MetaSEMAP: Usability Evaluation for Metadata Annotation Across the Mapping Lifecycle*

Sarah Alzahrani[1,2,*,†], Declan O'Sullivan[3,†]

[1]School of Computer Science and Statistics, Trinity College Dublin, Ireland
[2] Imam Mohammad Ibn Saud Islamic University (IMSIU), Saudi Arabia
[3]ADAPT Center, Trinity College Dublin, Ireland

Abstract

Comprehensive and well-structured metadata annotation and documentation are essential for supporting the reuse and maintenance of declarative mappings throughout their lifecycle. This paper presents a usability evaluation of MetaSEMAP, a tool designed to facilitate the annotation of declarative mappings, including uplift mappings, ontology alignment, and interlinking. While MetaSEMAP and its underlying metadata model are still under active development, the focus of this work is on evaluating how users interact with the tool and interpret the metadata concepts it introduces. The evaluation investigates MetaSEMAP's ability to support metadata annotation using real-world scenarios such as the Virtual Record Treasury of Ireland. Participants provided feedback on the tool's usability and their preferences for metadata representation, including RDF-star and Named Graph. The results reveal both strengths and areas for improvement, offering valuable insights for the development of more effective mapping annotation tools. This work supports efforts to improve interoperability and sustainability in mapping practices, with alignment to FAIR principles as a longer-term goal.

Keywords

Metadata Annotation, Mapping Annotation, Declarative Mappings, Mapping Lifecycle, RDF-star, Named Graph, FAIR Principles, Usability

1. Introduction

Declarative mappings are essential for interoperability across diverse data ecosystems, helping to resolve issues related to semantic heterogeneity and varying data structures [1]. Such diverse data ecosystems range from the deployment of semantic web technologies into infrastructure-type approaches, e.g., event-based networking [2], right through to complex multi-domain applications, e.g., building information management [3]. These mappings typically fall into three categories: ontology alignment, uplift mapping, and interlinking. However, while declarative mappings provide a useful means of linking and transforming data, challenges arise in managing, understanding, and reusing these mappings. Their lifecycles are complex, involving multiple stakeholders, evolving formats, and shifting requirements, all of which complicate their reuse and long-term sustainability [1].

The importance of metadata in addressing these challenges cannot be overstated. Metadata provides the critical context needed to understand the purpose, domain, contributors, and technical characteristics of declarative mappings. However, existing approaches often lack standardized metadata, offer limited queryability, and provide only partial coverage of the mapping lifecycle. This makes essential tasks such as reuse, maintenance, versioning, and governance difficult to carry out effectively.

To address these limitations, we introduce MetaSEMAP, a metadata-driven tool designed to help users annotate declarative mappings in a structured and consistent way. MetaSEMAP is built upon a metadata model that formalizes the lifecycle of mapping development, comprising five distinct phases: Analysis,

First International Workshop on Users and Knowledge Graphs (UKG), co-located with SEMANTiCS'25: International Conference on Semantic Systems, September 3–5, 2025, Vienna, Austria

*Corresponding author.
†These authors contributed equally.
salzahra@tcd.ie (S. Alzahrani); declan.osullivan@tcd.ie (D. O'Sullivan)
https://www.scss.tcd.ie/Declan.OSullivan/ (D. O'Sullivan)
0000-0003-2169-7006 (S. Alzahrani); 0000-0001-7770-2117 (D. O'Sullivan)

Design, Development, Testing, and Maintenance. Each phase is associated with specific metadata fields relevant to its role in the lifecycle. The proposed metadata model not only improves consistency and completeness but also enables automated validation and machine-readable documentation, supporting the FAIR principles, especially in terms of interoperability and reusability.

MetaSEMAP provides a simple, web-based interface through which users can annotate uplift mapping, ontology alignment, or interlinking mappings using a guided form based on the proposed metadata model. The tool supports both RDF-star and Named Graph representations to offer flexibility in metadata expression. To evaluate the effectiveness of MetaSEMAP, we conducted a usability study with 50 participants, assessing perceived usability, task completion time, and collecting qualitative feedback on the user experience.

While MetaSEMAP and the underlying metadata model are still under development, this paper does not aim to evaluate their completeness or technical implementation. Instead, the primary contribution is an exploratory usability evaluation, focusing on how users engage with the tool and interpret the metadata concepts it introduces. The results offer insights into user needs and guide future improvements to both the tool and the model.

The remainder of the paper is structured as follows. Section 2 reviews related work on metadata for declarative mappings and annotation tools. Section 3 introduces the different types of declarative mappings and outlines the proposed mapping lifecycle. Section 4 presents the metadata model that supports annotation across lifecycle phases. Section 5 describes the MetaSEMAP tool and its implementation. Section 6 outlines the evaluation design and the metrics used to assess usability and performance. Section 7 presents the results and insights from the study, including quantitative findings, qualitative feedback, and representation preferences. Section 8 discusses the broader implications for metadata design and usability. Finally, Section 9 concludes the paper with a summary of key findings and future directions.

2. Related Work

Efforts to manage and share metadata for declarative mappings often focus on specific mapping types or particular stages of the mapping lifecycle, rather than offering comprehensive solutions that support documentation, traceability, and reusability across all phases.

