Visualizing Mappings Between Pairwise Ontologies - An Empirical Study of Matrix and Linked Indented List in Their User Support During Class Mapping Creation and Evaluation

Bo Fu[0000-0001-9874-9551], Allison Austin, Max Garcia

California State University Long Beach, California, CA 90840, USA
Bo.Fu@csulb.edu; {Allison.Austin|Max.Garcia}@student.csulb.edu

Abstract. Visual support designed to facilitate human interaction with semantic data has largely focused on visualizing entities such as classes and relationships within an ontology. Comparatively speaking, less attention has focused on visualizing mappings established between independent ontologies and determining the effectiveness of mapping visualizations. This paper presents a user study of the matrix and the linked indented list visualization in their visual support for users during creation and evaluation of class mappings between pairwise ontologies. A total of 81 participants took part in a task-based controlled experiment, with the aim of assessing the extent to which a given visualization supports recognition of visual cues, validation of existing mappings, and creation of new results. Based on empirical evidence collected from the participants in their speed to complete the tasks, their success in answering various questions, as well as their physiological sensory data such as eye gaze, we aim to quantify user performance and visual attention demanded in the use of the two aforementioned mapping visualizations. The experimental results indicate that the linked indented lists and the matrix visualization are comparable in terms of effectiveness and efficiency when assisting users in the given task scenarios with marginal differences. However, linked indented lists are likely to demand less effort from the users' visual perceptual systems with statistically significant differences found in several gaze measures, including the physical distances needed to locate relevant visual information, number of fixations and time required to process visual cues, and the overall efforts in scanning the visual scene.

Keywords: Ontology Mapping Visualization, Matrix, Linked Indented List, Eye Tracking.

1 Introduction

A growing body of research aimed to overcome challenges in providing user-centered tools and systems to support and enhance human interaction with semantic data and semantically empowered systems has pushed advances in the Semantic Web to beyond simply improving the efficiencies and accuracies of various algorithms. While leveraging semantic technologies and computational intelligence in problem-solving remains a critical component in the Semantic Web, cohorts of research in interfaces for the Semantic Web as well as semantically empowered systems have contributed to recent innovations in interactive Semantic Web. Within this context, the *Interactive Semantic Web* (ISW) refers to the design, development, and refinement in all aspects of human interaction with the Semantic Web, whereby the integration of interfaces and interactions bridged between human users and semantic content and technologies contribute to enhanced and improved use of the Semantic Web and all its constituents.

Well-known Semantic Web tools and systems designed to support humans (experts and non-experts alike) in the loop often incorporate interactive visualizations and visual analytics in order to best assist users, such as Protégé[1]/WebProtégé[2] for ontology editing, the Wikidata Query Service[3] for querying semantic content, and Punya[4] for building semantically enriched mobile applications, to name just a few. Other efforts to assist novice users in the navigation and querying of knowledge graphs is also demonstrated in [1], where traditional SPARQL queries are simplified by means of visual queries. A more recent example of a knowledge graph empowered search and visualization system is presented in [2], which aims to facilitate researchers in their explorations of relevant scientific publications by leveraging thematic rule-based associations, networks of co-publications, and co-occurring topics. To accelerate biomedical research and discovery in the fight against COVID-19, a tool to automatically extract and visualize argumentative graphs of clinical articles is demonstrated in [3], whereby information extraction is enhanced by a continuously enriched knowledge graph to assist with human decision-making in healthcare.

Evidentially, appropriate integrations of interaction and visualization techniques lie at the center of ISW systems, where visual cues have continued to provide the necessary means for users to better understand, explore, correct, and modify semantic data and semantically empowered systems. As such, assessing the effectiveness of interactive visual support in their capacity to support users in the ISW has been investigated in prior work such as [4-6], where commonly used tools and techniques to visualize ontologies have been evaluated extensively. In comparison, less research focus has been placed on relationships established between independent ontologies, such as ontology mapping visualization. To this end, this paper focuses on a relatively unexplored area in mapping visualizations between pairwise ontologies. Through an empirical user study involving 81 participants in the context of creating and evaluating class mappings between pairwise ontologies, we report on the effectiveness and efficiency of two popular techniques in mapping visualization, namely the linked indented list and matrix visualization. In addition, we analyzed user gaze to investigate how users divided their visual attention when interacting with the aforementioned mapping visualizations, and report on a number of observations in gaze tendencies in each visualization technique through eye tracking. A main contribution of this paper lies in the generation of new

[1] `protege.stanford.edu`, last accessed 07/31/2023.

[2] `webprotege.stanford.edu`, last accessed 07/31/2023.

[3] `query.wikidata.org`, last accessed 07/31/2023.

[4] `punya.mit.edu`, last accessed 07/31/2023.

knowledge on the user experience with two interactive mapping visualization techniques frequently found in current tools and systems. With knowledge graph and ontology development environments amongst some of the most impactful and widely adopted technologies derived from the Semantic Web community, the empirical findings collected from a reasonably large user group shown in this paper will likely inform future design and development of these environments to be more user-centered with viable visual and interactive support for the average user.

2 Related Work

User involvement and how best to support human users remain a critical research area in ontology mapping given that automatically generated mappings typically require further refinement and user intervention. An example system aimed to support users in ontology and entity matching is presented in [7], which provides an environment for users to generate mappings based on element descriptions and system-generated matching suggestions. Visual assistance is provided in the form of interactive indented lists of the source and target ontologies, while visualizations of the mappings themselves are not supported. A review of mapping systems[5] designed to support various stages of ontology mapping is presented in [8]. Amongst the eight systems surveyed, five [9-18] use connecting links to visualize mappings between ontological nodes belonging to different ontologies, whereby the ontologies themselves are often arranged as indented lists in the visualization. Such techniques to visualize mappings (as links) between ontologies (as indented lists) are referred to as *linked indented lists* (LIL) in this paper. LIL remains a popular technique to visualize relationships between structured datasets as also seen in the RBA tool [23], which is designed to facilitate mappings between relational databases and ontologies. Subsequent research has also focused on reducing visual clutter in LIL when a large number of mappings is present, through layout techniques such as edge bundling [24].

Later systems have utilized matrices or grids when visualizing pairwise mappings, whereby mappings are illustrated as occurrences of vectors in a 2D plane, and the associated ontological entities are displayed along horizontal and vertical axes. In this paper, such visualization techniques are referred to as *matrix* visualization. An example of a matrix visualization is shown in [25] with mappings visualized at the ontology level as well as at the node level. To facilitate visual search, a user can sort the source and target ontologies/nodes displayed along the axes in alphabetical order. Another example can be found in ProvenanceMatrix [16], where the user can sort the matrix axes in breadth-first, depth-first, and similarity orderings. In addition to 2D matrices, later research has investigated adding a third dimension to utilize 3D cubes when visualizing mappings such as [27] with the goal of providing various levels of granularity when exploring and evaluating ontology mappings. Other visualization systems have utilized multiple views when presenting mappings to the user, in an effort to better support

[5] The review includes the following systems: AgreementMaker [9-11, VOILA 2015], AlViz [12], AML [13, 14], CogZ/Prompt [15-17], COMA [18], LogMap [19], SAMBO [20, 21], and RepOSE [22].

various viewpoints during visual search such as VOAR 3.0 [28] that visualizes mappings in both LIL and node-link diagrams, AlignmentVis [29] that uses LIL, matrices, scatterplots, and parallel coordinates to visualize mappings and related statistics in a mapping scenario, as well as BioMixer [30] that visualizes mappings as timelines and node-link diagrams with multiple coordinated layouts.

