

User Interaction Patterns for Linked Data

Mariana Aguiar [a,b], Sérgio Nunes [a,b,*] and Bruno Giesteira [a,c]
[a] *INESC TEC, Porto, Portugal*
[b] *Faculty of Engineering, University of Porto, Portugal*
E-mails: up201605904@edu.fe.up.pt, sergio.nunes@fe.up.pt
[c] *Faculty of Fine Arts, University of Porto, Portugal*
E-mail: bgiesteira@fba.up.pt

Abstract. Linked Data is often still perceived as data that will only be consumed by machines, and not by humans as well. As a result, Linked Data applications still often use more traditional visualisations that come with usability issues. However, alternative user interaction approaches have been developed and evaluated, many of which have proven to be better solutions for inexperienced users. One technique to formalise and document these user interaction techniques and best practices is in the form of patterns. While many pattern collections in the literature cover user interaction problems, there are no works targeted at user interaction problems for Linked Data interfaces. A survey conducted with Linked Data researchers and developers proved this need for user interaction patterns in the community, with over 90% of participants rating them as useful. Here, we propose a pattern collection of 20 novel user interaction patterns for Linked Data as a result of the abstraction of common problems and solutions for visualising, searching, browsing, and authoring Linked Data. The proposed patterns are the combined result of the most common problems reported by members of the community, the knowledge and experience gathered in 10 pattern mining interviews, and the recurrent approaches collected through a literature review of the solutions for user interaction with Linked Data. To validate the structure of the pattern collection, we followed the Pattern Classification method and, to evaluate the adoption and quality of each proposed pattern, we conducted a pattern adoption survey. From this evaluation, we obtained positive scores for 18 out of the 20 patterns. We believe that the proposed pattern collection can be a valid and helpful tool for developers, regardless of experience, to improve the user interfaces of their Linked Data applications.

Keywords: Linked Data, User Interaction, Pattern, Linked Data User Interaction

1. Introduction

With the rise of interest in Linked Data, we start to see more applications and platforms that use this data modeling paradigm. This interest is fueled by the benefits that semantically annotated and machine-readable information can have in many systems, like search engines and recommendation systems. However, Linked Data is still perceived as data only to be created and consumed by machines and robots, and not by humans as well. As a result Linked Data applications and tools often use the more traditional graph-based visualisations, that come with usability issues and are not familiar to the users, especially when associated with large and highly interconnected datasets. Thus, the web of data is still not commonly used by users without knowledge of Linked Data and previous experience with the related technologies. However, alternative visualisation techniques and user interaction approaches have been developed, studied, and evaluated in the literature since the beginning of the Semantic Web. Many of these less typical approaches have proven to be better solutions for inexperienced users regarding user interface usability.

However, research in the Linked Data field is lacking formalised and documented guidelines on how to design user interaction for Linked Data. One technique to formalise and document these user interaction techniques and best

[*] Corresponding author.

practices is in the form of patterns and pattern languages. Pattern languages are used to document and summarise best practices in a field of expertise, by abstracting common problems and solutions found in the literature and in the industry for that topic. While many pattern collections in the literature cover user interaction problems in more generic or specific contexts such as web, mobile, and social interfaces, there aren't works targeted at user interaction problems for Linked Data interfaces. We believe a pattern collection in this topic can support both experienced and inexperienced developers and designers in creating Linked Data applications that are intuitive and familiar to new users, motivating the use of the Semantic Web.

Section 2 presents an overview on pattern languages, including works that propose other user interaction pattern collections. Section 3 presents related work on the topic of interaction with linked data. Section 4 describes the methodology conducted to create the pattern collection proposed in this work. Section 5 details the steps in the pattern mining phase and briefly presents their results. Section 6 presents and describes the proposed pattern collection for user interaction for Linked Data. Section 7 validates the pattern collection by evaluating its structure, applicability, and the adoption of each pattern. In Section 8, the main results of this study are described, and future work is outlined.

2. User Interaction Patterns

The first use and formalisation of the term pattern, as a way to document and describe common problems and their proven solutions, was by Christopher Alexander in his work "A Pattern Language" [1]. Alexander introduces this notion of patterns as the description of recurring solutions to common problems in a way that allows the solutions to be reused in different cases. Alexander introduced this idea of pattern language as a mechanism, not only for experts to formalise and share their knowledge, but for laypeople to use as a vocabulary to express their ideas and designs. However, not all sets of patterns available can be considered pattern languages. Pattern catalogues are smaller sets of patterns that only cover a fraction of the common solutions for a certain domain. Pattern collections are sets of related or unrelated patterns that can be grouped together from other languages, collections or catalogues. Finally, pattern repositories are systems that provide access to a set of patterns from different pattern languages or catalogues.

Nowadays, the value of a piece of software is not entirely dependent on its quality, but also on how well it is interpreted and interacted with by users. So, studying best practices to model the user interaction with different types of systems is becoming more and more important to ensure that users can extract valuable information for the presented content, and quickly and correctly perform important tasks. The study of user interaction with more complex and uncommon data types, like Linked Data, is even more crucial, as they are an unfamiliar way of structuring data for most users and often require special and novel mental models.

One way of formalising and documenting the knowledge acquired studying human computer interaction is to create patterns and pattern languages. Collecting experiences and best practices, and composing a pattern language in user interaction is a great way to share knowledge in this area. Patterns are capable of presenting problems and explaining possible solutions in the form of an abstraction, providing the pattern's user with freedom to incorporate the solution in its design in a unique way. In "A Pattern Approach to Interaction Design", Borchers [2] provides a formal description of what a pattern in the HCI community should look like and justifies why the use of pattern languages aids not just a single design phase, but the whole design process. Borchers also proposes to apply the use of patterns and pattern languages to different and new application domains, explaining that they provide a vocabulary and ease communication in interdisciplinary

Over the years, a substantial number of works with suggestions of user interaction patterns were published. Here, we highlight 7 works considered more relevant from the literature on user interaction design patterns, a summary of these works can be consulted in 1. "Designing Interfaces: Patterns for Effective Interaction Design" [3] by Jenifer Tidwell is one of the first pattern collections published in the HCI community, with its first edition dating back to 1999. This pattern collection is organised in several sections, covering different parts or types of interaction with a system, from page organisation and layout, navigation and search, user input, actions and commands, to more specific topics like interaction with complex data, social media and mobile applications. In its first edition, this collection was targeted mostly at desktop applications, like email clients and browsers, however with the growth of the web, more recent editions include new patterns specific for web-based applications and mobile design. Tidwell's work features many patterns that are commonly encountered in numerous applications, and thus may now seem trivial. Nevertheless they once were not as widely used as right now, which comes to prove their importance and impact.

In "Interaction patterns in user interfaces", van Welie and Traetteberg [4] propose a pattern collection for interaction with user interfaces. van Welie and Traetterberg take on a different approach, compared to Tidwell's to identify patterns and deem them as "good" and "accepted" solutions. The patterns proposed are focused on the user perspective, meaning that they are targeting problems faced by end-users, as opposed to designers. In order to do this, the pattern collection is organised in categories representing the kinds of problems users have while interacting with a system: visibility, affordance, natural mapping, constraints, conceptual models, feedback, safety and flexibility [10]. Then, each pattern is focused on improving the usability of a system, and so, each one states the usability indicator it tries to improve. van Welie and Traetteberg defend that a user interaction design pattern must improve at least one usability usage indicator, otherwise it is not considered one.

Another pattern collection was published by Yahoo in 2012, "Yahoo Design Pattern Library" [5], it isn't only focused on user interaction but also on the design of user interfaces. It features several patterns that we can now identify in many recent websites and applications. Another work focused on web design is "Designing Social Interfaces: Principles, Patterns, and Practices for Improving the User Experience" [6] which is a pattern collection focused on gathering best practices and principles from applications and websites with social features or focus.

Other works like Laakso's [7] and Toxboe's [9] also propose pattern collections for interaction and interface design problems. These two works are not focused on a specific medium, as patterns can be applied to all types of applications, desktop, web-based or mobile. However, the pattern collection presented in "Designing Web Interfaces" [8] is targeted at tackling common problems in the design of web interfaces.

While some of the previously presented works differ on the types of systems that they are focused on, none are targeted at a certain data type. Most of the patterns featured in the pattern collections presented could be applied to systems with different data types. However, more complex data paradigms, such as Linked Data, emerge with more difficult interactions that could be confusing and not intuitive to users when not modelled correctly. Thus, a pattern collection on user interaction with Linked Data could be a useful tool for developers.

There isn't yet a pattern language or pattern collection formalised for interacting with Linked Data. However, we can point out one work proposing patterns for other LD related problems. In "Linked Data Patterns: A pattern catalogue for modelling, publishing, and consuming Linked Data" [11], a pattern catalogue is proposed for tackling problems when modelling, publishing and consuming LD. The pattern catalogue presents a section for application patterns, these are targeted at architectural and code design problems while developing the system, rather than visualisation and user interface design problems.

3. Interaction with Linked Data

In this section, we provide an overview of the relevant surveys in topic of interaction with Linked Data. While we do not focus on presenting a thorough survey of specific solutions for interacting with Linked Data, we have

Table 1

Overview of the reviewed user interaction pattern collections.

