concepts (fig. 6-a). This is the case with the top-left corner cell we noticed earlier. There are at least two nested cells between *confOf:Contribution* & *ekaw:Document* (fig. 6-a 1).

We expand both the *confOf:Contribution* and *ekaw:Document* concepts, as well as *ekaw:Paper*, which is a subconcept of *ekaw:Document*. We then switch to the concept networks projection view in *similarities* mode (fig. 6-b). We can now explore the matchers' behavior (matchers have become columns in this view) while hovering over the con-

Table 1: Functionality for alignment evolution systems.

Category	Functionality
I	F1 show an alignment version
I, E	F2 show different alignments in evolution graph
I, E	F3 show changes/diff between alignments (with change types)
I	F4 show summary of changes
I	F5 show specialized view of changes
I	F6 show change history of a mapping / concept in mappings
I, E	F7 show provenance information
I	F8 show information about/context of concepts in mappings
I, E	F9 compare different alignment versions
I	F10 search and query alignment
I	F11 query old versions using terminology of new version
I, E	F12 discover trends
I, E	F13 discover volatile and stable regions

cepts from the *ekaw* ontology on the right of the matrix. Observing the mappings in RA (the left-most column) around *confOf:Contribution*, we can see that the mappings for *confOf:Contribution* (①) and *confOf:Poster* (②) are only found in RA, and that there is no mapping under *confOf:Paper* (③) in RA (while it is found by all other matchers). To better observe where these mappings are, we switch to the concept-networks small-multiples view (not shown) and can see that the missing mappings are *confOf:Poster & ekaw:Poster Paper* and *confOf:Contribution & ekaw:Paper*. The wrong mapping is *confOf:Paper & ekaw:Paper*.

5.2 Ontology Alignment Evolution

We now describe how Alignment Cubes satisfy requirements for ontology alignment evolution systems (closely connected to UC3).

A five-steps process for ontology evolution was proposed in [23]. It can be generalized to ontology alignment evolution. The first step of the process focuses on the exploration of the ontologies. This is where an interactive visualization tool is most needed. Requirements to support the first step have been identified in [14]. In Table 1, we show an adaptation of these requirements for step 1, targeted at the alignment evolution. The category indicates whether we deal with inspection (I) or explanation (E).

An Alignment Cube addresses requirements F1–F9, F12–F13. It visualizes every alignment (F2) as a separate, labeled (F1) matrix showing both ontologies on its sides (F8). The cube, projections and small multiple views support visual identification and comparison of alignment revisions (F9), trends (F12), volatile and stable regions (F13) and changes between alignments (F3). The latter allows for deriving a history of changes made to mappings (F6). The alignment topology projection can be particularly useful for showing a summary of changes (F4). Additional matrices or small-multiples views can be introduced to show specialized views of changes (F5) and provenance information (F7).

6 Discussion and Future Work

Validation of the approach: We demonstrated some of the analytic tasks that Alignment Cubes support by describing a scenario in Section 5.1. Making such observations took only a couple of minutes with Alignment Cubes. We could identify the common trends between 12 alignments, and found five mappings missing from all 11 alignments and one wrong mapping in all 11 alignments. Furthermore, we found one incorrect mapping in five alignments (LogMap2011 and LogMapLite2012–2015). While we could likely have found missing and incorrect mappings via other means, this was achieved here easily, and in a very short time. The conclusions we draw come from integrating several observations, including all alignments after an exploration guided by the observed regions. Some observations could be made in more than one view, which allowed interpreting them in different contexts, and yielded additional opportunities to observe them if we had missed them earlier.

In Section 5.2, we showed that Alignment Cubes address many of the requirements expressed about understanding the evolution of ontology alignments. Alignment Cubes also support requirements in the *inspection* and *partial explanation* dimensions from [12]: representing dense region, ontologies, mappings and mappings suggestions, filtering, and providing a starting point for exploration. It also covers visual analytics and grouping aspects from [7]. It could further support the analytic tasks identified in [1]—evaluation of matcher performance, exploration and comparison and diagnosis (which comprises iterations of the other tasks).