For ontology alignment, Thomas et al. proposed OM2R, a metadata model capturing the full lifecycle of ontology mappings to support their reuse and management [4]. While effective in the alignment context, it was not designed to generalize across other mapping types such as interlinking or uplift. The Simple Standard for Sharing Ontological Mappings (SSSOM) is a community-developed specification for representing ontology and terminology alignments [5]. It defines a set of rich metadata fields such as predicate, confidence score, and justification to describe individual mapping assertions. While SSSOM effectively supports documenting the results of the alignment process, it offers limited coverage of metadata related to the broader mapping lifecycle, such as stakeholder roles, tool usage workflows, or quality assurance procedures. Although some of these aspects may be accommodated through community-driven extensions or through integration with complementary vocabularies such as PROV-O, they are not explicitly addressed in the core SSSOM specification.

In the interlinking domain, tools like Silk [6] and LIMES [7] support link discovery across datasets. Although they are effective for generating RDF links, these tools do not incorporate features for documenting metadata about the linking methodology, stakeholder decisions, or validation strategies.

A recent initiative, FAIR-IMPACT, has advocated for improving the FAIRness of mapping documentation. It promotes the use of structured metadata, persistent identifiers, and semantic enrichment to enhance the findability and reusability of mapping artefacts across their lifecycle [8, 9]. As part of this effort, FAIR-IMPACT recommends the use of the Simple Standard for Sharing Ontological Mappings (SSSOM) and has proposed extensions to support additional metadata fields [8]. While these extensions improve the ability to describe mappings with provenance and justifications, the focus remains largely on output-level assertions. The proposed updates still do not fully capture the broader mapping lifecycle, such as stakeholder involvement, design rationale, tool usage, or validation steps. As a result, manual annotation support and integration into end-to-end mapping workflows remain limited.

Complementing these initiatives, Toledo et al. proposed RMLdoc, a tool that generates human-readable documentation for RML mapping files [10]. While effective in increasing the transparency of uplift mappings, RMLdoc does not extend to other mapping types such as ontology alignment or interlinking, nor does it support full lifecycle documentation.

MetaSEMAP addresses these limitations by providing a unified, implementation-level metadata model and annotation tool that supports ontology alignment, interlinking, and uplift mappings. It enables users to document activities across all phases of the mapping lifecycle, such as stakeholder involvement, design decisions, testing processes, and maintenance practices.

3. Declarative Mappings: Types and Lifecycle

Declarative mappings are central to achieving interoperability on the Semantic Web and, as previously noted, can be grouped into three categories: ontology alignment, interlinking, and uplift mapping. These processes support the transformation, linking, and semantic integration of heterogeneous data. Figure 1 provides an illustrative representation of these mapping types.

This study addresses all three. Uplift mappings used to convert data into RDF, e.g. [11]. Interlinking identifies relationships between entities across datasets [12], while ontology alignment establishes correspondences between concepts in different ontologies to enable semantic interoperability [13].

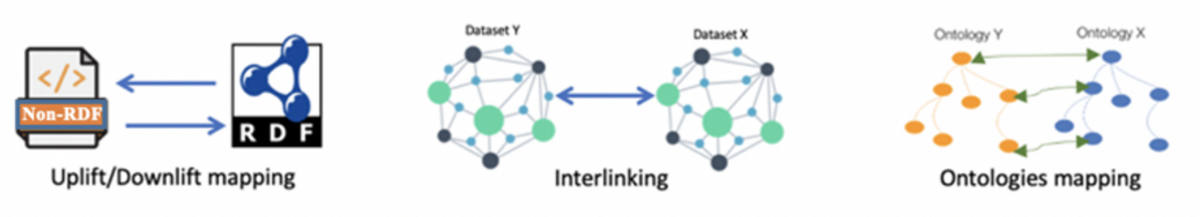

Figure 1: The three categories of mappings.

To structure these activities, we propose a mapping lifecycle inspired by prior models [14, 4, 1], with the addition of a dedicated testing phase. The resulting lifecycle includes five stages: Analysis, Design, Development, Testing, and Maintenance. This framework is designed to be applicable across all mapping types, with phase relevance varying by context.

4. Metadata Model

Building on the proposed mapping lifecycle, we developed a metadata model to document key activities across all phases. Metadata fields were designed to capture decisions and actions specific to each phase of the lifecycle [15]. Table 1 summarizes the structure of the model, aligning metadata fields with corresponding lifecycle stages.

To validate the model's structure and the relevance of its fields, a survey was conducted with participants from the Semantic Web and Linked Data community [16]. The goal was to assess the applicability of the proposed metadata for two mapping-related tasks. Results showed strong agreement with the selected fields, and participants also suggested additional useful metadata elements.

In the initial tool implementation (see next section), the model reused widely adopted vocabularies such as FOAF (for stakeholder roles), and DCMI (for metadata about inputs like source, format, and creator), along with a custom namespace for domain-specific fields (e.g., requirements, tools).