	Problem	Medium	Number of Patterns	Related Patterns
Tidwell [3]	Interface design	Digital interfaces	125	Yes
van Welie and Traetteberg [4]	Interaction Digital	Digital interfaces	22	Yes
Yahoo [5]	Web design	Web interfaces	59	No
Crumlish and Malone [6]	Social interface design	Social web interfaces	92	Yes
Laakso [7]	Interface design	Digital interfaces	21	Yes
Scott and Neil [8]	Web design	Web interfaces	65	No
Toxboe [9]	Web design	Web interfaces	90	Yes

gathered existing surveys on this topic to highlight their main conclusions. This will provide readers with a broader understanding of the trends and challenges in the field of Linked Data interaction.

The first relevant survey identified is from 2011, in "Approaches to Visualising Linked Data: A Survey" Dadzie and Rowe [12] focused on the identification of key requirements for Linked Data interaction to lower the technical barrier for end users and make the web of data accessible for all. The surveyed approaches were divided into two types, text-based browsers and browsers with visualisation options, with the former being mostly used for fine grained exploration, and the latter used for macro analysis (i.e. big picture). A total of 14 different tools were analysed under two sets of usability criteria, namely guidelines for visual representation and analytics, and guidelines for Linked Data consumption. The authors highlighted the limited number of LD browsers and attributed that to the infancy of the domain (in 2011).

In "Survey of Linked Data Based Exploration Systems" (2014), Marie and Gandon [13] organised their survey of 16 tools in three categories of systems, Linked Data browsers (divided in textual, visual, faceted, and others), recommender systems using Linked Data (divided in domain specific, cross-domain, undisclosed), and exploratory search systems (divided in view-based and algorithm-based). The authors observed a growing trend in research, from a focus on LD browsers (between 2001 and 2009) to a focus on LD-supported search (since 2010 to the date of the survey). The authors' vision that future web search experience would be highly impacted by the massive use of Linked Data proved to be correct given the impact that structured data and knowledge bases have on modern search systems [14].

In a survey focused on data exploration and visualisation systems for big data, Bikakis and Sellis [15] conducted an analysis on systems developed for the Web of Data. The survey is organized in six types of systems – browsers, generic visualisation, domain-specific, graph-based visualisation, ontology visualisation, and visualisation libraries. The survey has a focus on "big data visualisation" and the authors highlighted the need for a focus on large scale analysis and performance under demanding data requirements.

In "Semantic Web user interfaces - A model and a review" (2018), Charalampidis and Keramopoulos [16] present a review of applications that provide a user interface for accessing Semantic Web data. The authors analyse 37 applications using a model organised in six dimensions – type of user, information need of the user, output type, input mode, application category (search engine, browser, SPARQL query builder, domain-specific application), and access mode. From the authors' findings, we highlight two: most systems are designed for tech-users, and most systems are designed for information seeking, i.e. less for processing or adding new information.

In "Survey of Tools for Linked Data Consumption" (2019), Klímek el al. [17] target their survey at tools that allow for the consumption of Linked Data, specifically tools that are able to load a LD dataset, allow for user interaction without detailed knowledge of LD technologies, and provide non-RDF output to users. From an initial set of 110 tools identified as potential candidates, the authors selected 16 for a detailed evaluation under 94 criteria. These criteria were organised under 5 requirements groups: dataset discovery and selection; data manipulation; data visualisation, data output; and developer and community support. In line with other surveys, the authors highlighted the lack of support for non-LD experts in existing tools, and the need for integration of these tools into bigger platforms.

In "A Comparative Study of State-of-The-Art Linked Data Visualization Tools' (2020), Desimoni et al. [18] survey 10 Linked Data visualisation tools and evaluate their support over a set of 16 use cases identified by the authors as representative of LD visual exploration activities. As conclusions, the authors highlight the heterogeneity of the tools, their variability according to the domain and task, and the reduced support for custom visualisations (e.g. filter by type or property range).

In "Interacting with Linked Data: A Survey from the SIGCHI Perspective", Santo et al. [19] survey LD interaction literature published on SIGCHI venues. From a total of 60 papers identified on a first stage, the authors analysed 43 works according to four categories, the targeted end-users, the task context (e.g. querying, browsing, visualisation), the evaluation methodology, and the contribution (theoretical or practical). One of the most interesting findings is that in the decade previous to the survey (2010-2020) the community research was mostly used for machine consumption (e.g. improve recommendation systems, semantic enhancements of existing systems). Additionally, the authors also found that most research supporting human end-users focus on querying, browsing, and visualising LD. Again, this work confirms the reduced number of contributions targeting lay-users.

The most recent survey is "A Survey on User Interaction with Linked Data" (2021), where Aguiar et al. [20] review 18 tools specifically targeted at users with no experience with LD technologies. The authors describe and compare the solutions according to: the interaction features provided, the visualisation techniques adopted, the implementation of search functionalities, and the evaluation methods used in the research. The authors identify a trend to deviate from more traditional visualisations techniques (e.g. graph-based), and an effort in developing tools targeted at lay users. Amongst the most common solutions encountered is the use of facets to support search and the use of multiple data visualisation options for viewing one or multiple LD resources.

4. Pattern Development Methodology

To reach this work's goals and create a pattern collection for user interaction with Linked Data, we structured this research into three stages: pattern mining, pattern writing, and validation. We started by performing a study of Linked Data tools and applications to analyse the solutions for user interaction. Next, we conducted an online survey, with members of the Linked Data community about their experience using and developing with Linked Data. Then, in order to mine the pattern seeds [21], we performed interview sessions with experienced Linked Data researchers and developers to gather knowledge on the main recurring problems and their possible solutions. To complement the results from these interviews, we relied on the study of LD applications to identify common solutions that could be abstracted. For the pattern writing phase we followed the best practices for writing patterns and pattern languages present in the literature, and chose a pattern format suited for user interaction patterns. To validate the resulting pattern collection we performed three main tasks. The first was focused on the evaluation of the structure of the pattern collection and the identification of its strong and weak points, where we followed the method Pattern Classification [22]. In the second task, we evaluated the adoption of the patterns and the quality of the solutions proposed through an online survey based on the method proposed by Guerra and Fernandes [23]. Finally, in the third task the applicability to a specific context of the proposed pattern collection was studied by performing an analysis of the user interfaces of the Wikidata [24] platform.

5. Pattern Mining

In this section, we describe the steps conducted in the information collection phase of this study. In Section 5.1, we describe the literature review on solutions for user interaction with Linked Data. In Section 5.2, we detail the conducted survey about the participants' experience with Linked Data. In Section 5.3 we describe the pattern mining sessions.

5.1. User Interaction with Linked Data Solutions

Focused on collecting information for the creation of the patterns, we started out by conducting a literature review of the solutions for user interaction design in Linked Data systems. We studied LD tools targeted at inexperienced users to comprehend which user interaction features, like visualisation techniques and search solutions, are more frequently provided. We can point out some recurring solutions, such as the use of faceted search interfaces and several data visualisation options when consulting a single or a group of resources. Surprisingly, graph-based visualisation was the least common approach encountered in the set of works reviewed. This shows that researchers have started to deviate from the traditional visualisation techniques when developing for inexperienced users and that systematic documentation of these alternative solutions is necessary to broaden their use.

5.2. Linked Data Experience Survey

Before starting the pattern mining phase of the development of user interaction patterns for Linked Data, we set out to survey the opinion of members of the Linked Data community on the usefulness of user interaction patterns for the field, and to find Linked Data experts and developers willing to participate in pattern mining interviews. In order to do this, we developed and published an online questionnaire surveying the experience with Linked Data of the participants, common problems when using Linked Data applications, and the usefulness of user interaction patterns.

This section describes the conducted web survey with over 30 Linked Data researchers and developers, highlighting the design of the questionnaire, the analysis of the results obtained, and the conclusions taken from it.

5.2.1. Survey Design and Execution

The goal of the online survey was to collect information on the common problems of user interaction with Linked Data by asking the participants to share their experience with Linked Data / RDF / Semantic Web applications. Additionally, our intent was also to gather contacts of experienced researchers and developers willing to participate in a more thorough interview about their experience and findings. The goal of these more in depth interviews is to perform pattern mining, which means the collection of knowledge that will originate possible user interaction patterns for Linked Data.

The questionnaire was composed of a total of 15 questions, to be answered in an average total time of 5 minutes. The first section comprised 7 questions about the participants' experience with Linked Data as users or as developers. Section 2, composed of 4 questions, surveyed the participants opinion on the importance and need for design patterns for user interaction with Linked Data. Section 3, with 5 questions, aimed to identify participants interested in participating in a pattern mining interview. Given that the participants were all volunteers, we took into account that their availability and willingness to spend time answering the survey was limited. Due to this, we made an effort in balancing the number of questions with the time required to answer them, allowing us to extract useful information without the questionnaire answer time becoming too long.

The questionnaire was set up using Google Forms and disseminated via email to Linked Data researchers, online forums, and mailing lists on the topic. Some of these mailing lists were the following: W3C's RDF-DEV Community Group[1], Linked Open Data[2] mailing list, and the Knowledge Graph Conference[3] community on Slack. The form was online from the 5th of March 2021 to the 30th of April 2021, during which time we got 32 responses.

5.2.2. Survey Results

This sections present a data analysis of the results obtained in the survey on experience with Linked Data. The results described next were collected from the 32 responses of Linked Data experts and developers. We will characterise the participants according to their experience with Linked Data, perform an analysis on the problems reported by the participants, and analyse the results on the usefulness of Linked Data interaction patterns and how they relate to the participants experience.