Recent research has also investigated the application of block metaphors that are frequently used in visual programming languages in the context of semantic mapping such as [31] and the use of pie charts in large scale ontology mappings [32]. Moreover, efforts to broaden ontology visualization to beyond visualizing hierarchies but to also include non-hierarchical relationships in large ontologies is proposed in [33], where icicle plots coupled with visual compression have been shown to improve space-efficiency and reduce visual structural complexity. Given that node-link diagrams are often used to visualize ontological entities as nodes in a network with connecting edges illustrating ontological relationships such as *is-a* relations, a natural expansion is to also include mappings amongst ontological entities, by inserting visual associations (i.e., more edges that can be visually distinguished, often by line color or style, from those of *is-a* relations) into the visualization to illustrate additional relationships such as mappings among otherwise isolated nodes. This node-link technique can be observed in a number of systems such as [33-37] and studied extensively in [4-6, 39-44]. Recent efforts to advance node-link diagrams in the Information Visualization (InfoVis) community include determining the effects of progressively increasing encoded information in node/node-link/node-link-group diagrams [39], comparing different methods to visualize long, dense, complex, and piecewise linear spatial trajectories [40], displaying clusters overlaid using node coloring, GMap, BubbleSets, and LineSets [41], encoding multivariate and continuous data on edges in 3D node-link diagrams [42], developing novel exemplar-based layout to adjust substructures [43], and improving readability via layered layout that considers crossing reduction, edge bendiness, nested and multi-layer groups simultaneously [44]. Other efforts to improve matrix visualization in the InfoVis community have investigated ways to enhance visual analysis with hierarchy matrix [45] and ordering effects within matrices [46]. A recent evaluation study [47] from the InfoVis community compares bipartite, node-link, and matrix-based network presentation in a range of tasks such as network class identification, cluster detection, network density estimation, to demonstrate overall network structures are best illustrated with bipartite layouts.

There is however limited research focusing on the evaluation of interactive visualizations designed for ISW. In the context of ontological data modelling, prior evaluations have largely focused on assessing user experience with visualizations of class hierarchies. For instance, the usability of indented lists and node-link diagrams when visualizing class relationships are assessed in [4] with eye tracking results [5] providing further insights on the strengths and weaknesses of each visualization technique. In the context of supporting users in large-scale ontology alignment, a study [48] using heuristic evaluations and feedback from 8 participants across three mapping systems (that all utilize LIL including CogZ [15], COMA [18], and SAMBO [20, 21]) aims to elicit system design requirements. Given the frequent use of LIL and matrices across a number of existing tools and systems as outlined above, and with a lack of evaluations of the two, there is a pressing need to assess these visualizations designed specifically for mappings. To this end, this paper presents a controlled, between-subject user study utilizing eye tracking in the assessment of the LIL and matrix visualizations in the context of human mapping creation and evaluation.

3 Experimental Setup

3.1 Tasks and Visualizations

The goal of the experiment is to compare if one visualization technique may be better suited for a particular type of mapping creation and/or evaluation tasks than the other. To simulate an environment where the user tasks would require human interaction and comprehension of the given visual cues, we asked participants to answer a series of 15 questions while supported by mapping visualizations between an ontology pair. Table 1 presents an overview of the questions in two domains. These questions can be categorized as i) *identification* tasks (Q.1-6), where successful completion requires a participant to recognize what is and is not already visually displayed in the mapping visualization; ii) *validation* tasks (Q.7-12), where successful completion requires participants to verify the accuracy of a mapping displayed or the lack thereof; and iii) *creation* tasks (Q.13-15), where successful completion requires a participant to generate new knowledge (i.e., new mappings) that is not already displayed in the visualization. These questions are not intended to be exhaustive, but as examples of typical scenarios during mapping creation and evaluation where a human user needs to comprehend visual cues in the process of establishing correct and complete mappings between pairwise ontologies. The goal of these questions is to simulate a necessary environment for the purpose of enabling comparative studies between the matrix and LIL visualization in the context of class mappings with a range of example conditions where different visual needs may be demanded during human decision-making.

Where appropriate, some questions are presented as multiple-choice questions with a dropdown menu containing 2 or 4 options with one correct answer (e.g., in Q.1, 4 numbers are shown in a dropdown menu where one of them is the correct answer; in Q.4, yes or no options are given in a dropdown menu), and others are presented as open-ended questions with textboxes to fill in (e.g., Q.15). Identification tasks direct users to decode a given visual cue (e.g., Q.3 requires a user to describe what a link or solid/dotted cell between two entities entails, and Q.4 requires a user to interpret what a non-existent link or empty cell entails), validation tasks require a user to assess existing mapping quality (is a link/solid cell correct or wrong), and creation tasks ask a user to generate new knowledge by creating additional mappings not already shown in the visualization (e.g., in Q.15, users can create new mappings if they believe there are absent mappings such as those prompted in Q.4 and Q.6). The same set of mappings (containing the same correct, incorrect, and incomplete results) were visualized in each domain, so that the only difference between user groups remain as the mapping visualization themselves as opposed to differing mapping results shown to the user, to ensure the comparison between LIL and matrix visualization is made fair.

Table 1. Mapping Creation and Evaluation Questions Used in the Study.

	Conference Domain	Anatomy Domain
Identification	1. How many mappings are shown in the visualization in total? 2. How many classes is *Author* (in the source ontology) mapped to? 3. What is *SlideSet* (in the source ontology) mapped to? 4. Can *Person* (in the source ontology) be mapped to another class (in the target ontology)? 5. What is *ConferenceDinner* (in the source ontology) mapped to (in the target ontology)? 6. Can *Workshop* (in the source ontology) be mapped to another class (in the target ontology)?	1. How many mappings are shown in the visualization in total? 2. How many classes is *Skin* (in the source ontology) mapped to? 3. What is *Viscera* (in the source ontology) mapped to? 4. Can *Joint* (in the source ontology) be mapped to another class (in the target ontology)? 5. What is *Skull* (in the source ontology) mapped to (in the target ontology)? 6. Can *Arm* (in the source ontology) be mapped to another class (in the target ontology)?
Validation	7. Is there a mapping between *AcademicEvent* (in the source ontology) and *Scientific_Event* (in the target ontology)? 8. Is *AcademiaOrganization* (in the source ontology) correctly mapped? 9. *SecurityTopic* (in the source ontology) is mapped to *Research_Topic* (in the target ontology). Is this correct? 10. *Place* (in the source ontology) is mapped to *Location* (in the target ontology). Is this correct? 11. *RejectedPaper* (in the source ontology) is mapped to *Assigned_Paper* (in the target ontology). Is this correct? 12. *IndustryOrganization* (in the source ontology) is mapped to *Organisation* (in the right ontology). Is this correct?	7. Is there a mapping between *Blood* (in the source ontology) and *blood* (in the target ontology)? 8. Is *Cartilage* (in the source ontology) correctly mapped? 9. *Urinary_System_Part* (in the source ontology) is mapped to *muscle* (in the target ontology). Is this correct? 10. *Cheek* (in the source ontology) is mapped to *cuticle* (in the target ontology). Is this correct? 11. *Skin* (in the source ontology) is mapped to *skin* (in the target ontology). Is this correct? 12. *Mucus* (in the source ontology) is mapped to *nasal mucus* (in the target ontology). Is this correct?
Creation	13. Which class could *Attendee* (in the source ontology) be mapped to (in the target ontology)? 14. Which class could *ConferenceDinner* (in the left ontology) mapped to (in the right ontology)? 15. Is there any other mapping(s) that should be created between the ontologies but is currently absent from the visualization?	13. Which class could *Heart* (in the source ontology) be mapped to (in the target ontology)? 14. Which class could *Lip* (in the source ontology) be mapped to (in the target ontology)? 15. Is there any other mapping(s) that should be created between the ontologies but is currently absent from the visualization?