A more complete ontology is under development[1], incorporating standard vocabularies including

[1]Project repository: https://github.com/SarahAlzahranitcd/MetaSEMAP-Metadata

Table 1
Proposed Metadata Fields Aligned with Mapping Lifecycle Phases

Lifecycle Phase	Key Metadata Fields
Analysis	Stakeholder details (URI, name, background, role, organization); mapping purpose (requirements, type, domain, assumptions, technical needs, risks); input descriptions (URI, name, source, type, creator, format)
Design	Final design decisions, justifications, and anticipated quality metrics
Development	Mapping details (URI, name, start/end date, tools used, mapping method, algorithm, format)
Testing	Testing metadata (type, timestamp, results)
Maintenance	Versioning information (publisher name, source, version number, date)

PROV-O [17], DQV [18], and more recent vocabularies such as MQV [19] to formally describe provenance, quality assessment, and validation metadata. A detailed specification describing each metadata field and its usage is also provided in the same repository. This ontology will be integrated into the next version of the tool and evaluated with knowledge engineers during the second development phase of this project.

5. The Initial MetaSEMAP Tool

MetaSEMAP evolved from MetaMap [20], initially focused on uplift mappings, and was renamed to reflect broader support and avoid naming conflicts. MetaSEMAP is a web tool that supports structured annotation of mappings. The homepage (Figure 2) offers two workflows: annotate new or reuse existing mappings.

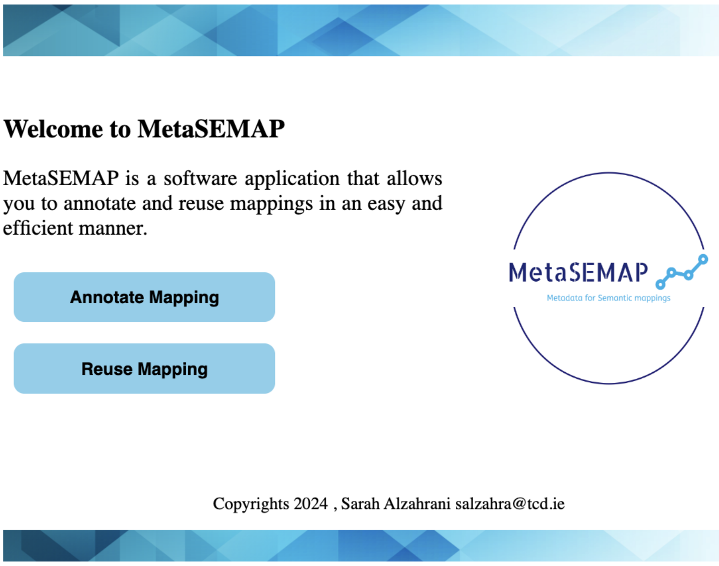

Figure 2: MetaSEMAP main interface

The annotation process in MetaSEMAP unfolds through four key steps. First,Users begin by uploading an RDF mapping file (Figure 3).

Next, users complete structured metadata fields aligned with the mapping lifecycle. These fields prompt users to document key aspects of the mapping, such as its purpose, input file, development method, and associated stakeholders. As shown in Figure 4, the form is organized into sections that correspond to each lifecycle phase. Some fields are scenario-specific and optional, while others are required to ensure metadata completeness and consistency.

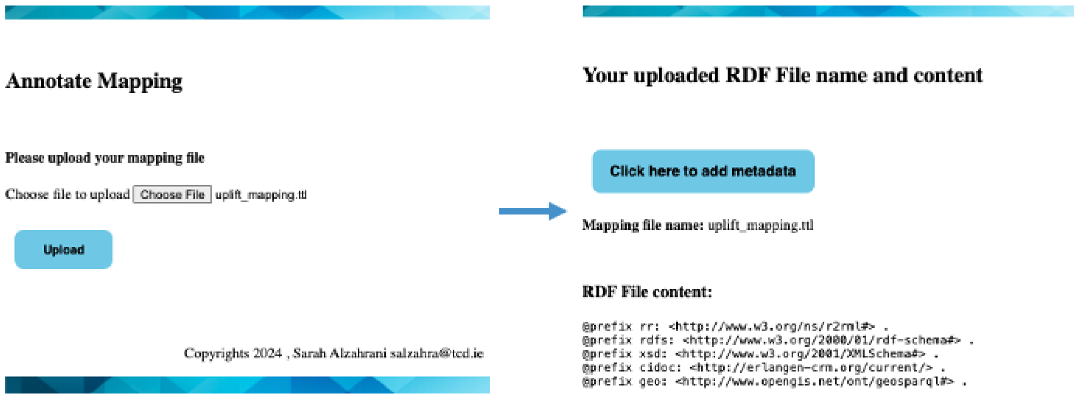

Figure 3: MetaSEMAP Step 1: Upload and preview of mapping file before annotation.

To assist users in understanding each metadata category, contextual help was provided via tooltips. As illustrated in Figure 5, hovering over a question mark icon revealed a short explanation about the purpose and expected input for the associated metadata section. This feature was especially useful for users and helped ensure more consistent and accurate annotations.