User Evaluation: Matrix Cubes has been evaluated with domain experts in Astronomy and Neurosciences, and have been found to be understandable after some initial learning [2]. We have performed pilot tests with ontology alignment experts using the smallest task in the LargeBio track in the OAEI⁷ and are currently planning more extensive usability tests with experts in ontology alignment in order to identify further use cases and to inform the design and development of future features.

Scalability: The ontologies in this paper contained less than 100 concepts, and individual alignments had no more than 20 mappings. The tool was successfully tested using the ontologies (3696 and 6488 concepts in each) and all 12 alignments from the smallest task in the LargeBio track in the OAEI. Alignment Cubes have provided a compact overview for all of our examples, and should remain relatively compact with a higher number of mappings. Interactively collapsing and expanding concepts has proven useful to cope with the size of the ontologies and the mappings. Due to the size and depth of the ontologies, the small-multiples views required some pan and zoom on a regular workstation monitor. This can be addressed by filtering, or to some extent by using a large very-high-resolution display. Using larger display surfaces to show multiple views simultaneously is one of the directions identified in [11].

Generalizability and Availability: Though we developed Alignment Cubes for ontology alignment, our adaptations to the original Matrix Cubes are generalizable to any (un)directed bi-graph. Examples include networks connecting authors to their publications, or documents to keywords, or proteins to certain biological functions.

⁷ <http://www.cs.ox.ac.uk/isg/projects/SEALS/oei/2016/>

Future Features: Drawing from our experience with the tool so far, we have identified several directions for future extensions. Immediate improvements to consider include: adding visual support for different types of mappings (subsumptions, asserted, derived, etc.); exploring clustering and reordering algorithms to further support trend and pattern discovery; further supporting comparisons by visualizing the results of set operations (union, intersection, complement) as matrices. As discussed in [12,7], providing explanations about why and how a mapping has been computed supports decision making. We thus plan to explore different ways to compactly present such information to users.

In the longer term, we are interested in investigating the integration of Alignment Cubes with the SEALS platform used in OAEI. This will also open the stage for investigating advantages, drawbacks and methodological issues around the two evaluation approaches: comparative visual exploration at a detailed level, and overall assessment of the quality of alignments.

7 Conclusions

This work aims to take the evaluation of ontology alignments' quality beyond general measures such as precision, recall and F-measure. It identifies several use cases and shared tasks where comparison and exploration of multiple alignments at a high level of detail is needed. As the number of approaches and algorithms grows, capturing and analyzing their similarities and differences at varying levels of granularity will facilitate the understanding of their features, and provide additional means for alignment evaluation, hopefully contributing to driving the field forward. We see the evaluation of ontology alignments as an exploratory task, and discuss several activities that could be efficiently supported by an interactive tool—interactive visual exploration at different levels of granularity to perform compare & contrast tasks. Drawing from the field of data visualization, we significantly adapt and extend a technique for dynamic network visualization. Our approach, *Alignment Cubes*, enables the interactive visual exploration of alignments and supports views at different levels of detail. We show their usefulness for the purpose of exploration of multiple alignments by describing a scenario where, in only a few minutes, we could identify several missing and incorrect mappings. This initial experience with the tool is encouraging, and we strongly believe that alignment evaluation should consider other means beyond precision, recall and F-measure, including visual exploration. We hope to integrate this tool with the SEALS platform used in the OAEI campaigns and evaluate its usefulness with developers and researchers involved and participating in them.

Acknowledgements. This work has been supported by the Swedish e-Science Research Centre (SeRC), the Swedish national graduate school in computer science (CUGS) and the EU project VALCRI (FP7-IP-608142).

References

1. J Aurisano, A Nanavaty, and I Cruz. Visual analytics for ontology matching using multi-linked views. In *VOILA*, pages 25–36, 2015.

