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

Valentina Ivanova[1], Benjamin Bach[2], Emmanuel Pietriga[3], and Patrick Lambrix[1]

[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. The quality of ontology alignments is evaluated by comparing to reference alignments and calculating general measures, i.e., precision, recall and F-measure. These measures, however, only provide an overall assessment of the alignments' quality, but do not reveal differences and commonalities between alignments at a finer-grained level such as, e.g., regions or individual mappings. This demands comparative exploration of alignments at different levels of granularity. Furthermore, reference alignments to compare to are often unavailable, which makes such comparative evaluation even more important. To address this issue, we introduce Alignment Cubes which supports efficient interactive visual exploration of multiple alignments at different granularity levels.

1 Motivation

Ontology matching is an active research area and many tools and approaches have been developed in the last 15 years. Tools are evaluated by comparing their alignments against *reference alignments - RAs* - and computing measures, such as precision, recall and F-measure. These measures, however, provide only an overall assessment of alignments' quality and cannot reveal in what aspects one tool outperforms another. Furthermore, RAs are often not available and these evaluation measures cannot be computed.

To understand the strengths and weaknesses of a tool and to compare it to others, we seek answers to questions such as: Which mappings are computed by all or most tools? Are there mappings which are rarely computed or not at all? Do tools compute mappings for the same regions of the ontologies? [4]. These questions demand flexible exploration and comparison of the alignments at different granularity levels and cannot be answered by the aforementioned measures. Practitioners resort to writing custom scripts, which can be error-prone, and time- and effort-consuming as such scripts have to be crafted for every single question. Besides, comprehending their results is cumbersome especially when the size and number of alignments grow.

In our main conference paper [5], we identify several evaluation use cases and discuss shared activities that could be efficiently supported through interactive visualization. To address these use cases, we present *Alignment Cubes*—an interactive visualization environment for the comparative exploration and evaluation of multiple alignments at different levels of granularity. This demo paper is a companion paper to this research

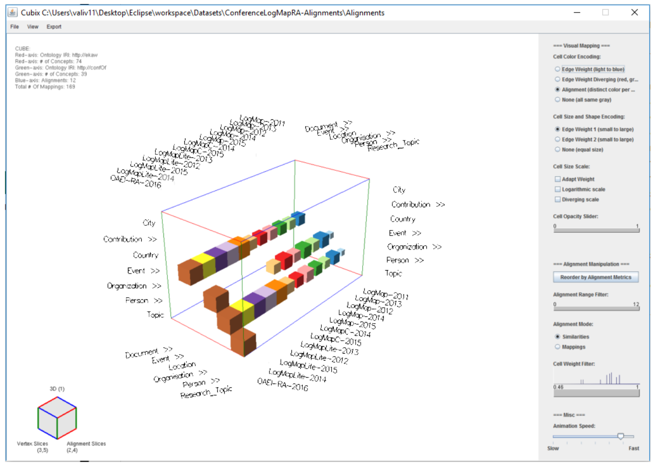

Fig. 1: Alignment Cubes User Interface

paper. As this is an interactive tool, it will be best demonstrated in an interactive setting where attendees can interact with it using their own datasets.

Related Work: Existing evaluation frameworks (SEALS[4], KitAMO [7] and AMC [8]) have focused mostly on back-end features, such as storing alignments, configuring and executing algorithms. Little attention has been devoted to the interactive exploration of alignments. Interactive visual approaches have the potential to efficiently support such interactive exploration by taking benefit from the humans' powerful visual perception system. Similarly to one of Alignment Cubes' small multiples view, AgreementMaker [2] visualizes several alignments as juxtaposed matrices. Another tool capable of presenting several alignments is VOAR [9]. However, the view quickly becomes cluttered as the size of ontologies and the number of alignments grow.

2 Alignment Cubes

In [5], we identify several analytic tasks shared by our use cases, that could be efficiently supported by interactive exploration at different granularity levels to perform, e.g., compare and contrast tasks. These high-level features together with our previous analysis of user interfaces for ontology alignment [6,3] served as guidelines for the tool we developed. We drew from state-of-the-art approaches in the field of dynamic network visualization, and identified Matrix Cubes [1] as a promising visual approach to serve as a foundation for our tool - Alignment Cubes[5].

[4] `http://seals-project.eu` — Semantic Evaluation At Large Scale

[5] `http://www.ida.liu.se/~patla00/publications/ISWC17` provides supplemental material to this submission: a screencast and a downloadable version of the tool itself.

2.1 Alignments Presentation

To make efficient use of the available screen real-estate, an alignment is presented as a matrix where the rows hold the concepts from one of the ontologies and the columns hold those from the other. The ontologies are depicted as expandable and collapsible indented lists (according to their taxonomic relationships) on both sides of the matrix. Cells denote existing mappings between concepts in the respective rows and columns. Stacking several matrices (alignments) on top of each other creates an alignment cube.

Fig. 1 shows the initial view of the tool. Two of the ontologies from the OAEI Conference track, *ekaw* (columns, 77 concepts) and *confOf* (rows, 38 concepts) are depicted on the red and green axes. 12 alignments are laid out along the blue axis—the RA for 2016 and alignments computed by the LogMap-family of systems from 2011 to 2015. Each alignment is color-coded to visually differentiate the mappings, by grouping the cells that belong to each of them using a pre-attentive variable.

2.2 Granularity Levels

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—we introduced *alignment modes*. In *similarities* mode a filled cell represents an existing mapping between a pair of concepts. In *mappings* mode, 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 where smaller cells indicate mappings deeper in the hierarchy.

2.3 Interactive Visual Exploration

The Alignment Cubes user interface provides a variety of interactions to support visual exploration, shown in fig. 1: 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. 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 and to facilitate comparison, the order of alignments (slices) in the cube can be changed based on precision, recall, F-measure or a matcher name.

2.4 Compare and Contrast

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 orthogonal faces of the cube, (c) side-by-side layout of the cube's slices (juxtaposition)—small multiples view, (d) in-place rotation of individual slices for quick preview. The 3D cube provides an overview of the number of alignments, and 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 (on fig. 1 we quickly notice that a mapping is present in the RA (the single brown cell) but not in the other alignments). 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. Side-by-side views provide the most detailed view onto the data by entirely decomposing the cube and showing each alignment in detail. 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 side-by-side view. Both projections/side-by-side 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.

3 Demonstration Scenario

During the Demo session we will demonstrate the different views and exploration features of the tool by conducting a comparative evaluation of several alignments. We would also like to invite participants to the Demo session to use their own datasets with Alignment Cubes[6].

References

1. B Bach, E Pietriga, and J-D Fekete. Visualizing dynamic networks with matrix cubes. In *CHI*, pages 877–886, 2014.
2. I Cruz, C Stroe, and M Palmonari. Interactive user feedback in ontology matching using signature vectors. In *ICDE*, pages 1321–1324, 2012.
3. 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.
4. Z Dragisic, V Ivanova, H Li, and P Lambrix. Experiences from the anatomy track in the ontology alignment evaluation initiative. 2017. submitted.
5. V Ivanova, B Bach, E Pietriga, and P Lambrix. Alignment cubes: Towards interactive visual exploration and evaluation of multiple ontology alignments. In *ISWC*, 2017.
6. 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.
7. P Lambrix and H Tan. A tool for evaluating ontology alignment strategies. *J Data Semantics*, VIII:182–202, 2007.
8. 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.
9. B Severo, C Trojahn, and R Vieira. A gui for visualising and manipulating multiple ontology alignments. In *ISWC (Posters & Demos)*, pages 37–48, 2015.

[6] If you are interested to do so, please contact the first author in advance.