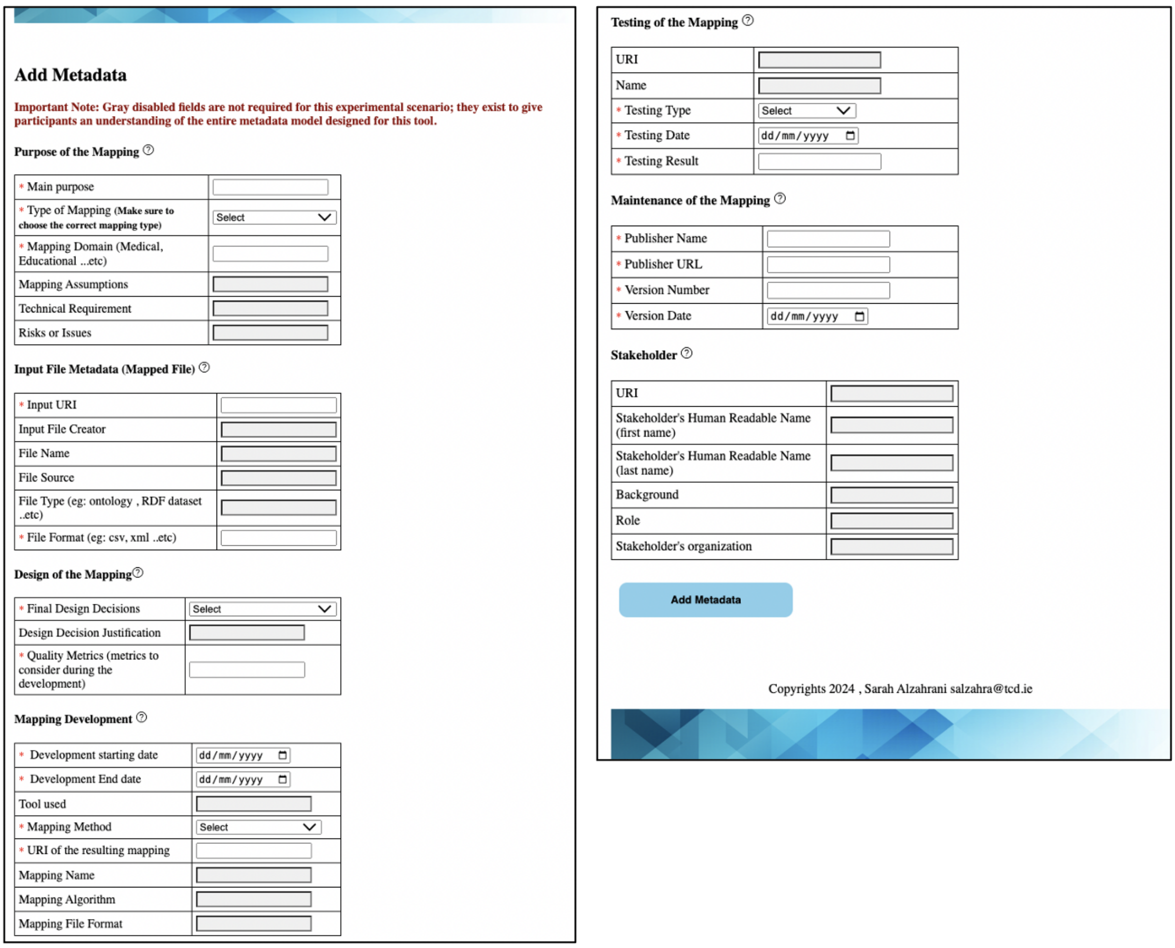

Figure 4: MetaSEMAP metadata annotation interface.

After completing the metadata form, users are presented with a review screen (Figure 6) summarizing all the fields they have entered. This confirmation step was designed to promote careful review and reduce input errors before submission. Once the confirmation is done, the tool generates machine-

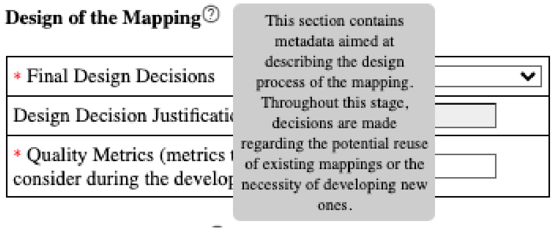

Figure 5: Tooltip example: Metadata field explanations appeared when hovering over category labels.

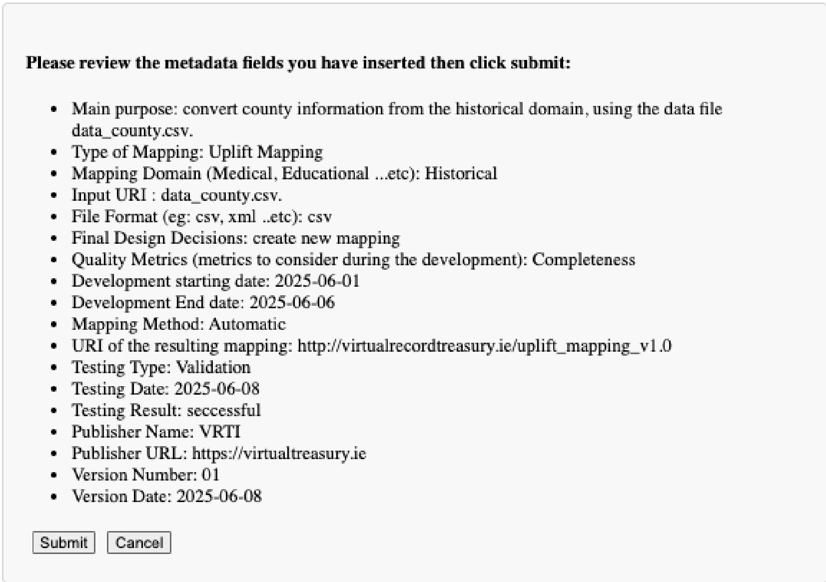

Figure 6: Review screen.

readable metadata in either RDF-star or Named Graph. Users can preview and download the output for further use or publication. Figure 7 illustrates the final interface that users see after annotation, showing the confirmation message and access to RDF-star and Named Graph outputs.