2. B Bach, E Pietriga, and J-D Fekete. Visualizing dynamic networks with matrix cubes. In *CHI*, pages 877–886, 2014.
3. F Beck, M Burch, S Diehl, and D Weiskopf. The state of the art in visualizing dynamic graphs. *EuroVis STAR*, 2, 2014.
4. S Carpendale. Evaluating information visualizations. In *Information Visualization*, pages 19–45, 2008.
5. I Cruz, C Stroe, and M Palmonari. Interactive user feedback in ontology matching using signature vectors. In *ICDE*, pages 1321–1324, 2012.
6. T Dang, N Franz, B Ludäscher, and AG Forbes. Provenancematrix: A visualization tool for multi-taxonomy alignments. In *VOILA*, pages 13–24, 2015.
7. Z Dragisic, V Ivanova, P Lambrix, D Faria, E Jiménez-Ruiz, and C Pesquita. User validation in ontology alignment. In *ISWC*, pages 200–2017, 2016.
8. Z Dragisic, V Ivanova, H Li, and P Lambrix. Experiences from the anatomy track in the ontology alignment evaluation initiative. 2017. submitted.
9. S Falconer and M-A Storey. A Cognitive Support Framework for Ontology Mapping. In *ISWC/ASWC*, pages 114–127, 2007.
10. T Isenberg, J Chen, M Sedlmair, and T Möller. A systematic review on the practice of evaluating visualization. *IEEE TVCG*, 19(12):2818–2827, 2013.
11. V Ivanova. Applications of large displays: Advancing user support in large scale ontology alignment. In *Doctoral Consortium @ ISWC 2016*, pages 50–57, 2016.
12. V Ivanova, P Lambrix, and J Åberg. Requirements for and evaluation of user support for large-scale ontology alignment. In *ESWC*, pages 3–20, 2015.
13. Z El Jerroudi and J Ziegler. imerge: interactive ontology merging. In *EKAW*, pages 52–56, 2008.
14. P Lambrix, Z Dragisic, V Ivanova, and C Anslow. Visualization for ontology evolution. In *VOILA*, pages 54–67, 2016.
15. P Lambrix and H Tan. A tool for evaluating ontology alignment strategies. *J Data Semantics*, VIII:182–202, 2007.
16. M Lanzenberger, J Sampson, and M Rester. Ontology visualization: Tools and techniques for visual representation of semi-structured meta-data. *J UCS*, 16(7):1036–1054, 2010.
17. Y Li, C Stroe, and I F. Cruz. Interactive visualization of large ontology matching results. In *VOILA*, pages 37–48, 2015.
18. L Otero-Cerdeira, F J Rodríguez-Martínez, and A Gómez-Rodríguez. Ontology matching: A literature review. *Expert Systems with Applications*, 42(2):949–971, 2015.
19. E Peukert, J Eberius, and E Rahm. Amc-a framework for modeling and comparing matching systems as matching processes. In *ICDE*, pages 1304–1307, 2011.
20. B Severo, C Trojahn, and R Vieira. A gui for visualising and manipulating multiple ontology alignments. In *ISWC (Posters & Demos)*, pages 37–48, 2015.
21. P Shvaiko and J Euzenat. Ontology Matching: State of the Art and Future Challenges. *J. Knowledge and Data Engineering*, 25(1):158–176, 2013.
22. T Von Landesberger, A Kuijper, T Schreck, J Kohlhammer, J J van Wijk, J-D Fekete, and D W Fellner. Visual analysis of large graphs: state-of-the-art and future research challenges. In *Computer graphics forum*, volume 30, pages 1719–1749. Wiley Online Library, 2011.
23. F Zablith, G Antoniou, M d’Aquin, G Flouris, H Kondylakis, E Motta, D Plexousakis, and M Sabou. Ontology evolution: a process-centric survey. *The knowledge engineering review*, 30:45–75, 2013.