General Experience with Linked Data We asked every participant to rate their general experience studying, using, and developing Linked Data from 0 to 5, 0 corresponding to no experience, and 5 corresponding to a high level of experience. We can point out that a large majority of about 80% of the participants rated their experience studying Linked Data with a high level (4 and 5). This leads us to assume that most of the surveys' participants are likely to be experienced researchers and have academic work developed in the Linked Data field. The results for using and developing Linked Data are quite similar, with about 50% of participants rating their experience as high (4 and 5), and about 20 to 25 percent of participants rating their experience as low (1 and 2). Even though a large majority of participants rated their experience studying Linked Data as high (4 and 5), this category is the only one out of the three where a very small percentage of participants reported having no experience. All in all, the results obtained about the participants experience with Linked Data are in line with what we expected and hoped to obtain, as a large majority of users reported having a medium to high level of experience for the three categories, studying, using, and developing.

[1]https://lists.w3.org/Archives/Public/semantic-web/
[2]https://lists.w3.org/Archives/Public/public-lod/
[3]https://www.knowledgegraph.tech

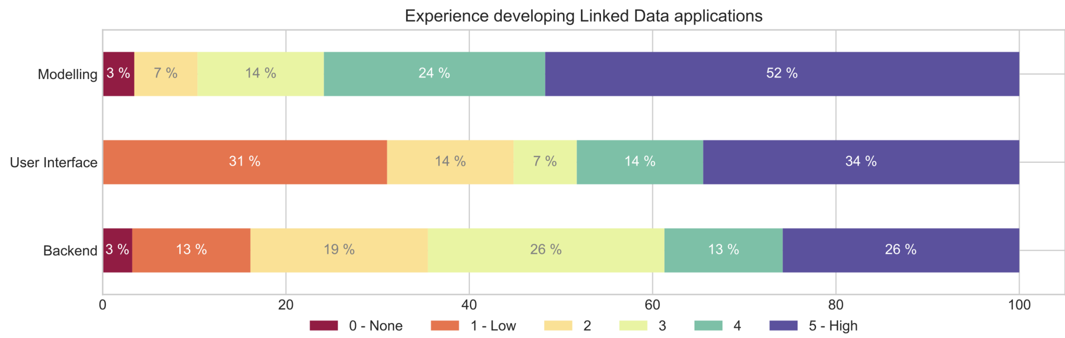

Fig. 1. Results of the participants experience developing Linked Data applications.

Experience Developing Linked Data Applications Next, we asked every participant to rate their specific experience in each area of the development of Linked Data applications, modelling, user interface, and backend. The results are presented in Figure 1, where we can consult the percentage of each level of experience for these three main areas. One of the reasons to collect information about the participants' experience developing with Linked Data, was to gather knowledge on the participants experience with user interfaces, and possibly with the design of user interaction. We can point out that modelling is the category with the highest percentage, of about 75%, of participants rating their experience as high (4 and 5), followed by user interface with about 50% of participants, and finally backend with 40% of participants. Even though user interface is not the category with the smallest percentage of participants with a high level of experience, it is by far the category with the largest percentage of participants with a low level of experience (1 and 2). The discrepancy of the results obtained for the user interface category, compared to the other two categories, are in line with the small number of research and studies regarding user interfaces and user interaction with Linked Data, and show that this is still a neglected aspect of the development of Linked Data applications.

Experience Using Linked Data Applications Lastly, we surveyed more in depth information about the participants experience using Linked Data applications, specifically in four main tasks performed in applications —- visualising, performing search tasks, browsing and navigating, and creating and editing data. The results obtained are not very different between the four tasks, but we can point out that browsing and navigating is the task with the highest percentage of users rating their experience as medium and high (3, 4, and 5) with about 75%. It is followed by creating and editing data with about 70%, and performing search tasks with close to 70% of participants. The task with the lowest percentage of participants is visualising with about 55%. These results inform us on which tasks are more important to the majority of users and will be helpful later on when gathering knowledge and creating the patterns.

Common Problems One of the goals of this survey is to gather information and knowledge on some of the most common problems, users and developers encounter in Linked Data applications. In the survey we included an open question where users could report problems encountered in Linked Data applications. We obtained a total of 28 answers with 25 considered valid, from which we extracted 47 problems. We performed a manual analysis of these problems and categorised them according to the three parts of the development of Linked Data applications – modelling, user interface, and backend. As the focus of this research is user interaction, we performed an additional analysis on the user interface problems. We can point out that user interface is the category with the largest number of problems associated, a total of 24, followed by backend (19) and then modelling (12). This reflects the fact that many Linked Data applications are still neglecting their user interface and user interaction problems, proving that this topic is still not sufficiently explored and analysed. These results also prove that applications, and consequently their users, will likely benefit from guidelines and studies performed to identify best practices when designing user interfaces and user interaction for Linked Data.

User Interface Problems Regarding the tasks associated to the user interface problems, we can clearly see that visualising has the vast majority of problems associated, with a total of 17. Next, we can point out browsing and navigating with 7 problems, and performing search tasks and creating and editing data with a single problem each. The distribution of problems amongst the four tasks does not align with the results obtained for the frequency and importance of problems, where visualising had the lowest percentage of high ratings. This could mean that participants rated more according to the importance of problems and less with their frequency, thus resulting in a higher number of visualising problems, but a lower importance associated to them. Regardless, these mixed results in the frequency of problems justify a pattern collection that tries to cover the four tasks as much as possible. As shown previously most of the reported user interface problems were related to the visualisation of Linked Data, and most of them report the difficulty of choosing visualisation techniques for the data, especially when the datasets are large and dense; the lack of mechanisms that provide the user with a clear overview of the available data and their vocabularies. Regarding the problems reported for browsing and navigating the dataset, these are focused on the fact that browsing new datasets is often hard and cumbersome due to the time it takes to understand the new ontology or vocabulary, and the difficulty to understand the data distribution. Another problem commonly reported is the fact that there are often too many links or irrelevant nodes in the interfaces; and that exploring relationships is difficult, especially when they are complex, causing a poor understanding of the data.

Usefulness of User Interaction Patterns Another goal of this online survey was to gather information on the opinion of members of the Linked Data and Semantic Web community about the usefulness of user interaction patterns for Linked Data. For this, we asked the participants to rate the need for guidelines on how to design user interaction with Linked Data, and to rate the usefulness of user interaction patterns for Linked Data. The results obtained for these two questions are presented in Figure 7.6. We can point out that, in both questions, a majority of the participants answered with a high rating (4 and 5), with the usefulness of patterns having the highest percentage of 85%, and the need for guidelines with 70%. We also observe that very few participants gave these questions a low rating (1 and 2), and that no participants gave the lowest rating to neither questions. We can conclude from these results that the lack of guidelines for this topic is a valid and relevant problem, and that the development and creation of user interaction problems to tackle it is a useful and needed solution.

Final Comments The final section of this online survey was dedicated to collect the contacts of participants willing to participate in a follow-up interview, part of the pattern mining sessions, and to give the space for the participants to leave comments and their opinions on the topic of this research. From the 32 participants we collected the contacts of a total of 18 people willing to participate in a follow-up interview. Regarding the comments and opinions of the participants of the survey, some mentioned that it is a very interesting research topic and that there is a need for more studies at the intersection of Linked Data and Human-Computer Interaction to better inform other researchers and professionals about this pertinent subject of Linked Data. Other participants mentioned that developers and designers generally have little experience in this topic, which tends to translate in a lack of capacity to create new functionalities for users that take advantage of the capabilities of Linked Data. Two participants mentioned a possible reason for the large amount of user interaction issues related to Linked Data. One argued that most of the these issues come from the fact that Linked Data based applications are built by a niche group of researchers and developers, that are more eager to push Linked Data capabilities on their users rather than focus on the users' tasks and only adding the needed modifications on top of more traditional UIs. This participant even suggested that we start to study the generic user interaction guidelines, and then analyse how they can be modified when taking in account Linked Data. The other participant stated that most of the Linked Data specific user interaction problems stem from the relatively complicated history of how the Semantic Web standards have been built. Only one participant rejected the core idea of this research, user interaction with Linked Data, arguing that it puts the storage mechanism and communication protocols at the heart of the interaction, and that is what has been holding back Linked Data. This participant defends that applications should be built to satisfy a user need and not for a researcher or developer to simply publish their data. This would result in the exploitation of Linked Data resources only because they have data that is required for an application or particular website. Essentially this participant argues that there are no new user interaction problems that come with the use of Linked Data, because it should not be transparent to the user interface. This was a very pertinent point that we furthered explored in the one-on-one interviews with Linked Data researchers and developers, by asking the participants' opinion on this topic

5.3. Pattern Mining Interviews

Pattern mining is the phase in the development and creation of patterns where the knowledge and information to write them is collected and gathered. There are a variety of methods in the literature on how to perform pattern mining, such as in work groups [21] or in workshops [25] with other pattern authors. However, these methods require a previous strong experience and knowledge on the domain by the author, as well as, a large group of people available at the same time and with experience on the domain. Due to the time restrictions, difficulty in finding experienced people on the field available to participate, and lack of deep knowledge and strong experience on the domain by the author, we decided to conduct expert interviews as a method of pattern mining. After disseminating the Linked Data experience survey and obtaining a solid number of responses and some people interested in participating in the pattern mining interviews, we contacted them and scheduled interviews. The goal of these interviews was to listen to researchers and developers experience with Linked Data, using, developing, or studying, and extract the common struggles and problems encountered regarding user interaction.

In this section we describe these pattern mining interviews, including the design of the interview guide, a thematic analysis of the results and their presentation, and a final reflection on the knowledge gathered and importance of these sessions for the development of the patterns.