6. Evaluation Design

To evaluate usability and the effectiveness of the metadata model for declarative mappings, we designed a controlled user evaluation of the initial MetaSEMAP implementation, guided by the following research questions:

- **RQ1:** How usable is the MetaSEMAP tool for annotating different types of declarative mappings?
- **RQ2:** How effectively can users complete metadata annotation tasks using the proposed metadata model?
- **RQ3:** Which metadata representation (Named Graph or RDF-star) do users find easier to read and interpret?

6.1. Participants

Fifty MSc students enrolled in a knowledge and data engineering course participated in the user evaluation study. They had introductory knowledge of mapping processes but varied in technical expertise. The study was approved by the institutional ethics committee, and participation was voluntary, anonymous, and online. Participants used the public MetaSEMAP website (MetaSEMAP Experiment) asynchronously.

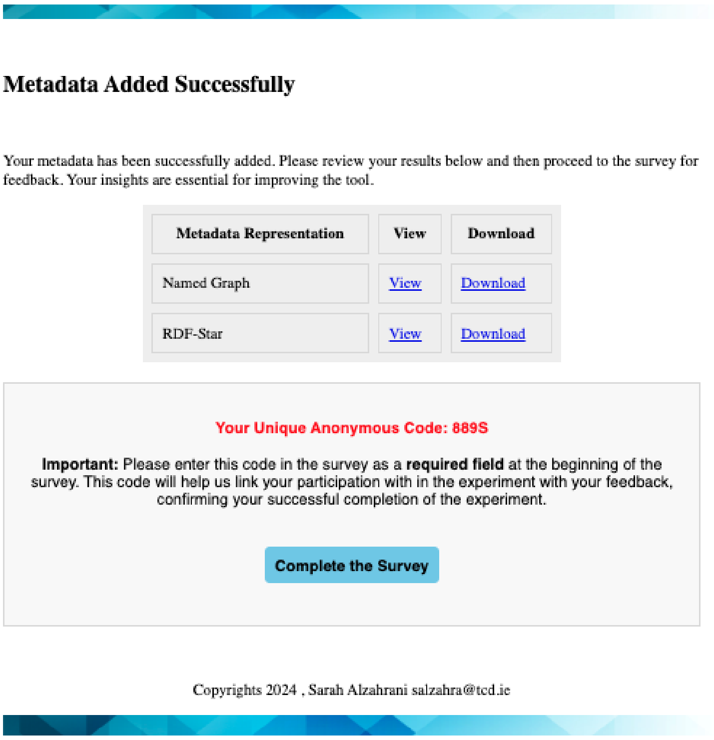

Figure 7: MetaSEMAP interface showing post-annotation confirmation.

6.2. Tasks and Scenarios

Each participant was randomly assigned one of three scenarios representing a different declarative mapping type: uplift mapping, ontology alignment, or interlinking. These scenarios were grounded in the Virtual Record Treasury of Ireland (VRTI) knowledge graph [21]. VRTI is a historical knowledge graph designed to digitally reconstruct Ireland's archival heritage, offering structured, domain-rich data suitable for metadata annotation experiments. Each scenario included a mapping file and contextual metadata. Participants were instructed to:

1. Read the scenario and download the provided mapping file.
2. Upload the file to MetaSEMAP.
3. Annotate the mapping using the tool's metadata fields, referencing the scenario.
4. Review the generated metadata in both RDF-star and Named Graph representations.
5. Complete a survey reflecting on usability and representation preference.

The scenario summaries are listed in Table 2.

6.3. Evaluation Metrics

To assess the study goals, we employed a mix of quantitative and qualitative metrics. These included:

- *System usability and satisfaction* measured via the Post-Study System Usability Questionnaire (PSSUQ), based on a 1–7 Likert scale survey.[2]
- *Task completion time* tracked from start to metadata submission.
- *Representation preference* captured through survey feedback.
- *Thematic feedback* from open-ended responses.

The anonymized survey responses and the mapping files annotated by participants during the study are available in the MetaSEMAP project repository 1[0].

[2]PSSUQ Survey: https://docs.google.com/forms/d/e/1FAIpQLSdRhBPxzCQwrQyi3zmMPdP-EeON7zqFAG5tR6aoSd0UnXuMXQ/viewform

Table 2
Summary of Scenarios in the Usability Experiment

Mapping Type	Summary of Scenario
Uplift Mapping	Convert county information from historical data into RDF triples. The data file `data_county.csv` was used to map county IDs, names, and geographic details to enrich the VRTI Knowledge Graph.
Ontology Alignment	Align person entities (e.g., historical figures) between the VRTI Knowledge Graph and external datasets. Features such as relationships, affiliations, and roles were matched using the file `person_alignment.rdf`.
Interlinking	Connect Irish historical figures to their corresponding Wikidata entries using `owl:sameAs` relationships. The project enriched VRTI records with biographical and contextual data using the file `interlinking.rdf`.