Two pairs of ontologies are used in this study. These ontologies and their respective mappings are based on the conference and biomedical tracks at the Ontology Alignment Evaluation Initiative (OAEI)[6]. Mappings in the OAEI gold standards have been

[6] oaei.ontologymatching.org, last accessed 07/31/2023.

modified in the study for the purpose of presenting a mixture of correct, incorrect, and incomplete mappings in the visualization, so that participants can complete a range of creation and evaluation scenarios as discussed above. An overview of the characteristics of the ontologies used in the study is presented in Table 2. The conference ontologies, taken as is from the OAEI, have 97 and 73 classes in the source and target ontology respectively. In order to ensure comparable sizes, for the biomedical ontologies, we used the portions that focus on the human anatomy, whereby there are 115 and 97 classes in the source and target ontology respectively. For each domain, the accompanying mapping visualizations show a total of 10 mappings between the ontology pair, amongst which, 5 are correct and 5 are incorrect. There are also 5 additional mappings that are missing from the visualization. The experiment variables such as the ontology size, task scenario, and the number of mappings were controlled in the study to ensure their potential impact on user performance is minimized so that a given mapping visualization (i.e., either the LIL or matrix visualization) can be assessed as the independent variable in the experiment. In other words, if a difference were to be found in user success or completion time, the underlying cause is likely attributed to the visualization used as opposed to a simple result of a smaller ontology pair or fewer mappings to inspect. The domains, ontologies, and mappings used in our study are not intended to be exhaustive, but as example scenarios aimed to provide the necessary experimental conditions to compare the LIL and matrix visualization. For those interested in specific domains or ontologies with certain characteristics, it would be necessary to target other mapping scenarios/domains not presented in this paper.

Table 2. Ontologies and Mappings Used in the Study. The longest path to root defines the depth of a class hierarchy, the largest sibling pool refers to the most number of subclasses for a given class, multiple inheritances refer to instances of classes with more than one superclass.

	Conference Ontologies		Anatomy Ontologies	
	Source	Target	Source	Target
Number of Classes	97	93	115	97
Longest path to root	4	6	6	5
Largest sibling pool	21	9	12	26
Multiple inheritance	-	1	2	-
Correct mappings	5		5	
Incorrect mappings	5		5	
Incomplete mappings	5		5	

The matrix and LIL visualization used in this study are implemented using the D3 JavaScript Library[7]. Fig. 1 presents a screenshot of the study interface showing how questions and mapping visualizations are presented to a participant in a Web browser. In the matrix visualization (Fig. 1a), the source and target ontologies are visualized as indented lists along the axes. Toggling a node on the axis will expand or collapse child nodes (if any), where solid blue triangles indicate subclasses that can be further revealed, hollow blue triangles indicate fully expanded child nodes, and blue dots indicate childless nodes. Mappings are illustrated as cells associating pairwise entities shown

[7] `d3js.org`, last accessed 07/31/2023.

along the vertical and horizontal axes. Solid blue cells indicate mappings of entities that are already fully visible in the visualization, e.g., *ConferenceEvent* is mapped to *Event*. Dotted blue cells indicate mappings exist amongst child nodes of the corresponding entities and are yet to be revealed in the visualization, e.g., there is at least one mapping amongst the child nodes of *ConferenceEvent* and *Scientific_Event*, and a user will need to toggle these entities to further reveal the exact mapping(s) at the child level. As nodes are expanded, the matrix grows in both length and width and a user can scroll vertically and horizontally during interaction, while the entity labels along the axes remain fixed in position to facilitate readability.

In the LIL visualization (Fig. 1b), the source and target ontologies are visualized as two separate indented lists, and mappings are visualized as links connecting pairwise entities belonging to different ontologies. Users can toggle nodes to expand or collapse an entity, with solid triangles indicating expandable nodes, hollow triangles indicating nonexpendable nodes, and dotted nodes indicating childless entities. Solid blue links denote mappings of two entities that are already fully visible in the visualization, e.g., *Mucus* is mapped to *nasal mucus*. Dotted links denote at least one mapping exists amongst the children of the associated entities and the user must toggle the associated entities to reveal the mappings beneath, e.g., *Lower_Extermity_Part* in the source ontology should be toggled to reveal the child node that is mapped to *leg* in the target ontology. Finally, users can scroll vertically and horizontally during the interaction as nodes are expanded and the indented lists grow in the visualization.

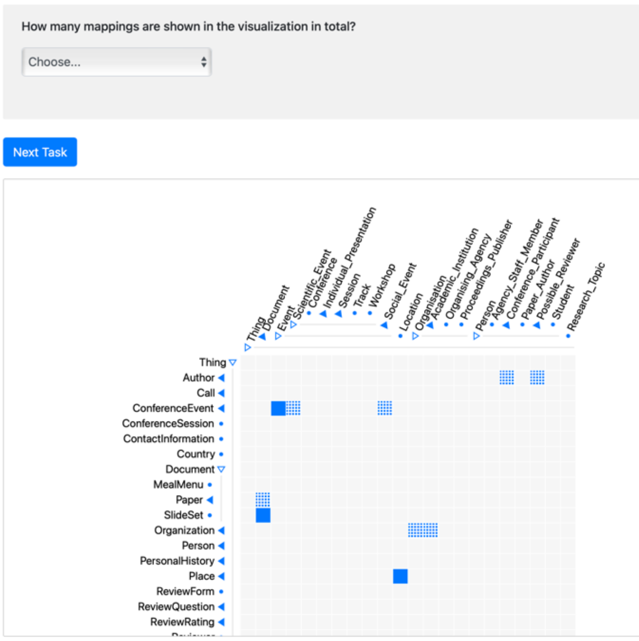

(a) Mapping Visualized in A Matrix in the Conference Domain

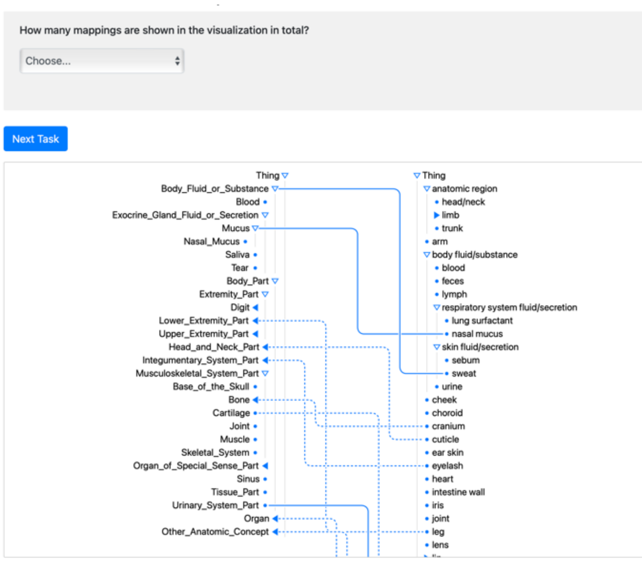

(b) Mapping Visualized in A Linked Indented List in the Anatomy Domain

Fig. 1. Study Interface Displaying Questions and Visualizations.