5.3.1. Interview Design and Execution

After researching and studying how to conduct expert interviews, and how to decide which type of interview to perform, open, semi-structured, or structured, we decided on semi-structured interviews for the pattern mining sessions. Semi-structured interviews have a predefined structure and the interviewer has a general guide [26]. However, they are open to follow and investigate any other issues or topics raised by the participants [27], and give the interviewers the flexibility to ask the questions in any order, omit questions, or ask questions that are not in the guide. This style of interview better suits to perform pattern mining, as it conducts the interviewees to focus on their experience and problems and solutions, while at the same time provides the space for them to highlight what they consider more relevant and important.

We carried out a total of 10 pattern mining interviews, with Linked Data researchers and developers from around the world, 5 from Europe, 3 from North America, and 2 from South America. The interview pool was made up of 8 men, 1 woman, and 1 participant identifying themselves as non-binary. The interviews had an average time of around 60 minutes and were conducted online, via Google Meets. The participants were encouraged to share any examples and applications using the screen sharing feature, as a big part of the interview consisted of their previous projects and works in the Linked Data field. All the interviews were audio and video recorded, upon previous consent from the interviewees, and later transcribed. After the interviews, a thematic analysis was conducted on the interviews' transcriptions and notes taken.

The interview guide is composed of 3 parts and designed for an expected duration of 45 to 60 minutes. The first section covers the participants' introduction, consent for recordings, interview and research context, and a brief overview of the data collected on the need and usefulness of patterns and guidelines for the topic at hand. The second part is composed of 6 main questions focused on the interviewee's experience with Linked Data, as a user and as a developer, and on the problems encountered. The third and final part of the interview is composed of 4 main questions covering the solutions for the problems mentioned previously; context, examples, and consequences of the solutions; and final tips and comments on the topic.

As is usual with semi-structured interviews, not all questions were asked in most of the in- terviews conducted. However, some questions were consistently asked across the sessions and played an important part in the extraction of the knowledge and information for the patterns. The main questions asked in all interview session are the following:

- Can you talk a little about your opinion on the impact, future, preset,or past, of Linked Data on the web?
- Can you tell me about how you started to work on or be interested in Linked Data?
- What are some of the projects or research studies that you have been involved with that you would like to highlight?
- Do you use Linked Data based applications in your day to day life, for personal or work related tasks?

- Can you point out some common problems you have faced when developing Linked Data applications? If possible in specific user interaction problems?
- What do you usually do to tackle them? What are your most recurrent approaches?

5.3.2. Thematic Data Analysis

After conducting the pattern mining phase in the form of expert interviews with Linked Data re- searchers and developers, we set out to analyse the qualitative data gathered in the interviews' transcriptions and notes. These pattern mining interviews are a tool to perform exploratory research on the topic of user interaction with Linked Data, and thus produce a large amount of qualitative attitudinal data that corresponds to people's thoughts, experience, and believes that are reported to the interviewer [28].

Aiming to extract clear and valuable information from the interviews, in order to feed possible pattern ideas and the patterns themselves, we decided to conduct a thematic analysis on the interviews' transcriptions and notes. Thematic analysis is a method of systematically breaking down and organising rich data produced by qualitative research methods, and consists of tagging and annotating observations and quotations with different codes, in order to discover significant themes amongst the several interviews [29]. The three most common methods of performing thematic analysis are the use of specialised software, journaling, or the use of affinity diagramming techniques. We decided to use the journaling method that consists of the manual annotation and highlighting of the data, followed by a record of the researchers' ideas and thought processes on the interviews' data.

Following the interview research questions defined and mentioned previously, we categorised the data according to any mention of possible problems when developing Linked Data applications, and solutions used to tackle them. We tagged the data as a problem, a solution, or a pair of problem and related solution. Additionally, we categorised these problems and solutions as regarding user interface, modelling, and backend.

5.3.3. Problems and Solutions

The thematic analysis performed on the 10 pattern mining interviews identified a total of 38 problems, 22 solutions, and 35 pairs of problems and solutions. As expected, the category with more problems and solutions associated is user interface with a total of 70, followed by backend with 15, and modelling with 7. We collected the problems and solutions that were more frequently mentioned and the ones that were considered more relevant and useful for the creation of patterns, and aggregated and summarised them in Table 2. We can observe that the set of problems and solutions is composed of 3 categorised as authoring, 7 as visualising, 2 as browsing, and 1 as searching. This set is also comprised of 8 problem and solution pairs, 4 solutions, and 1 problem.

The information collected regarding the authoring of data in Linked Data applications is focused on aiding, specially inexperienced users, in the creation and edition of data, by providing familiar features, such as input areas for each part of the triple or resource with clear labels that lead the users to enter the correct information in the correct place, as can be consulted in Problem 1. Another solution frequently mentioned, detailed in Problem 3, is the inclusion of valid recommendations for each input area, to help the user choose and enter valid and correct data. Regarding authoring solutions, such as the one in Problem 2, some interviewees also mentioned that it is important to provide the user with the option to download the data that is presented to them, specially when one of the goals of the data collection is to be used by others in the Linked Data community.

Regarding the problems and solutions targeted at the visualisation of Linked Data, they are mainly focused on trying to help the user to understand the data better and more quickly. These include solutions like in Problem 4, defining a set of the most relevant and important information about a resource and the outline of how to present it to the user, in order to reduce the large number of triples displayed in a single page; and in Problem 6, aggregating a subset of triples to present the user with complete information that may be needed for a user task or need. Other solutions like Problem 5, also include the use of primitive formats of visualisation for each data type, such as maps for geographical data, and calendars for temporal data; in Problem 8, the ability by the user to choose from one of these provided visualisations to better suit their needs; and the inclusion of raw data in the visualisation options to allow the users to analyse it further if needed, in Problem 9. To help the users gain a better and more complete understanding of the data collections, some solutions were highlighted, such as providing an overview of the structure of the data, and the definitions of all internal and external resources and properties, detailed in Problem 10 and 7.

In regards to browsing the dataset, some interviewees mentioned that providing the users with links to the internal and external related resources is an easy to use solution, as detailed in Problem 12. However, other interviewees mentioned a problem that arises from it, due to the fact that everything is linked, it is easy for the users to loose track of where they are and where they started, while browsing the dataset via links, making them feel lost in the data collection, as can be consulted in Problem 11.

Finally, only one solution was highlighted regarding the ability to perform search tasks in the data collection, as is presented in Problem 13. Some interviewees mentioned the importance of providing a SPARQL query editor, in addition to more simple solutions like keyword-based approaches. This would allow the experienced users to create more complex queries when the more traditional approaches are insufficient. It was also mentioned that it should be possible to visualise the corresponding SPARQL query of the search task previously performed, to allow the users to edit it, or save it to be used in another occasion.

5.3.4. Other Topics

In addition to the problems and solutions identified previously, by performing a thematic analysis we identified other topics regarding the presentation of Linked Data, and the validity of this research study that we found worth mentioning. Several interviewees identified the open world assumption as one of their biggest struggles while developing Linked Data applications, as it isn't a user interface problem it was not mentioned in the previous section. According to these interviewees, the open world assumption is very difficult to work around when an application uses mainly local data, and when it doesn't the developers don't know what to expect regarding the structure and the volume of the data. This problem is aggravated when the application needs to present the unknown data to the user in a meaningful way, which is difficult to deal as a developer or as a designer.

One topic discussed during one of the interviews concerns moral issues related to the use and creation of Linked Data. This interviewee exposed the problem of the lack of diversity of voices in the Linked Data field, as its research tends to be concentrated on the Northern Hemisphere, specially Europe and North America. While at the same time there is a belief that Linked Data collections and datasets are a mirror of reality that people create and put online to be explored by everyone. However, as this interviewee reported, all the knowledge comprised in Linked Data collections is only partly true, due to the fact that it is contributed by individual people, with individual bias and individual geographical contexts. Even though, this is a problem hard to solve and tackle, it is important to create awareness of developers and data scientists about this issue, and the need to make clear when displaying data to the user that it represent a situated knowledge.

Following the comment identified in the online rejecting the core idea of this research studies, we posed this question to some of the researchers interviewed, by asking if they agree that the use of Linked Data poses new user interaction problems. Some of the interviewees reported that it depends, stating that in regards to user interfaces many components could be reused from traditional interfaces when using Linked Data and that in most cases it won't stray from the use of tables and lists, because that is what people are used to and familiar with. It was also mentioned that it depends on what the application wants to show to the user, it could be an application in some domain that chooses to present the data in a more traditional way, without it ever being transparent to the user that the application uses Linked Data. However, if the application wants to provide tools to explore large knowledge graphs, or Linked Data collections, new issues are going to arise from that. Another interviewee argued that new user interaction problems only arise if the application wants to show to the user that it is using Linked Data, and that it makes the most of its capabilities. The other half of interviewees reported that they believe Linked Data poses new user interaction problems, for example in providing ways to browse the dataset due to the fact that everything is linked, which makes it differ from more traditional approaches. It was also argued that Linked Data poses new problems, because developers and designers have to deal with the fact that they don't know what to expect when accessing the data, and that there isn't always immediate access to the description of everything on the web of data. However, it was also reported that even though it poses new problems, we can use the same heuristics to solve them, by taking the traditional approaches in as inspiration.

5.3.5. Conclusions

We can identify the conducted pattern mining interviews as the most important phase in the collection and gathering of knowledge to create the user interaction patterns. The results obtained from these interviews are rich and full of problems and solutions reported by developers and researchers in the field, and were very important to deciding which problems to cover and which solutions to propose in the patterns. We can conclude that the chosen interview format, as well as the developed guide, helped in extracting the right information. One aspect of the interviews, that proved to be more useful than initially thought was the sharing of concrete examples. This allowed us to view and explore new examples that we hadn't come in contact with yet, as well as, listen to the rationale behind some of the features and user interaction decisions.