7. Results and Discussion

7.1. System Usability and User Satisfaction

System usability was assessed using a modified version of the Post-Study System Usability Questionnaire (PSSUQ), adapted to include a subset of questions inspired by the System Usability Scale (SUS). The final instrument consisted of 14 items covering system usefulness, information quality, and interface quality, rated on a 1–7 Likert scale (lower is better). Negatively phrased questions (e.g., Q8, Q9, Q11, Q13) were reverse-coded prior to analysis, in line with standard usability survey practices. Out of 50 participants, 46 completed the questionnaire. The overall average score was 2.49, indicating good usability. Around 70% rated the tool positively, citing ease of use and effective support for metadata annotation.

System Usefulness scored 2.4, reflecting user confidence and successful error recovery. Information Quality received a score of 3.4, suggesting moderate satisfaction, particularly with error message clarity. Interface Quality was rated at 2.3, pointing to a generally intuitive and consistent layout. Table 3 summarizes these scores and their associated questions.

Table 3
Average Scores by PSSUQ Category and Associated Questions

Category	Avg. Score	Questions Used
System Usefulness	2.4	Q4, Q5, Q6, Q7
Information Quality	3.4	Q1, Q2, Q8, Q9
Interface Quality	2.3	Q3, Q10, Q11, Q13
Overall Usability	2.49	All of the above + Q12, Q14

7.2. Task Completion Time

Participants completed annotation tasks based on one of three mapping types. The initial average completion time was 44 minutes. To ensure the results reflected realistic task durations, outliers were excluded using empirical thresholds: submissions under 5 minutes (likely rushed or incomplete) and those exceeding 200 minutes (suggesting prolonged interruptions or technical issues). After removing these outliers, the adjusted averages were: 32 minutes for Uplift Mapping, 39 minutes for Ontology Alignment, and 33 minutes for Interlinking.

7.3. Preferred Metadata Representation

Participants reviewed both RDF-star and Named Graph outputs generated from their annotations. A majority (74.4%) expressed a preference for Named Graphs. In contrast, 25.6% preferred RDF-star.

One potential influencing factor was the representation setup: RDF-star annotations were included in the same file as the mapping (see Listing 1), and participants may have found them harder to interpret without prior exposure to the syntax. Named Graph metadata, on the other hand, was presented in a separate file (see Listing 2). This physical separation may have enhanced the perceived readability and organization of Named Graphs. Additionally, inconsistencies in prefix usage between the two representations may have further affected participant preferences. These aspects will be addressed in future iterations to ensure more balanced comparison conditions.

Listings 1 and 2 show an example of RDF-star and Named Graph annotations for an uplift mapping, where the `rr:TriplesMap` triple was the subject of annotation. For ontology alignment, RDF-star annotations targeted the `align:Alignment` instance URI, while for interlinking, the RDF-star subject was the SPARQL query resource used to perform linking between VRTI and Wikidata. In all cases, the Named Graph representation followed a consistent pattern, where metadata was written in a separate file using a `metag:subject` placeholder and reused the same structured fields as RDF-star.

```
@prefix b2022: <https://ont.virtualtreasury.ie/ontology#> .
@prefix cidoc: <http://erlangen-crm.org/current/> .
@prefix geo: <http://www.opengis.net/ont/geosparql#> .
@prefix ql: <http://semweb.mmlab.be/ns/ql#> .
@prefix rml: <http://semweb.mmlab.be/ns/rml#> .
@prefix rr: <http://www.w3.org/ns/r2rml#> .
@prefix xsd: <http://www.w3.org/2001/XMLSchema#> .

<http://example.com/ns##COUNTY> a rr:TriplesMap ;
   rml:logicalSource [
      rml:query "SELECT␣VRTI_ID,␣NAME_TAG␣FROM␣'data_county.csv'" ;
      rml:referenceFormulation ql:CSV ;
      rml:source "data_county.csv"] ;
   rr:subjectMap [
      rr:template "https://kg.virtualtreasury.ie/place/county/{NAME_TAG}/{VRTI_ID}" ;
      rr:class cidoc:E53_Place] ;
   rr:predicateObjectMap [
      rr:predicate geo:hasCentroid ;
      rr:objectMap [ rr:template "https://kg.virtualtreasury.ie/place/county/centroid/{NAME_TAG}/{
          VRTI_ID}" ]] .

<< <http://example.com/ns##COUNTY> <http://www.w3.org/1999/02/22-rdf-syntax-ns#type> rr:TriplesMap >>

   <http://example.com/metag/purpose> "convert␣county␣information␣from␣the␣historical␣domain,␣using␣
       the␣data␣file␣data_county.csv." ;
   <http://example.com/metag/mappingType> "Uplift␣Mapping" ;
   <http://example.com/metag/mappingDomain> "Historical" ;
   <http://example.com/metag/startDate> "2025-06-01" ;
   <http://example.com/metag/endDate> "2025-06-06" ;
   <http://example.com/metag/mappingMethod> "Automatic" ;
   <http://example.com/metag/testingType> "Validation" ;
   <http://example.com/metag/testingResult> "seccessful" ;
   <http://example.com/metag/versionNumber> "01" ;
   <http://example.com/metag/versionDateTime> "2025-06-08" ;
   <http://xmlns.com/foaf/0.1/organization> "VRTI" .
```