3.2 Participants, Protocol, and Data Collection

A total of 81 participants (with approximately 20% of them having taken a semester long graduate level introductory course on the Semantic Web) took part in this study, including a mixture of undergraduate and graduate students majoring in Computer Science, Computer Engineering, Mechanical Engineering, Applied Math, and Political Science. We recruited these participants to present a reasonably sized sample of novice users. We focused on novice users in this study as this user population would likely need most help and support. Though there is the opportunity to include expert users, it is a frequently debatable topic what qualifies one as a true expert, especially when expertise is often self-reported. In addition, variables such as user expertise will likely have an impact on user performance (e.g., if a person is more successful or faster at a given task, is such an outcome due to this person's expertise, or the result of a more effective and efficient visualization), and thus controlled in our experiment in order to keep the visualization type as the independent variable. For those interested in eliciting differences between distinct user groups (such as novice vs. expert users), it would be necessary to also include domain/ontology/visualization expertise in their experimental design. Since the goal of our study is to compare two mapping visualizations, one user sample with novices would be sufficient for this purpose. Furthermore, it is unlikely the case that real-world users of ISW technologies are mainly expert users, nor should ISW technologies require users to hold advanced degrees, intricate technical background, or expert knowledge in order to use.

Each participant completed a short tutorial on ontology mapping and the interactive features in a given visualization. They were then randomly assigned to a visualization and completed the associated questions one domain at a time. Participants were informed that all questions have correct answers and were instructed to do their best while being as fast as they can. To minimize training and ordering effects, a participant used one visualization exclusively in a study session, while completing tasks with varied orderings of the domain. Overall, two separate groups of participants completed the same questions, but supported by two distinct visualization types. In other words, this between-subject design ensures minimal learning effects (compared to a within-subject design) on the given visualization.

We collected a range of data from each participant, shown in table 3. Time on task can be further categorized by the type of questions (T_i, T_v, T_c) such as the three categories outlined in section 3.1, in addition to the overall time it takes to complete all 15 questions (T_o). Likewise, success can be determined by the question type (S_i, S_v, S_c). For the open-ended question (i.e., Q.15 in each domain) where participants were asked to generate new mappings, we computed additional metrics to quantify the precision, recall, and f-measure of the answers produced by a participant in that question (S_p, S_r, $S_{f\text{-}m}$), in the same way automated mapping algorithms are evaluated, whereby the correctness, completeness, and overall quality of the new mappings were measured against known missing mappings between the ontology pair.

Table 3. Data Collected from Each Participant.

Data Type		Description
Time on Task	T_i:	Time it takes a participant to complete all *identification* tasks.
(min)	T_v:	Time it takes a participant to complete all *validation* tasks.
	T_c:	Time it takes a participant to complete all *creation* tasks.
	T_o:	Overall time it takes a participant to complete all questions.
Task Success	S_i:	Correct answers contained in all *identification* tasks as a ratio.
(0,1)	S_v:	Correct answers contained in all *validation* tasks as a ratio.
	S_c:	Correct answers contained in all *creation* tasks as a ratio.
	C_p:	The precision of the new mappings generated in Q.15.
	C_r:	The recall of the new mappings generated in Q.15.
	$C_{f\text{-}m}$:	The f-measure of the new mappings generated in Q.15.
	S_o:	Overall success in all 15 questions.

To quantify how participants divided their visual attention during their interaction with a given mapping visualization, we used physiological sensors to collect eye gaze from each participant. More specifically, we used a Gazepoint GP3 HD eye tracker[8] with a 150Hz sample rate and a 24" Full HD Dell monitor[9] with 1920×1080p at 165Hz and 1ms response time to collect gaze data. Each participant completed a 9-point calibration before each eye tracking session to ensure maximized gaze data accuracy. Participants were seated on non-wheeled and non-swiveled office chairs and maintained relatively unchanged distances to the eye tracker, which tolerates a range of 35cm (horizontal) × 22cm (vertical) × ±15cm (depth) movement. Based on the raw gaze data generated from the eye tracker, a set of descriptive gaze measures (DGM) are computed

[8] `gazept.com/product/gp3hd`, last accessed 07/31/2023.

[9] `www.dell.com/en-us/shop/dell-24-gaming-monitor-g2422hs/apd/210-bdpw/monitors-monitor-accessories`, last accessed 07/31/2023.

for each participant after the person has completed all 15 questions using a given visualization. The goal of the DGM is to capture one's ability in selecting and maintaining awareness of specific locations and attributes of a visual scene presented. To compute the DGM, the raw gaze data produced from the eye tracker underwent a cleaning process, whereby invalid (per validity codes reported by the eye tracker, negative numbers, and when entries are off-screen), incomplete (e.g., when only one eye was captured, missing x and y coordinates), and corrupted entries (e.g., pupil dilation exceeding possible ranges, whereby anisocoria or asymmetric pupils is rarely greater than 1mm [49], and normal pupil size in adults typically varies from 2-4mm in diameter in bright light and 4-8mm in the dark [50]) were discarded.

In this paper, we report on the most notable DGM that reflect distinct behaviors for participants using one visualization versus the other such as those related to fixations and scanpaths. Fixations refer to those moments where the eyes are relatively stationary and holding a vision of focus in place, which are typically understood as information processing behaviors as a person stops scanning the visual scene at large but concentrating on extracting information from the targeted visual cues. Descriptive statistics of fixations are typically measured as sums and durations, which are indicative of the number of fixations required and the time needed to process information [51]. As a person scans for various fixations to focus on in a visual scene, the rapid eye movements between fixations are captured as saccades, which are typically understood as information searching behaviors. Saccades are typically quantified through counts and durations. In addition, a useful measure to capture the distances between successive fixations is saccadic length (in pixel). As a person searches and processes visual information, a sequence of fixations and saccades can be captured, whereby the sum of all saccadic lengths (known as scanpath length) is typically used to reflect the complete visual journey commenced [51]. In this paper, average saccadic lengths and their standard deviations (StDev) are used to describe gaze behaviors sampled from the participants. Furthermore, fixation count, scanpath length, and the StDev of fixation durations are correlated to participant success.

4 Results

To prevent data distortion, out of all data collected from 81 participants, we discarded those from 8 participants who experienced various issues during the experiment, such as reflective wear (e.g., earrings) being sampled in gaze, leaning too close (e.g., placing elbows on desk) or too far (e.g., beyond the tolerated range as discussed in section 3.2) from the eye tracker. The findings shown in this paper are aggregated results of the remaining 73 individuals, whereby 35 participants completed the given tasks using LIL and 38 participants completed the same tasks using the matrix visualization.