From the 10 people interviewed in this pattern mining phase, we found some more useful than others, usually due to the number of examples discussed during the interview, and the type of examples. We found very useful to interview Linked Data developers that work in more commercial companies, in addition to academic researchers. The former group tends to work more with applications targeted at end-users, and not applications or tools solely developed to post and share datasets or data collections. In future and additional pattern mining sessions for this pattern collection, we would consciously make an effort to include more Linked Data developers that work outside of the academic world.

Even though the pattern mining interviews resulted in a collection of information rich in solutions, topics, and themes, it still portrays the opinions and experience of a very small group of people, that may not be representative of the Linked Data community. As mentioned previously, an effort was made to include in these sessions people from different locations and with different backgrounds, which was constrained by the time limitations and number of people available. When handling the information gathered, there needs to exist an understanding that it may be biased and not represent the general opinion of the Linked Data community. As such, thorough evaluations need to be conducted on the results that come from or are based on the knowledge and information gathered in these interviews.

6. Pattern Collection

This section introduces a pattern collection for user interaction with Linked Data, where we propose twenty novel patterns divided in four different categories: visualising, browsing, searching, and authoring. The patterns proposed in this section are a result of the previously presented study of the state of the art of Linked Data applications and tools, the common problems reported in the online survey, and the pattern mining sessions conducted with Linked Data researchers and developers.

The proposed patterns use a structure based on the classical pattern structure proposed in the literature [30, 31], which includes the context, problem, forces, and solution. We made some alterations to this proposed pattern structure to better fit the needs of the user interaction patterns. The pattern structure we opted for is composed of the following sections:

Context. Describes the situation in which the problem occurs, and details the circumstances that constraint the solution. Sometimes the context is described in terms of the patterns that have already been applied.

Problem. Presents the specific problem that needs to be solved.

Solution. Presents the proposed solution to the problem, describes the implementation details that need to be considered, and explains some of the rationale behind it.

Diagram. Shows a simple illustration that helps the reader to understand the pattern more quickly.

Examples. Presents concrete examples, aligned with the context, that illustrate how to apply the pattern.

An overview of the structure of the pattern collection can be consulted on Figure 2, where the patterns' names, categories, and relationships are present. The visualising category comprises 10 patterns and is focused on problems and solutions regarding the visualisation and display of Linked Data. The searching category is composed of 3 patterns and covers problems when performing search tasks in Linked Data datasets and data collections. The browsing category is also composed of 3 patterns and targets problems that arise from browsing and navigating Linked Data. The final category, authoring, comprises 4 patterns and covers problems and solutions for creating and editing Linked Data. A list of the patterns' patlets [32], a term known as a short description of a pattern, can be consulted in Table 3 and Table 4.

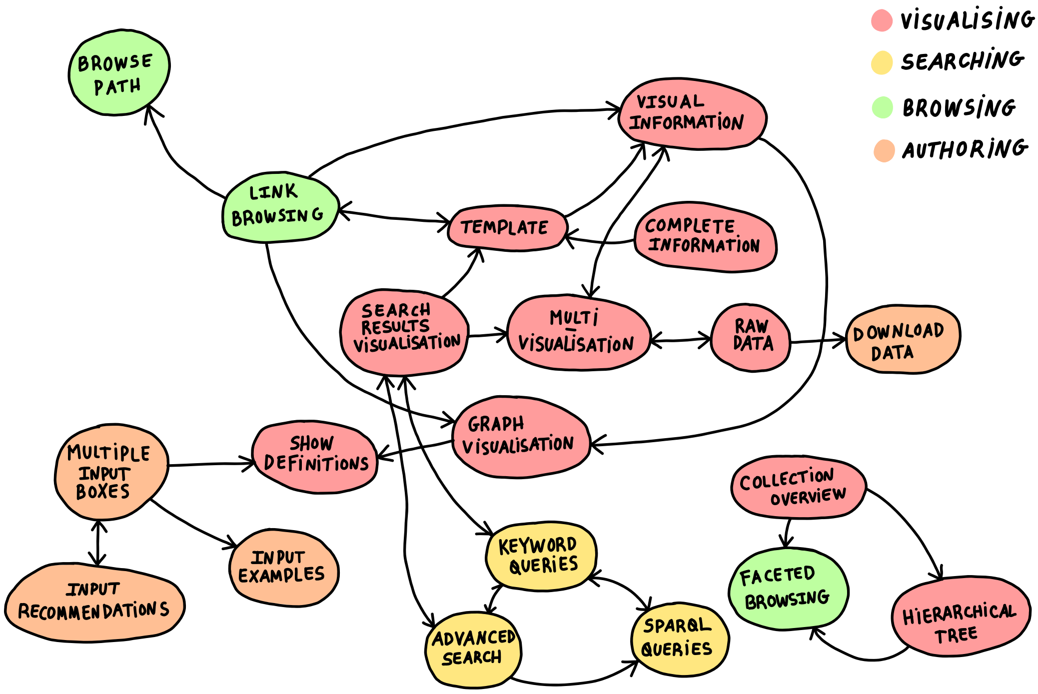

Fig. 2. Pattern collection structure overview.

7. Pattern Collection Validation

This section covers the methods followed to validate the pattern collection on user interaction with Linked Data and their results. In Section 7.1, we start out by conducting an analysis on the structure of the pattern collection, in order to identify its weak and strong points. In Section 7.2, we describe the questionnaire conducted to survey the proposed patterns quality and adoption by the Linked Data community. In Section 7.3, we studied the applicability of the pattern collection to a specific context, using the platform Wikidata [24] as an example.

7.1. Pattern Classification

Pattern Classification is a method proposed by Seidel [22] that aims to identify possible strong and weak points in the topics covered by a pattern language or pattern collection. To perform this classification method, a set of rules or categories on how to classify the patterns is needed. Thus, the quality of the results of the classification depend on the quality of the chosen qualification system. Seidel recommends the use of a classification system that is grounded on scientific theory, or the use of a model and classifications derived from the literature. The results of this analysis are the distribution of the patterns on each category, and are more valuable when compared to existing classifications of related pattern languages or collections. Based on the application of this pattern classification method provided by Seidel, we used as the categories for the classification of the patterns the nine layers of the "Universal Model of User Interface" [33]. This classification system was chosen due to the fact that it has already been showed by Hübscher et al. [34] that it is capable of describing user interaction pattern languages, like Tidwell's [3] and van Welie and Traetteberg's [4].

The classification system based on the "Universal Model of a User Interface" is composed of 9 different categories or layers, divided in three tiers. The first tier is composed of the lowest levels of the user interface: conceptual model, task flow, and organisation model; that form the conceptual and organisational basis of the user experience. The second tier comprises the middle layers of the model: viewing and navigating, editing and manipulation, user assistance. These describe the interactive qualities, users actions, and the system's reaction to them. Lastly, the final tier is composed of the presentation layers: layout, style, and text; that describe the specific visual and textual expressions of the interface.

7.1.1. Results and Discussion

Seidel recommends that the classification process is performed by someone familiar with the pattern collection to be analysed, thus the classification of the user interaction patterns for Linked Data was conducted by the author of the patterns. For each pattern, the author identified a matching category from the previously mentioned categorisation system. The results from this classification can be consulted in Table 7. The strong emphasis on patterns for viewing and navigation is not surprising, as for all Linked Data based applications, regardless of domain, navigation features are crucial. Either regarding performing search tasks or browsing the data in the application, these features pose new problems when Linked Data is used, due to the fact that everything is linked and to the difficulty in providing alternative approaches accessible to all users. Especially ones that are as powerful as the use of SPARQL queries, when searching data.

Following the viewing and navigation layer, we have two layers with the same percentage of patterns associated, user assistance and layout, both aspects important when designing user interaction for Linked Data applications. One of the main goals of the proposed pattern collection is to gather good solutions that will aid the user when interacting with Linked Data, by providing information that will help them understand the data and its structure. Thus, the moderately high percentage of patterns associated with this layer is in agreement with the pattern collections' goals. Regarding the layout layer, we can justify its percentage with the aim to help designers and developers make informed decisions when choosing what data to present and how to do it in a way that is accessible to all users of the application. This is an important aspect to cover in the pattern collection, because usually Linked Data is presented to the user with approaches that are not accessible to all and that pose user interaction problems.

Many of the problems and aspects that would be covered by the patterns in the organisational model layer, would not emerge from the use of Linked Data, but from the domain and goals of the application. Solutions that are accessible to all users for problems regarding editing and manipulation are important to be covered by a pattern collection for Linked Data. However, the proposed collection only has 1 pattern associated to this layer. We can justify this with the difficulty to find solutions for this problem that are accessible to all users. Nevertheless, we believe it is a matter that needs to be explored more in depth in future iterations of this pattern collection.

We can justify the lack of patterns associated with the conceptual model and task flow layers with the fact that the proposed pattern collection is focused on new user interaction patterns that arise with the use of Linked Data, regardless of the domain or technical environment of the applications. Patterns concerned with the style layer have not been considered, because their problems are unrelated to the use of Linked Data, and can be constrained by style guides or manuals of the company or organisation of the application. When problems related to the style layer come up in Linked Data based applications, the general solutions covered by collections like Tidwell's [3] can be applied. Similarly, problems regarding the text layer are unrelated to the use of Linked Data, thus the existing solutions can be used when these problems arise.