Listing 1: RDF-star Annotation with Mapping Statement (shortened)

```
@prefix dcmi: <http://purl.org/dc/terms/> .
@prefix foaf: <http://xmlns.com/foaf/0.1/> .
@prefix metag: <http://example.com/metag/> .
@prefix xsd: <http://www.w3.org/2001/XMLSchema#> .

metag:subject
   metag:purpose "convert␣county␣information␣from␣the␣historical␣domain,␣using␣the␣data␣file␣
       data_county.csv." ;
```

```
    metag:mappingType "Uplift␣Mapping" ;
    metag:mappingDomain "Historical" ;
    metag:startDate "2025-06-01"^^xsd:date ;
    metag:endDate "2025-06-06"^^xsd:date ;
    metag:mappingMethod "Automatic" ;
    metag:testingType "Validation" ;
    metag:testingResult "seccessful" ;
    metag:versionNumber "01" ;
    metag:versionDateTime "2025-06-08"^^xsd:date ;
    foaf:organization "VRTI" .
```

Listing 2: Named Graph Annotation (shortened)

Note: These examples are shortened to highlight the structure and differences between RDF-star and Named Graph metadata representations. Complete examples are available in the project repository 1[0].

7.4. Metadata Annotation Observations

In addition to usability metrics, we conducted a light-touch review of the metadata submissions to observe common patterns and challenges in annotation. Participants generally completed structured fields (e.g., start dates, URIs, publisher names) with high accuracy, suggesting that these fields were intuitive and well-supported by the interface. However, free-text fields such as *purpose* and *qualityMetrics* exhibited greater variability in both content and specificity.

Some responses conveyed the intended meaning, but lacked precision or used overly general terms. For example, several participants filled in the *mapping purpose* field with generic phrases like 'to describe the data' or 'link information', rather than specifying the actual goal of the project or the user need that motivated the mapping. More conceptual fields, such as distinguishing between manual and automatic mapping methods, also introduced occasional ambiguity.

These patterns indicate that while the tool successfully supports basic annotation tasks, certain metadata elements may benefit from additional in-tool guidance. Features such as example values, tooltips, or predefined templates could help users provide more precise and consistent entries, particularly in scenarios involving complex or abstract metadata concepts.

7.5. Qualitative Feedback

Open-ended responses were analyzed to extract key themes:

- **Ease of Use and Simplicity:** MetaSEMAP was generally described as intuitive and non-technical.
- **Error Handling:** Users requested clearer error messages and better field validation, especially for URIs.
- **Field Guidance:** Participants asked for more in-tool help (e.g., tooltips or explanations) to understand metadata fields.
- **Interface Improvements:** Suggestions included larger input fields, better text wrapping, and persistent scenario visibility.
- **Additional Features:** Users proposed templates, collaboration support, and enhanced consistency checks.

User feedback emphasized priorities for improvement, such as enhanced error messages, better field-level guidance, and improved layout design. These findings will inform future enhancements of MetaSEMAP to support both usability and clarity in metadata annotation workflows.

8. Discussion

The evaluation results show that MetaSEMAP is generally usable across different mapping types, with high scores for system usefulness and interface quality. However, areas such as error message clarity and guidance for free-text fields remain improvement priorities. Participants strongly preferred Named Graphs over RDF-star, which appears to be influenced by presentation structure and familiarity. Named Graphs were delivered in a separate file and followed consistent prefixing, which likely enhanced readability and perceived structure. The current annotation strategy focused on a project-level scope, (e.g., alignment file, SPARQL query, or RML mapping). Although this proved sufficient in our controlled scenarios, it may be less reliable or too coarse-grained in larger use cases, particularly when mappings involve many modular components. Future work will explore more fine-grained annotation. In addition to usability, the collected metadata can support mapping reuse, quality assurance, and contextual understanding by other users. Fields such as purpose, assumptions, and stakeholder roles help assess whether a mapping suits a new use case, while testing metadata and versioning support maintenance and trust. Although some metadata elements require effort to complete, they provide long-term value in ensuring mappings are interpretable, reusable, and aligned with FAIR principles. Overall, the evaluation confirms the value of lifecycle-aware metadata annotation and highlights practical lessons for designing usable tooling in this space.

9. Conclusion and Next Steps

This paper's primary contribution is the usability evaluation of MetaSEMAP, rather than a full validation of the metadata model or complete tool functionality, both of which are still under development. Nevertheless, the feedback offers concrete guidance for improving metadata design, user support, and metadata representation in future iterations of the tool.

An open question emerging from this work is whether structured metadata remains necessary when large language models are able to generate or infer declarative mappings independently. On the other hand, incorporating contextual metadata about the mapping project such as domain, purpose and assumptions into LLM prompts may improve the quality, relevance, and explainability of the generated mappings.

Future work will investigate how structured metadata, when used as contextual input in prompt-based mapping generation, affects the quality and reliability of LLM outputs. This direction aims to assess how human-in-the-loop metadata practices and AI-assisted approaches can work together to support sustainable knowledge graph construction.

Acknowledgments

The first author acknowledges financial support from Al Imam Mohammad Ibn Saud Islamic University and the Saudi government — represented by the Royal Embassy of Saudi Arabia — Cultural Bureau in Dublin. The second author is partially supported by the SFI ADAPT Research Centre (grant number 13/RC/2106_P2).

Declaration on Generative AI

During the preparation of this work, the author used ChatGPT (GPT-4) and Grammarly for the purposes of grammar improvements. All AI-generated content was thoroughly reviewed, edited, and validated by the author, who takes full responsibility for the final manuscript and all its content.