Fig. 2 presents the findings of participant success by task type. Across domains and irrespective of the type of task, participants were marginally more successful using the LIL (Fig. 2a-d). When creating new mappings, as shown in Fig. 3, those who used LIL produced higher quality mappings, which is evident in the higher precision (Fig. 3a), recall (Fig. 3b), and f-measure (Fig. 3c) compared to the group of participants who used the matrix visualization. Fig. 4 presents the time it takes a participant to complete the series of given tasks. Those who used LIL were found to be marginally faster compared to the others who used the matrix visualization in the identification and validation task scenarios (Fig. 4a and Fig. 4b). Notably, when creating new mappings, those who used the matrix visualization were found to be faster at task completion (Fig. 4c). Overall, time on task is comparable irrespective of the visualization used, though a greater difference is evident in the Conference domain, whereby those who used the matrix visualization were faster at task completion (Fig. 4d). In the identification tasks, participants who used LIL were faster at completing the given tasks and yielded higher success in the Conference domain. This is also echoed in the validation tasks across both domains, whereby the participants supported by LIL achieved higher success while needing less time. In the creation tasks however, participants who used the matrix visualization were faster at completing the given questions while being more successful in the Conference domain. When evaluating the quality of the new mappings participants generated in the open-ended question Q.15, those who were supported by LIL had generally produced more correct and complete results, yielding to better overall f-measure scores consequently. The difference found across all aspects of success, time, and the quality of new mappings produced are marginal and relatively comparable. This result suggest that the LIL and the matrix visualization are relatively comparable to one another, since irrespective of the visualization used, the participants performed almost equally well in all questions across both domains. All differences shown in Fig. 2-4 reported greater than 0.05 p-value, suggesting the differences found are not statistically significant. One notable finding is that in the Anatomy domain, participants were equally successful in the identification tasks irrespective of the visualization used (shown in Fig. 2a, where p>0.05), though those who used the LIL were faster (shown in Fig. 4a, where p>0.05).

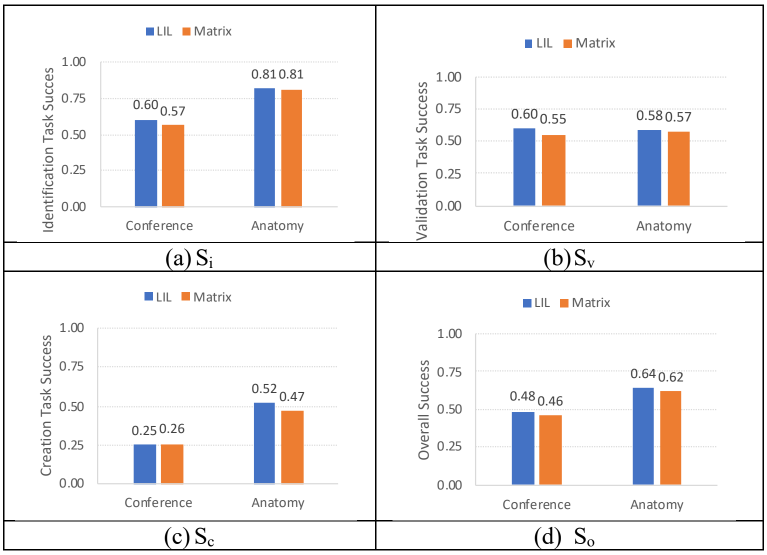

Fig. 2. Task Success by Task Type (p>0.05).

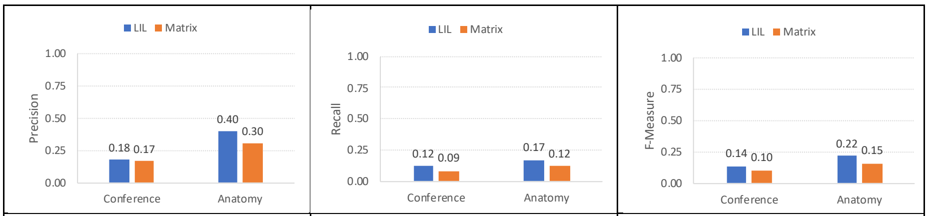

(a) C_p (b) C_r (c) $C_{f\text{-}m}$

Fig. 3. Precision, Recall, and F-Measure of the New Mappings Created (p>0.05).

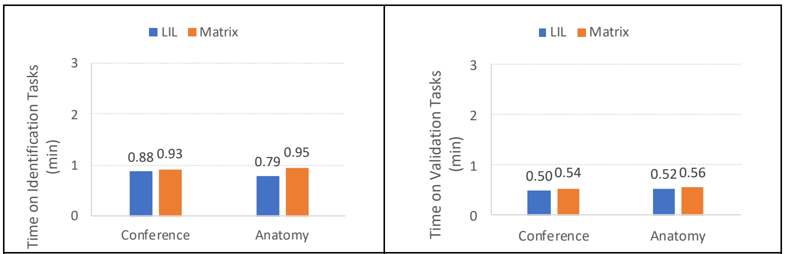

(a) T_i (b) T_v

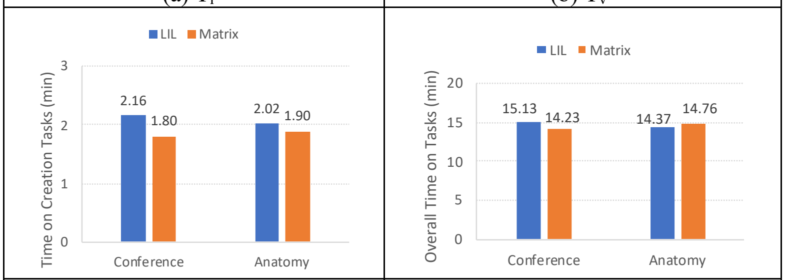

(c) T_c (d) T_o

Fig. 4. Task Time by Task Type (p>0.05).

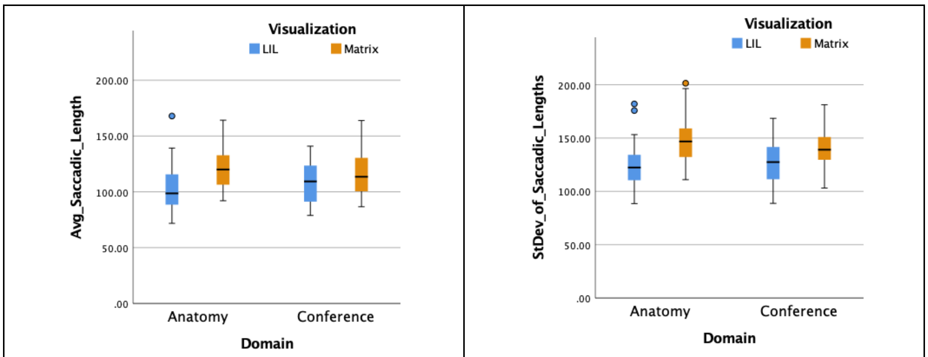

(a) Average Saccadic Length Across Domains by Visualization (b) StDev of Saccadic Lengths Across Domains by Visualization

Fig. 5. Saccadic Length Mean and Dispersion (p<0.01).