7.2. Pattern Adoption Survey

To evaluate the patterns' adoption and usage amongst the Linked Data community, we conducted an online survey with researchers and developers from the field. Evaluating each pattern is important to understand if the problems and solutions documented are commonly encountered and used, and if the solutions are considered as good by the community. In order to perform this evaluation, we based our questionnaire on the method proposed by Guerra and Fernandes [23]. The method proposed is composed by a 3 star rating system, that combines several characteristics of the patterns and pattern collection. The first star is attributed to a pattern if at least three known use cases are identified for it, the second and third stars are received by a pattern if it is considered a recurrent and good solution amongst experienced and inexperienced developers in the field. From this method, we conducted the usage and adoption questionnaire for experienced Linked Data developers and researchers.

The authors outline a set of options that the participants of the questionnaire must choose between regarding their experience and opinion on the problem and solution documented by each pattern. These answer options allow the author of the patterns to understand if the participants have used the pattern or not, and if they consider it a good solution for the problem that it documents. Each response option is associated with an evaluation and with a number of points, that are used to calculate the final score for each pattern. Guerra and Fernandes state that for a pattern to receive an overall positive score and be recognised as a recurrent and good solution, it must obtain a score higher than half of the participants' total number.

7.2.1. Design and Execution

Given that the participants were all volunteers, we took into account that their availability and willingness to spend time answering the survey was limited. Thus, instead of presenting the participants with a full version of each pattern, we decided to only include a summarised version, present in Table 3 and Table 4, to minimise the total time of the survey, and consequently increase the number of responses. In addition to the 20 questions where the participants could report their opinion and usage of each pattern, we included an open text question for each pattern category, where the participants could leave any comments and thoughts on the patterns. After designing the questionnaire we set it up using Google Forms and disseminated it via email to the Linked Data researchers interviewed in the pattern mining session, and additional mailing lists in the Linked Data community. Some of these mailing lists include the W3C's RDF-DEV Community Group[4] and Linked Open Data[5] mailing list, the Knowledge Graph Conference community on Slack[6], the Wikidata[7] mailing list, and the Europeana Tech[8] mailing list. The form was online from the 21st of May to the 14th of June, during which time we got 20 responses.

7.2.2. Survey Answers

From the 20 obtained answers to the pattern usage and adoption survey, we received from every participant a rating for almost every proposed pattern. Apart from the COLLECTION OVERVIEW, SPARQL QUERIES, INPUT RECOMMENDATIONS, DOWNLOAD DATA, MULTIPLE INPUT BOXES, and INPUT EXAMPLES patterns with 19 ratings, every pattern received a total of 20 ratings. The answers for each proposed pattern are combined in Figure 3. We can point out from the obtained results that 13 out of the 20 proposed patterns have a 50 or higher percentage of participants reporting to have already used them, exactly as they are presented or with some variations. We can also highlight the LINK BROWSING pattern as the pattern with the highest percentage of participants stating to have used it, with 95%, and INPUT EXAMPLES with the lowest percentage of about 20%. In regards to the answers that reported not using or not being sure if they have used the pattern, we can point out the SEARCH RESULTS VISUALISATION pattern with the highest percentage of participants considering it a good solution regardless, followed closely by INPUT EXAMPLES, COMPLETE INFORMATION and INPUT RECOMMENDATIONS patterns.

Regarding the comments by the survey participants, we can point out, related to the visualising patterns, a participant expressing the need for new visualisation techniques. In addition to the commonly used ones documented in the proposed patterns, this participant states that visualisation techniques that help users transcend complexity and improve contextual understanding are needed. We can also highlight an issue reported, when visualising Linked Data choosing between aggregation, where the resources are shown with their related information, or navigation, show- ing hyperlinks for the related resources. The participant states that the later results in a simple and consistent way of presenting the data, but can leave users lost when navigating. However, the former can result in overloading the user with information. This reported issue is one we try to tackle with the proposed pattern COMPLETE INFORMATION, but that we recognise it still needs special attention in a future iteration of the collection. Another participant pointed out that the visualisation of the dataset's metadata is important and should be covered by the collection, as it can help provide context and ease the interpretation of the data. Finally, a participant expressed their

[4] https://lists.w3.org/Archives/Public/semantic-web/
[5] https://lists.w3.org/Archives/Public/public-lod/
[6] https://www.knowledgegraph.tech
[7] https://lists.wikimedia.org/hyperkitty/list/wikidata@lists.wikimedia.org/
[8] https://www.europeana.eu/en

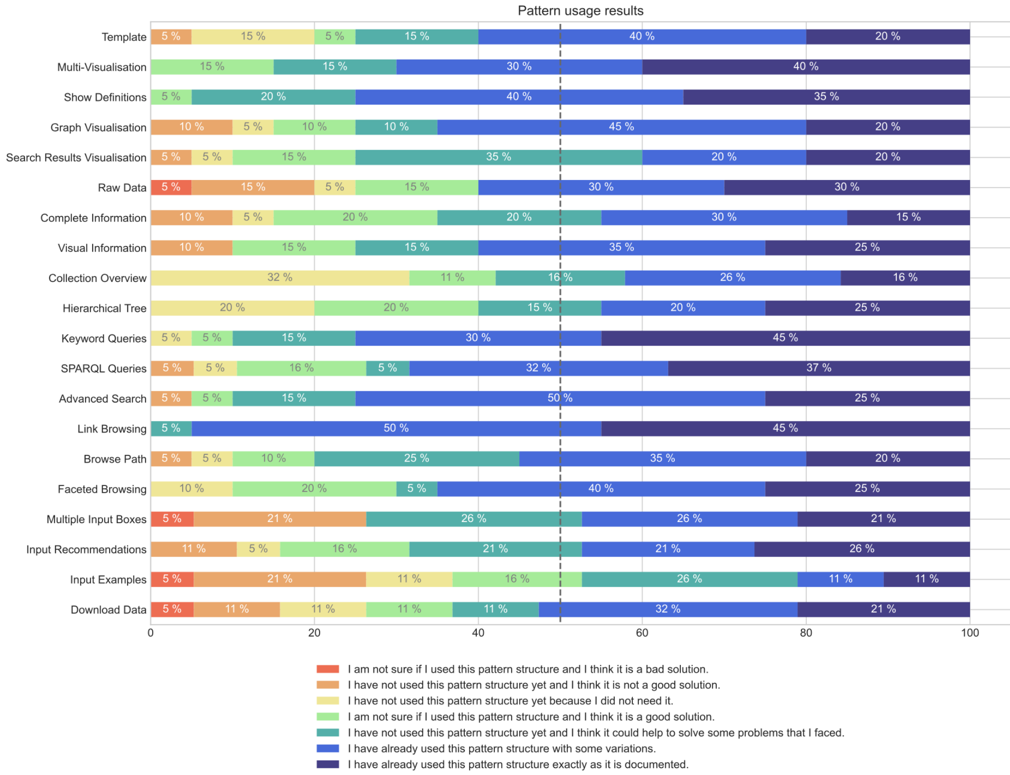

Fig. 3. Pattern usage results.

concern with the fact that graph visualisations don't scale very well, and can quickly become overwhelming to the users. This is something we covered and tried to tackle while writing the GRAPH VISUALISATION pattern.

Regarding the comments on the searching patterns, we can highlight one participant that states that while they usually use a variation of the SPARQL QUERIES pattern, they try to use instead augmented search techniques that tailor semantics to the persona and jobs the individual is trying to perform. We can also point out a problem reported by a participant, that should be considered in a future interaction of the pattern collection. This participant reports that in practice it is difficult to find the right balance of heuristic searching and producing too many results, as it would be useful to have the system match close results when a user misspells terms. However, this can greatly increase the number of results and consequently make hard for the user to find the right result.

For the browsing patterns, we can highlight two comments, one states that the participant found that keeping a browsing history does not help the user, according to their evaluations. The other comment states that while faceted browsing is a powerful feature that users tend to enjoy, it can cause a heavy load on the backend query engine.

Finally, regarding the authoring patterns, the majority of comments expressed the fact that users don't typically think of the underlying structure as triples, but rather in the domain language. For this reason, some participants tend to prefer to have the user interface do the work of translating between domain terms, that the user understands, and the underlying structure. We can also point one participant that stated the preference to encourage saving the Uniform Resource Identifiers instead of letting the user download the data.

7.2.3. *Pattern Usage Results*

After briefly analysing the unprocessed answers to the pattern usage survey, we followed the method outlined by Guerra and Fernandes to calculate the scores for each pattern resulting from the obtained answers. The final patterns' scores can be consulted on Table 8. Following the recommendation by Guerra and Fernandes to consider a pattern score as positive if its absolute number is larger than half of the total number of ratings, we can conclude that 18 out of the 20 proposed patterns have a positive score.

Looking at the results on Table 8, we can observe that the patterns MULTIPLE INPUT BOXES and INPUT EXAMPLES have a score lower than half of their total ratings, which means that they have an overall negative rating and need to be revised. We can also point out that 3 patterns, MULTI-VISUALISATION, SHOW DEFINITIONS, and LINK BROWSING, have only received positive ratings, and obtained a perfect score of 20. In addition to the patterns with negative scores, we can also identify some other patterns that might benefit from a revision of their content or structure, namely 3 patterns, COLLECTION OVERVIEW, RAW DATA, and DOWNLOAD DATA, due to their very low scores.

We can also highlight from the patterns' results that the four patterns with the lowest scores propose a solution where the internal structure of Linked Data is displayed to the user. Either by displaying data in RDF in the RAW DATA and DOWNLOAD DATA patterns or exposing the triple structure in the MULTIPLE INPUT BOXES and INPUT EXAMPLES patterns. This is one aspect that needs further attention on future iterations of the pattern collection, requiring revision of the already proposed patterns and an analysis on additional similar solutions.