References

[1] C. Debruyne, B. Walshe, D. O'Sullivan, Towards a project centric metadata model and lifecycle for ontology mapping governance, in: Proc. 17th Int. Conf. on Info. Integration and Web-based Applications & Services (iiWAS '15), ACM, 2015, p. Article 50. doi:10.1145/2837185.2837201.

[2] J. Keeney, D. Roblek, D. Jones, D. Lewis, D. O'Sullivan, Extending siena to support more expressive and flexible subscriptions, in: Proceedings of the 2nd International Conference on Distributed Event-Based Systems (DEBS), 2008.

[3] K. McGlinn, R. Brennan, C. Debruyne, A. Meehan, L. McNerney, E. Clinton, P. Kelly, D. O'Sullivan, Publishing authoritative geospatial data to support interlinking of building information models, Automation in Construction 124 (2021) 1–15.

[4] H. Thomas, R. Brennan, D. O'Sullivan, Moom – a prototype framework for management of ontology mappings, in: Proc. Int. Conf. on Adv. Information Networking and Applications (AINA), 2011, pp. 548–555. doi:10.1109/AINA.2011.55.

[5] N. Matentzoglu, J. P. Balhoff, S. M. Bello, M. H. Brush, S. Carbon, M. Courtot, A. D. Diehl, D. Dooley, H. J. Drabkin, M. A. Harris, M. Haendel, C. J. Mungall, D. A. Natale, D. Osumi-Sutherland, N. Perry, K. Shefchek, B. Smith, N. Vasilevsky, N. Vasilevsky, J. A. McMurry, A simple standard for sharing ontological mappings (sssom), Database 2022 (2022) baac035. URL: https://doi.org/10.1093/database/baac035. doi:10.1093/database/baac035.

[6] R. Volz, et al., Silk: A link discovery framework for the web of data, in: Proceedings of the International Semantic Web Conference, 2009.

[7] A.-C. N. Ngomo, S. Auer, Limes - a time-efficient approach for large-scale link discovery on the web of data, in: Proceedings of the 22nd International Joint Conference on Artificial Intelligence (IJCAI), 2011, pp. 2312–2317.

[8] Z. Martinková, P. Escobar, M. Galkin, V. Rodríguez-Doncel, et al., D4.5: Guidelines for enriching mappings and crosswalks with FAIR-enabling metadata, Technical Report, FAIR-IMPACT Project, 2024. URL: https://zenodo.org/records/15111167, zenodo, DOI:10.5281/zenodo.15111167.

[9] FAIR-IMPACT Consortium, Semantic artefact mappings, 2024. URL: https://fair-impact.eu/semantic-artefact-mappings, accessed June 2025.

[10] J. Toledo, A. Iglesias-Molina, D. Chaves-Fraga, D. Garijo, Rmldoc: Documenting mapping rules for knowledge graph construction, in: ESWC24 Posters and Demos Track, 2024.

[11] Rdf mapping language (rml), https://rml.io/specs/rml/, 2023. Accessed: Mar. 29, 2023.

[12] A. Ferrara, A. Nikolov, F. Scharffe, Data linking for the semantic web, Int. J. Semantic Web Inf. Syst. 7 (2011) 46–76.

[13] J. Euzenat, An api for ontology alignment, in: S. A. McIlraith, D. Plexousakis, F. van Harmelen (Eds.), The Semantic Web – ISWC 2004, Springer Berlin Heidelberg, 2004, pp. 698–712.

[14] D. O'Sullivan, V. Wade, D. Lewis, Understanding as we roam, IEEE Internet Comput. 11 (2007) 26–33. doi:10.1109/MIC.2007.50.

[15] S. Alzahrani, D. O'Sullivan, Towards a unified metadata model for semantic and data mappings, in: Ontology Matching Workshop at ISWC, 2022.

[16] S. Alzahrani, D. O'Sullivan, What metadata is needed for semantic and data mappings?, in: 2023 IEEE 17th International Conference on Semantic Computing (ICSC), 2023, pp. 312–316. doi:10.1109/ICSC56153.2023.00064.

[17] L. Moreau, et al., Prov-o: The provenance ontology, 2013. URL: https://www.w3.org/TR/prov-o/.

[18] C. L. F. Rocha, et al., Data quality vocabulary (dqv), 2015. URL: https://www.w3.org/TR/vocab-dqv/.

[19] M. Randles, et al., Mapping quality vocabulary (mqv): A domain-specific model for capturing metadata and provenance relating to mapping quality assessment, refinement, and validation, Journal of Semantic Web 15 (2021) 305–323. URL: https://doi.org/10.1007/s11220-021-00325-4. doi:10.1007/s11220-021-00325-4.

[20] S. Alzahrani, D. O'Sullivan, Metamap: Facilitating mapping reuse and lifecycle documentation, in: 2024 IEEE 18th International Conference on Semantic Computing (ICSC), 2024, pp. 323–328. doi:10.1109/ICSC59802.2024.00058.

[21] A. Randles, L. McKenna, L. Kilgallon, B. Yaman, P. Crooks, D. O'Sullivan, The knowledge graph explorer for the virtual record treasury of ireland, in: VOILA Workshop - International Semantic Web Conference (ISWC), CEUR, 2024, pp. 1–15.