When comparing how participants divided their visual attention while interacting with the two visualizations, one notable observation is the statistically significant differences found in the participants' saccadic lengths and their StDev, as shown in Fig. 5. Across both domains, the participants exhibited consistently smaller saccadic lengths (Fig. 5a) as well as smaller StDev in the dispersion of these saccadic lengths (Fig. 5b) when using the LIL. This result indicates that when using the LIL, participants generally fixated on visual cues that are closer by and that the disparity of various points of interest was smaller. In other words, this finding suggests that the matrix visualization in comparison requires visual searches of fixations that are located further apart and that there is a greater dispersion among various visual cues relevant to the participants. As such, the matrix visualization likely demands greater efforts from one's visual perceptual system in completing the given tasks shown in this paper.

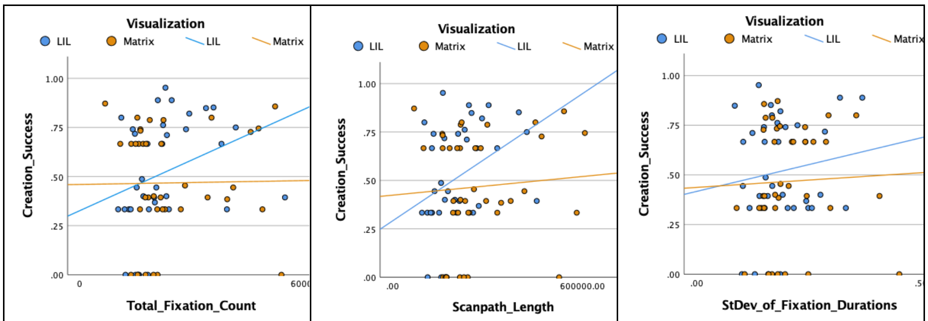

(a) LIL: r = 0.31, matrix: r = 0.01	(b) LIL: r = 0.36, matrix: r = 0.18	(c) LIL: r = 0.13, matrix: r = 0.04

Fig. 6. Correlations in the Anatomy Domain (p<0.05).

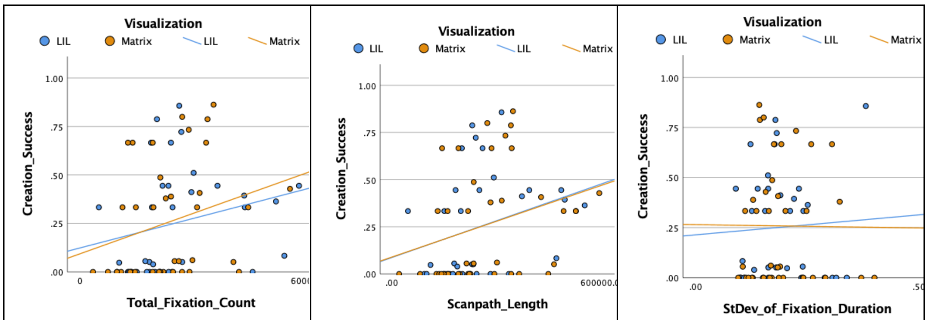

(a) LIL: r = 0.24, matrix: r = 0.26 (p<0.05)	(b) LIL: r = 0.31, matrix: r = 0.28 (p<0.05)	(c) LIL: r = 0.05, matrix: r = -0.08 (p>0.05)

Fig. 7. Correlations in the Conference Domain.

In order to examine relationships between one's performance and this person's visual attention spent during the interaction, we performed correlation coefficient tests as shown in Fig. 6 and Fig. 7. In the Anatomy domain, a participant's creation success is found to be correlated with this person's total fixation count (Fig. 6a), scanpath length (Fig. 6b), and the StDev of fixation durations (Fig. 6c). The r values indicate weak but statistically significant positive relationships between a person's success in creation tasks (Q.13-15) and the number of fixations this person sampled during the interaction, coupled with how extensive the entire scanpath was and how dispersed the person divided their attention processing information at various fixations. Notably, to achieve the same level of success, a participant using the LIL typically generated fewer fixations, shorter scanpath, and more variation in the time spent on processing information. This finding suggests that in the context of mapping creation, the LIL may require less effort from the user while assisting them in accomplishing the same level of success compared to the matrix visualization, which likely demands greater numbers of fixations to be sampled and longer scanpaths from the user.

In the Conference domain, similar weak correlations are found though with reduced differences between the two visualizations as well as the correlative relationships themselves. Creation success and scanpath length continue to demonstrate a positive relationship (Fig. 7b), although the two visualization techniques exhibit almost exactly the same degree of correlations with a negligible difference. In contrast to the findings shown in Fig. 6a, participants with the same level of success in the creation tasks sampled fewer fixations (Fig. 7a) when using the matrix visualization. In addition, participants who achieved higher success in the creation tasks exhibited greater consistency in the time spent on processing visual information at various fixations (Fig. 7c) when using the matrix visualization, though this result is not statically significant and the relationship is almost linear. This finding suggests that in the Conference domain when creating new mappings, there is marginal differences demanded from the two visualizations, though the matrix visualization is likely to demand slightly less effort in comparison to the LIL.

5 Conclusions and Future Work

The findings shown in this paper add to the existing body of knowledge in interactive visualization for semantic data and particularly mappings between pairwise ontologies. Through controlled user studies and based on empirical results, we demonstrate that the LIL and matrix visualization are highly comparable in the task scenarios investigated in this paper, since they independently led to very similar user performance (i.e., success in various types of mapping task, and time needed to complete them) in a between-subject experimental setting across domains. In other words, the participants were as successful and efficient as one another irrespective of which visualization was used. However, there are notable differences in visual attention demanded in each visualization group and how the participants arrived at those performance outcomes. Firstly, across domains, the LIL visualization is shown to have consistently demanded shorter physical travels of the eyes compared to the matrix visualization. This is likely due to their closer displays of the ontological entities that are visualized side-by-side. Secondly, in a less familiar scenario such as the organization of academic conferences (containing ontological concepts that are generally new to the participants in this experiment), it makes little difference to the user irrespective of the visualization used when creating new mappings, considering that all participants exhibited similar gaze behaviors to achieve the same level of success and task speed. Notably, in a more familiar scenario such as the human anatomy domain (with ontological concepts that are generally recognizable and understandable to the participants in this experiment), the LIL indicates less visual effort demanded from the users when creating new mappings. This finding suggests that in more challenging scenarios (be it domain related or otherwise), it is likely that the LIL may be more appropriate to support novice users than the matrix visualization.

While this study focuses on the mapping visualizations of ontology classes, future research may broaden the scope to include visualizations of instance and property associations made between two or more ontologies, as well as investigating effective visualizations for other types of mapping relationships beyond equivalences (e.g., part-of, disjoints, etc.) and one-to-one mappings (e.g., one-to-many). In addition, while there are infinite numbers of scenarios and domains that can also be investigated, further efforts focusing on specific types of mapping scenarios such as mismatches at the language level (e.g., due to syntax, expressiveness, etc.) vs. structure level (e.g., due to differing ontology modeling convention, paradigms, granularity, coverage, etc.) may be useful in identifying effective visual primitives that are helpful in overcoming targeted issues for the average user.

The ontologies and mapping sets used in this study are relatively small in size, as such, we have not investigated how the LIL and matrix visualization would perform in large-scale ontology mappings. There is an opportunity to measure user experience in the context of creating and evaluating large sets of ontology mappings in future studies. Furthermore, though we did not find any statistically significant differences in Computer Science vs. non-Computer Science participants or results across the two domains in this particular study, it may be speculated that distinct user backgrounds, domain knowledge, and personal preferences are potential attributing factors dictating various visual needs, future research focusing on eliciting gaze trends in novice vs. expert user groups in experiments designed specifically to extract differing behaviors from various user groups are needed in the refinement of mapping visualizations.