7.3. Wikidata Analysis

To understand the applicability of the proposed pattern collection, we analysed Wikidata's [24] user interfaces and identified which patterns are already implemented, and which patterns are not. From the proposed 20 user interaction patterns for Linked Data, we identified 11 patterns in the Wikidata's interfaces, 3 visualising patterns, 3 searching patterns, 1 browsing pattern, and 4 authoring patterns. We can point out that all the proposed searching and authoring patterns are present in the user interface, and that the visualising category is the one with the lowest percentage of patterns implemented, with 30%, followed by the browsing category with 33% of patterns. The remaining 9 patterns were identified as possible to be implemented in order to better the users experience in the application.

During this exercise, we found that a deep and thorough understanding of the vocabularies used is needed, as well as, the knowledge about which tasks are more frequently executed, and which information is typically more important for the majority of users. This information is critical to understand if and how some patterns can be applied, such as the TEMPLATE, MULTI-VISUALISATION, COMPLETE INFORMATION, COLLECTION OVERVIEW, and FACETED BROWSING, where the implementation relies on choices of what information should be displayed, and how to do it. This was the only difficulty encountered in the execution of this analysis, where almost all the patterns proved able to be implemented regardless of vocabulary, or style choices of the application.

From this analysis, we can conclude that the proposed patterns are easily integrated and applied to existing user interfaces and Linked Data solutions, and that they work together to tackle a vast number of user interaction problems. Additionally, we were able to present to the user a more complete example of how the proposed abstract solutions can be applied together to a specific context.

8. Conclusions and Future Work

With the rise of interest in Linked Data, we start to see more applications and platforms that use this data modeling paradigm. This interest is fueled by the benefits that semantically annotated and machine-readable information can have in many systems, like search engines and recommendation systems. However, Linked Data is still perceived as data only to be created and consumed by machines and robots, and not by humans as well. As a result Linked Data applications and tools often use the more traditional graph-based visualisations, that come with usability issues and are not familiar to the users, especially when associated with large and highly interconnected datasets. Thus, the web of data is still not commonly used by users without knowledge of Linked Data and previous experience with the related technologies. However, alternative visualisation techniques and user interaction approaches have been developed, studied, and evaluated in the literature since the beginning of the Semantic Web. Many of these less typical approaches have proven to be better solutions for inexperienced users regarding user interface usability.

One technique to formalise and document these user interaction techniques and best practices is in the form of patterns and pattern languages. While many pattern collections in the literature cover user interaction problems in more generic or specific contexts such as web, mobile, and social interfaces, there aren't works targeted at user interaction problems for Linked Data interfaces. We believe a pattern collection in this topic can support both experienced and inexperienced developers and designers in creating Linked Data applications that are intuitive and familiar to new users, motivating the use of the Semantic Web.

To collect and gain knowledge needed in the creation of the patterns we conducted an online questionnaire with members of the Linked Data community, to survey their experience with Linked Data, collect common problems while developing LD based applications, and to analyse their opinion on the usefulness of user interaction patterns for Linked Data. This survey allowed us to collect over 40 problems, with a clear focus on visualising data for the user interface problems reported. This study also proved the need for user interaction patterns, with over 90% of participants giving a positive rating to their usefulness. Following the methods described in the literature on how to create a pattern language, we conducted pattern mining sessions from second-hand contributions, in the form of expert interviews. We executed 10 pattern mining interviews with Linked Data researchers and developers, whose contacts were collected with the help of the previously described online questionnaire. We consider these interviews as the most important phase in the development of the proposed pattern collection, as they allowed us to gather over 70 problems and solutions regarding user interfaces for Linked Data. The information collected proved very useful when deciding which problems to cover and which solutions to propose in the pattern collection.

The pattern collection proposed here is the combined result of the most common problems reported by members of the community in the conducted survey, the knowledge and experience gathered in the 10 pattern mining interviews with Linked Data researchers and developers, and the recurrent approaches collected through a literature review of the solutions for user interaction with Linked Data. The proposed pattern collection is composed of 20 novel user interaction patterns for Linked Data, and is focused on four main aspects of user interaction, containing patterns concerning visualising, searching, browsing, and authoring problems. We can also point out the information collected from the experience with Linked Data survey, the results of the pattern collection evaluation, and the findings of the Wikidata analysis as contributions of this work.

The future goals for the proposed pattern collection include the improvement of the currently written patterns, and the addition of new ones to cover all the sub-problems encompassed in the design of user interaction for Linked Data, in order to come to a complete pattern language for this topic. To revise the already proposed patterns, additional evaluation sessions for the proposed pattern collection should be conducted, to understand why some patterns received less positive scores and how they can be improved. Different formats of evaluation methods should also be taken into consideration in the future, to allow for the evaluation of the full version of the patterns. To expand the pattern collection, additional pattern mining interviews can be conducted to gather even more solutions to user interaction problems that were not covered by the current version of the pattern collection.

Acknowledgements

This work was supported by National Funds through the Portuguese funding agency, FCT – Fundação para a Ciência e a Tecnologia within project EPISA DSAIPA/DS/0023/2018.

References

[1] C. Alexander, *A Pattern Language: Towns, Buildings, Construction*, Oxford University Press, USA, 1977, p. 1216. ISBN 0-19-501919-9. https://archive.org/details/patternlanguage00chri/page/1216.

[2] J.O. Borchers, A Pattern Approach to Interaction Design, in: *Proceedings of the 3rd Conference on Designing Interactive Systems: Processes, Practices, Methods, Techniques, DIS 2000, New York City, NY, USA, August 17-19, 2000*, D. Boyarski and W.A. Kellogg, eds, ACM, 2000, pp. 369–378. doi:10.1145/347642.347795.

[3] J. Tidwell, *Designing interfaces - patterns for effective interaction design (2. ed.)*, O'Reilly, 2011. ISBN 978-1-449-37970-4. http://www.oreilly.de/catalog/9781449379704/index.html.

[4] M. van Welie and H. Traetteberg, Interaction Patterns in User Interfaces., in: *7th Pattern Languages of Programs Conference (PLoP 2000).*, 2000, PLoP 2000 ; Conference date: 01-01-2000 Through 01-01-2000.

[5] Yahoo, Yahoo Design Pattern Library, 2015.

[6] C. Crumlish and E. Malone, *Designing Social Interfaces - Principles, Patterns, and Practices for Improving the User Experience*, O'Reilly, 2009. ISBN 978-0-596-15492-9. http://www.oreilly.de/catalog/9780596154929/index.html.

[7] S.A. Laakso, User Interface Design Patterns, 2003.