Finally, empowered with the knowledge of user gaze during an interaction such as how one is searching and processing visual cues in real time, intelligent visualizations with adaptative features can be developed in the advancement of visualizations for the ISW, whereby we can envision adaptive ISW visualization systems (such as predicting user success and failure in real time based on gaze data during user interaction with ontology visualizations to recognize those potential moments to provide visual intervention [52, 53]) suggesting differing visualization techniques and modifying visual primitives in an existing visualization in order to tailor to users' changing needs in a visual scene with the overall goal of improving user success and reducing cognitive efforts.

Supplemental Material Statement: Source code of the ontology pairs, their mapping visualizations, and the user tasks are available at `https://github.com/TheD2Lab/OntoMapVis`*. Source code used to analyze the participants' eye gaze data is available at* `https://github.com/TheD2Lab/Eye.Tracking.Data.Analysis.For.Tobii.2150`*. The raw user data generated from the study is available upon request.*

References

1. Vargas, H., Buil-Aranda, C., Hogan, A., López, C.: RDF Explorer: A Visual SPARQL Query Builder. ISWC 2019. LNCS 11778, 647-663, Springer, Cham. (2019).
2. Toulet, A., Michel, F., Bobasheva, A., Menin, A., Dupré, S., Deboin, M-C., Winckler, M., Tchechmedjiev, A.: ISSA: Generic Pipeline, Knowledge Model and Visualization Tools to Help Scientists Search and Make Sense of a Scientific Archive. ISWC 2022. LNCS 13489, 660-677, Springer, Cham. (2022).
3. Michel, F., Gandon, F., Ah-Kane, V., Bobasheva, A., Cabrio, E., Corby, O., Gazzotti, R., Giboin, A., Marro, S., Mayer, T., Simon, M., Villata, S., Winckler, M.: Covid-on-the-Web: Knowledge Graph and Services to Advance COVID-19 Research. ISWC 2020, LNCS 12507, 294-310, Springer, Cham. (2020).
4. Fu, B., Noy, N. F., Storey, M.-A.,: Eye Tracking the User Experience - An Evaluation of Ontology Visualization Techniques. Semantic Web Journal – Interoperability, Usability, Applicability, 8 (1), 23-41, IOS Press (2017).
5. Fu, B., Noy, N. F., Storey, M.-A.,: Indented Tree or Graph? A Usability Study of Ontology Visualization Techniques in the Context of Class Mapping Evaluation. ISWC 2013, LNCS 8218, 117-134, Springer, Berlin, Heidelberg, (2013).
6. Asmat, M. R. A., Wiens, V., Lohmann, S.: A Comparative User Evaluation on Visual Ontology Modeling using Node-Link Diagrams, VOILA@ISWC 2018, CERU-WS Vol-2187, 25-36 (2018).
7. Karampatakis, S., Bratsas, C., Sváb-Zamazal, O., Filippidis, P., Antoniou, I.: Alignment: A Hybrid, Interactive and Collaborative Ontology and Entity Matching Service. Information, 9, 281. (2018).
8. Dragisic, Z., Ivanova, V., Lambrix, P., Faria, D., Jiménez-Ruiz, E., Pesquita, C.: User Validation in Ontology Alignment. ISWC 2016. LNCS 9981, 200-217, Springer, Cham. (2016).
9. Cruz, I., Stroe, C., Palmonari, M.: Interactive user feedback in ontology matching using signature vectors. In IEE 28th International Conference on Data Engineering, 1321-1324, (2012).
10. Cruz, I., Antonelli, F., Stroe, C.: Agreementmaker: efficient matching for large real-world schemas and ontologies. In Proceedings of the VLDB Endowment, 2(2): 1586-1589 (2009).
11. Cruz, I., Sunna, W., Makar, N., Bathala, S.: A visual tool for ontology alignment to enable geospatial interoperability. Visual Languages and Computing, 18(3): 230-254 (2007).
12. Lanzenberger, M., Sampson, J., Rester, M., Naudet, Y., Latour, T.: Visual ontology alignment for knowledge sharing and reuse. Knowledge Management, 12(6): 102-120 (2008).
13. Faria, D., Martins, C., Nanavaty, A., Oliveira, D., Sowkarthiga, B., Taheri, A., Pesquita, C., Couto, F. M., Cruz, I.: AML results for OAEI 2015. OM@ISWC 2015 (2015).
14. Pesquita, C., Faria, D., Santos, E., Neefs, J., Couto, F. M.: Towards Visualizing the Alignment of Large Biomedical Ontologies. In Proceedings of the International Conference on Data Integration in the Life Sciences, vol 8574, 104-111, Springer, Cham. (2014).
15. Falconer, S., Bull, R., Grammel, L., Storey, M.-A. Creating Visualizations through Ontology Mapping. Proceedings of the International Conference on Complex, Intelligent and Software Intensive Systems, 688-693 (2009).
16. Falconer, S., Storey, M.-A,: A Cognitive Support Framework for Ontology Mapping. In ISWC 2007/ASWC 2007, LNCS 4825, 114-127, Springer, Berlin, Heidelberg (2007).
17. Noy, N., Musen, M.: Algorithm and Tool for Automated Ontology Merging and Align-ment. In AAAI 2000, 450-455, (2000).

18. Aumueller, D., Do, H., Massmann, S., Rahm, E.: Schema and ontology matching with COMA++. Proceedings of the ACM SIGMOD International Conference on Management of Data, 906-908 (2005).
19. Jiménez-Ruiz, E., Grau, B.C., Zhou, Y., Horrocks, I.: Large-scale Interactive Ontology Matching: Algorithms and Implementation. Frontiers in Artificial Intelligence and Applications 242: 444-449 (2012).
20. Lambrix, P., Kaliyaperumal, R.: A Session-Based Approach for Aligning Large Ontologies. ESWC 2013. Lecture Notes in Computer Science, LNCS 7882, 46-60, Springer, Berlin, Heidelberg (2013).
21. Lambrix, P., Tan. H., SAMBO - a system for aligning and merging biomedical ontologies. Journal of Web Semantics, 4(3): 196-206 (2006).
22. Ivanova, V., Lambrix, P.: A unified approach for aligning taxonomies and debugging taxonomies and their alignments. In ESWC 2013, LNCS 7882, 1-15, Springer, Berlin, Heidelberg (2013).
23. Neto, L.E.T., Vidal, V.M.P., Casanova, M.A., Monteiro, J.M.: R2RML by Assertion: A Semi-automatic Tool for Generating Customised R2RML Mappings. ESWC 2013 Satellite Events. LNCS 7955, 248-252, Springer, Berlin, Heidelberg (2013).
24. Nasir, M., Hoeber, O., Evermann, J.: Supporting Ontology Alignment Tasks with Edge Bundling. In Proceedings of the 13th International Conference on Knowledge Management and Knowledge Technologies, Article 11, 1-8, ACM (2013).
25. Voyloshnikova, E., Fu, B., Grammel, L., Storey, M.-A.: BioMixer: Visualizing Mappings of Biomedical Ontologies. The 3rd International Conference on Biomedical Ontology, CEUR-WS Vol. 897 (2012).
26. Dang, T., Franz, N., Ludäscher, B., Forbes, A. G.: ProvenanceMatrix: A Visualization Tool for Multi-taxonomy Alignments, Proceedings of the International Workshop on Visualizations and User Interfaces for Ontologies and Linked Data, co-located with ISWC 2015, CEUR-WS Vol-1456, 13-24 (2015).
27. Ivanova, V., Bach, B., Pietriga, E., Lambrix, P.: Alignment Cubes: Towards Interactive Visual Exploration and Evaluation of Multiple Ontology Alignments. ISWC 2017, LNCS 10587, 400-417, Springer, Cham. (2017).
28. Severo, B., Trojahn, C., Vieira, R.: VOAR 3.0: A Configurable Environment for Manipulating Multiple Ontology Alignments. Proceedings of the ISWC 2017 Posters & Demonstrations and Industry Tracks, CEUR-WS Vol-1963 (2017).
29. Aurisano, J., Nanavaty, A., Cruz, I., Visual Analytics for Ontology Matching Using Multi-linked Views, Proceedings of the International Workshop on Visualizations and User Interfaces for Ontologies and Linked Data, co-located with ISWC 2015, CEUR-WS Vol-1456, 25-36 (2015).
30. Fu, B., Grammel, L., Storey, M.-A.: BioMixer: A Web-based Collaborative Ontology Visualization Tool. In Proceedings of the 3rd International Conference on Biomedical Ontology, CEUR-WS Vol.897, ISSN 1613-0073 (2012).
31. Junior, A.C., Debruyne, C., O'Sullivan, D.: An Editor that Uses a Block Metaphor for Representing Semantic Mappings in Linked Data. ESWC 2018 Satellite Events. LNCS 11155, 28-33, Springer, Cham. (2018).
32. Li, Y., Stroe, C., Cruz, I.: Interactive Visualization of Large Ontology Matching Results, VOILA@ISWC 2015, CEUR-WS Vol-1456, 37-48 (2015).
33. Heyvaert, P., Dimou, A., De Meester, B., Seymoens, T., Herregodts, A.-L., Verborgh, R., Schuurman, D., Mannens, E.: Specification and implementation of mapping rule visualization and editing: MapVOWL and the RMLEditor, Journal of Web Semantics, Vol. 49; 31-50 (2018).

34. Sicilia, A., Nemirovski, G., Nolle, A.: Map-On: A web-based editor for visual ontology mapping. Semantic Web – Interoperability, Usability, Applicability. 8. 969-980 (2017).
35. Heyvaert, P., Dimou, A., Herregodts, A.-L., Verborgh, R., Schuurman, D., Mannens, E., Van de Walle, R., RMLEditor: A Graph-based Mapping Editor for Linked Data Mappings, ESWC 2016, LNCS 9678, 709-723, Springer, Cham. (2016).
36. Lembo, D., Rosati, R., Ruzzi, M., Savo, D. F., Tocci, E.: Visualization and Management of Mappings in Ontology-based Data Access. Informal Proceedings of the 27th International Workshop on Description Logics, CERU Vol-1193, 595-607 (2014).
37. Falconer, S.M., Callendar, C., Storey, M.-A., A Visualization Service for the Semantic Web. Knowledge Engineering and Management by the Masses. EKAW 2010. LNCS 6317, 554-564, Springer, Berlin, Heidelberg (2010).
38. Liu, G. T., Volpe, N. J., Galetta, S.: Neuro-Ophthalmology Diagnosis and Management, 3rd edition, Chapter 13, ISBN: 9780323340441, Elsevier (2018).
39. Spector R.H.: The Pupils. Clinical Methods: The History, Physical, and Laboratory Examinations, 3rd edition, Chapter 58, Boston: Butterworths (1990).
40. Goldberg, J. H., Kotval, X. P.: Computer interface evaluation using eye movements: methods and constructs, International Journal of Industrial Ergonomics, 24(6), 631-645 (1999).
41. Ivanova, V., Lambrix, P., Åberg, J., Requirements for and Evaluation of User Support for Large-Scale Ontology Alignment, The 12th European Semantic Web Conference (ESWC 2015), LNCS 9088, 3-20 (2015).
42. Fu, B., Steichen, B., McBride, A., Tumbling to Succeed: A Predictive Analysis of User Success in Interactive Ontology Visualization. In Proceedings of the 10th International Conference on Web Intelligence, Mining and Semantics (WIMS 2020), 78-87, ACM (2020).
43. Fu, B., Steichen, B., Impending Success or Failure? An Investigation of Gaze-Based User Predictions During Interaction with Ontology Visualizations. In Proceedings of the International Conference on Advanced Visual Interfaces (AVI 2022), Article No.7, 1-9, ACM (2022).
44. Y. Yang, M. Wybrow, Y.-F. Li, T. Czauderna, Y. He, OntoPlot: A Novel Visualisation for Non-hierarchical Associations in Large Ontologies, in IEEE Transactions on Visualization and Computer Graphics, 26 (1), 1140-1150 (2020).
45. Saket, B., Simonetto, P., Kobourov, S., Borner, K., Node, Node-Link, and Node-Link-Group Diagrams: An Evaluation. IEEE Transactions on Visualization and Computer Graphics, 20 (12), 2231-2240 (2014).
46. Netzel, R., Burch M., Weiskopf, D., Comparative Eye Tracking Study on Node-Link Visualizations of Trajectories, in IEEE Transactions on Visualization & Computer Graphics, 20 (12), 2221-2230 (2014).
47. Jianu, R., Rusu, A., Hu, Y., Taggart, D., How to Display Group Information on Node-Link Diagrams: An Evaluation. IEEE Transactions on Visualization and Computer Graphics, 20 (11), 1530-1541 (2014).
48. Büschel, W., Vogt, S., Dachselt, R., Augmented Reality Graph Visualizations: Investigation of Visual Styles in 3D Node-Link Diagrams. IEEE Computer Graphics and Applications, 39 (3), 29-40 (2019).
49. Pan, J., Chen, W., Zhao., X., Zhou, S., Zeng W., Zhu, M., Chen., J., Fu, S., Wu., Y., Exemplar-based Layout Fine-tuning for Node-link Diagrams, in IEEE Transactions on Visualization and Computer Graphics, 27 (2), 1655-1665 (2021).
50. Bartolomeo, S. di, Riedewald, M., Gatterbauer W., Dunne, C., STRATISFIMAL LAYOUT: A modular optimization model for laying out layered node-link network visualizations, in IEEE Transactions on Visualization and Computer Graphics, 28 (1), 324-334 (2022).

51. Shimabukuro, M., Zipf, J., El-Assady M., Collins, C., H-Matrix: Hierarchical Matrix for Visual Analysis of Cross-Linguistic Features in Large Learner Corpora, 2019 IEEE Visualization Conference, 61-65 (2019).
52. Beusekom, N. van, Meulemans W., Speckmann, B., Simultaneous Matrix Orderings for Graph Collections, in IEEE Transactions on Visualization and Computer Graphics, 28 (1), 1-10 (2022).
53. Abdelaal, M., Schiele, N. D., Angerbauer, K., Kurzhals, K., Sedlmair, M., Weiskopf, D., Comparative Evaluation of Bipartite, Node-Link, and Matrix-Based Network Representations, in IEEE Transactions on Visualization & Computer Graphics, 29 (1), 896-906 (2023).