[8] B. Scott and T. Neil, *Designing Web Interfaces - Principles and Patterns for Rich Interactions*, O'Reilly, 2009. ISBN 978-0-596-51625-3. http://www.oreilly.de/catalog/9780596516253/index.html.
[9] A. Toxboe, User Interface Design Patterns, 2021.
[10] D.A. Norman, *The Design of Everyday Things*, Basic Books, Inc., USA, 2002. ISBN 9780465067107.
[11] L. Dodds and I. Davis, Linked Data Patterns: A pattern catalogue for modelling, publishing, and consuming Linked Data, 2012.
[12] A. Dadzie and M. Rowe, Approaches to visualising Linked Data: A survey, *Semantic Web* **2**(2) (2011), 89–124. doi:10.3233/SW-2011-0037.
[13] N. Marie and F.L. Gandon, Survey of Linked Data Based Exploration Systems, in: *Proceedings of the 3rd International Workshop on Intelligent Exploration of Semantic Data (IESD 2014) co-located with the 13th International Semantic Web Conference (ISWC 2014), Riva del Garda, Italy, October 20, 2014*, D. Thakker, D. Schwabe, K. Kozaki, R. García, C. Dijkshoorn and R. Mizoguchi, eds, CEUR Workshop Proceedings, Vol. 1279, CEUR-WS.org, 2014. http://ceur-ws.org/Vol-1279/iesd14_8.pdf.
[14] K. Balog, *Entity-Oriented Search*, The Information Retrieval Series, Vol. 39, Springer, 2018. ISBN 978-3-319-93933-9. doi:10.1007/978-3-319-93935-3.
[15] N. Bikakis and T.K. Sellis, Exploration and Visualization in the Web of Big Linked Data: A Survey of the State of the Art, in: *Proceedings of the Workshops of the EDBT/ICDT 2016 Joint Conference, EDBT/ICDT Workshops 2016, Bordeaux, France, March 15, 2016*, T. Palpanas and K. Stefanidis, eds, CEUR Workshop Proceedings, Vol. 1558, CEUR-WS.org, 2016. http://ceur-ws.org/Vol-1558/paper28.pdf.
[16] C.C. Charalampidis and E. Keramopoulos, Semantic Web user interfaces - A model and a review, *Data Knowl. Eng.* **115** (2018), 214–227. doi:10.1016/j.datak.2018.04.003.
[17] J. Klímek, P. Skoda and M. Necaský, Survey of tools for Linked Data consumption, *Semantic Web* **10**(4) (2019), 665–720. doi:10.3233/SW-180316.
[18] F. Desimoni, N. Bikakis, L. Po and G. Papastefanatos, A Comparative Study of State-of-The-Art Linked Data Visualization Tools, in: *Proceedings of the Fifth International Workshop on Visualization and Interaction for Ontologies and Linked Data co-located with the 19th International Semantic Web Conference (ISWC 2020), Virtual Conference (originally planned in Athens, Greece), November 02, 2020*, V. Ivanova, P. Lambrix, C. Pesquita and V. Wiens, eds, CEUR Workshop Proceedings, Vol. 2778, CEUR-WS.org, 2020, pp. 1–13. https://ceur-ws.org/Vol-2778/paper1.pdf.
[19] A.D. Santo and A. Holzer, Interacting with Linked Data: A Survey from the SIGCHI Perspective, in: *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems, CHI 2020, Honolulu, HI, USA, April 25-30, 2020*, R. Bernhaupt, F.F. Mueller, D. Verweij, J. Andres, J. McGrenere, A. Cockburn, I. Avellino, A. Goguey, P. Bjøn, S. Zhao, B.P. Samson and R. Kocielnik, eds, ACM, 2020, pp. 1–12. doi:10.1145/3334480.3382909.
[20] M. Aguiar, S. Nunes and B. Giesteira, A Survey on User Interaction with Linked Data, in: *Proceedings of the Sixth International Workshop on the Visualization and Interaction for Ontologies and Linked Data co-located with the 20th International Semantic Web Conference (ISWC 2021), Virtual Conference, 2021*, P. Lambrix, C. Pesquita and V. Wiens, eds, CEUR Workshop Proceedings, Vol. 3023, CEUR-WS.org, 2021, pp. 13–28. https://ceur-ws.org/Vol-3023/paper5.pdf.
[21] A. Sasabe, T. Kaneko, K. Takahashi and T. Iba, Pattern Mining Patterns: A Search for the Seeds of Patterns, in: *Proceedings of the 23rd Conference on Pattern Languages of Programs*, PLoP '16, The Hillside Group, USA, 2016.
[22] N. Seidel, Empirical Evaluation Methods for Pattern Languages: Sketches, Classification, and Network Analysis, in: *Proceedings of the 22nd European Conference on Pattern Languages of Programs, EuroPLoP 2017, Irsee, Germany, July 12-16, 2017*, ACM, 2017, pp. 13:1–13:24. doi:10.1145/3147704.3147719.
[23] E.M. Guerra and C.T. Fernandes, An evaluation process for pattern languages, in: *Proceedings of the 8th Latin American Conference on Pattern Languages of Programs, SugarLoafPLoP 2010, Salvador, Bahia, Brazil, September 23-27, 2010*, R. Wirfs-Brock and U. Kulesza, eds, ACM, 2010, pp. 18:1–18:11. doi:10.1145/2581507.2581525.
[24] D. Vrandecic and M. Krötzsch, Wikidata: a free collaborative knowledgebase, *Commun. ACM* **57**(10) (2014), 78–85. doi:10.1145/2629489.
[25] Y. Akado, S. Kogure, A. Sasabe, J. Hong, K. Saruwatari and T. Iba, Five patterns for designing pattern mining workshops, in: *Proceedings of the 20th European Conference on Pattern Languages of Programs, EuroPLoP 2015, Kaufbeuren, Germany, July 8-12, 2015*, C. Link and V. Eloranta, eds, ACM, 2015, pp. 9:1–9:13. doi:10.1145/2855321.2855331.
[26] K. Pernice, User Interviews: How, When, and Why to Conduct Them, 2018.
[27] M. Rosala, Writing an Effective Guide for a UX Interview, 2021.
[28] M. Rosala, How to Analyze Qualitative Data from UX Research: Thematic Analysis, 2019.
[29] G. Guest, K. MacQueen and E. Namey, *Applied Thematic Analysis*, Thousand Oaks, California, 2012. https://www.doi.org/10.4135/9781483384436.
[30] G. Meszaros and J. Doble, A Pattern Language for Pattern Writing, *Pattern Languages of Program Design* **3** (1997), 529–574. ISBN 0-201-31011-2.
[31] T. Wellhausen and A. Fiesser, How to write a pattern?: a rough guide for first-time pattern authors, in: *16th European Conference on Pattern Languages of Programs, EuroPLoP 2011, Irsee, Germany,July 13-17, 2011*, A. Fießer and P. Avgeriou, eds, ACM, 2011, p. 5. doi:10.1145/2396716.2396721.
[32] J.O. Coplien and N.B. Harrison, *Organizational Patterns of Agile Software Development*, Prentice-Hall, Inc., USA, 2004. ISBN 0131467409.
[33] B. Baxley, Universal Model of a User Interface, in: *Proceedings of the 2003 Conference on Designing for User Experiences*, DUX '03, Association for Computing Machinery, New York, NY, USA, 2003, pp. 1–14–. ISBN 1581137281. doi:10.1145/997078.997090.
[34] C. Hübscher, S.L. Pauwels, S.P. Roth, J.A. Bargas-Avila and K. Opwis, The organization of interaction design pattern languages alongside the design process, *Interact. Comput.* **23**(3) (2011), 189–201. doi:10.1016/j.intcom.2011.02.009.

[35] S. Brin and L. Page, The anatomy of a large-scale hypertextual Web search engine, *Computer Networks and ISDN Systems* **30** (1998), 107–117. http://www.yahoo.com.
[36] T. Berners-Lee, Linked Data - Design Issues, 2006.
[37] C. Bizer, T. Heath and T. Berners-Lee, Linked data - The story so far, *International Journal on Semantic Web and Information Systems* **5** (2009), 1–22. doi:10.4018/jswis.2009081901.
[38] L. Po, N. Bikakis, F. Desimoni and G. Papastefanatos, *Linked Data Visualization: Techniques, Tools, and Big Data*, Morgan & Claypool Publishers LLC, 2020, pp. 1–157. ISSN 2160-4711. doi:10.2200/s00967ed1v01y201911wbe019.
[39] T. Berners-Lee and R. Cailliau, WorldWideWeb: Proposal for a HyperText Project, 1990.
[40] T. Berners-Lee, J. Hendler and O. Lassila, The Semantic Web, *Scientific American* (2001), 28–37. doi:10.1038/scientificamerican0501-34.
[41] T. Heath and C. Bizer, *Linked Data: Evolving the Web into a Global Data Space*, Synthesis Lectures on the Semantic Web, Morgan & Claypool Publishers, 2011. doi:10.2200/S00334ED1V01Y201102WBE001.
[42] R.W. Group, Resource Description Framework (RDF), 2014.
[43] G. Schreiber and Y. Raimond, RDF 1.1 Primer, 2014.
[44] T. O'Reilly, What is Web 2.0, 2005.
[45] D. Brickley and R.V. Guha, RDF Schema 1.1, 2014.
[46] E. Prud'hommeaux and A. Seaborne, SPARQL Query Language for RDF, 2008.
[47] T.W.S.W. Group, SPARQL 1.1 Overview, 2008.
[48] P. Hitzler, M. Krötzsch, B. Parsia, P.F. Patel-Schneider and S. Rudolph, OWL 2 Web Ontology Language Primer (Second Edition), 2012.
[49] F. Desimoni and L. Po, Empirical evaluation of Linked Data visualization tools, *Future Gener. Comput. Syst.* **112** (2020), 258–282. doi:10.1016/j.future.2020.05.038.
[50] T. Iba and T. Isaku, Holistic Pattern-Mining Patterns A Pattern Language for Pattern Mining on a Holistic Approach, in: *Proceedings of the 19th Conference on Pattern Languages of Programs, PLoP 2012*, 2012.
[51] D.E. Delano, *Patterns Mining*, Cambridge University Press, USA, 1998, pp. 87–95–. ISBN 0521648181.
[52] A. Seffah, The Evolution of Design Patterns in HCI: From Pattern Languages to Pattern-Oriented Design, in: *Proceedings of the 1st International Workshop on Pattern-Driven Engineering of Interactive Computing Systems*, PEICS '10, Association for Computing Machinery, New York, NY, USA, 2010, pp. 4–9–. ISBN 9781450302463. doi:10.1145/1824749.1824751.
[53] H. Glaser, H. Alani, L. Carr, S. Chapman, F. Ciravegna, A. Dingli, N. Gibbins, S. Harris, M.M.C. Schraefel and N. Shadbolt, CS AKTiveSpace: Building a Semantic Web Application, in: *The Semantic Web: Research and Applications, First European Semantic Web Symposium, ESWS 2004, Heraklion, Crete, Greece, May 10-12, 2004, Proceedings*, C. Bussler, J. Davies, D. Fensel and R. Studer, eds, Lecture Notes in Computer Science, Vol. 3053, Springer, 2004, pp. 417–432. doi:10.1007/978-3-540-25956-5_29.
[54] T. Berners-Lee, Y. Chen, L. Chilton, D. Connoll, R. Dhanaraj, J. Hollenbach, A. Lerer and D. Sheets, Tabulator: Exploring and Analyzing linked data on the Semantic Web, in: *Proceedings of the 3rd International Semantic Web User Interaction*, 2006. http://swui.semanticweb.org/swui06/papers/Berners-Lee/Berners-Lee.pdf.
[55] S. Davies, J. Hatfield, C. Donaher and J. Zeitz, User Interface Design Considerations for Linked Data Authoring Environments, in: *Proceedings of the WWW2010 Workshop on Linked Data on the Web, LDOW 2010, Raleigh, USA, April 27, 2010*, C. Bizer, T. Heath, T. Berners-Lee and M. Hausenblas, eds, CEUR Workshop Proceedings, Vol. 628, CEUR-WS.org, 2010. http://ceur-ws.org/Vol-628/ldow2010_paper17.pdf.
[56] D. Quan, D. Huynh and D.R. Karger, Haystack: A Platform for Authoring End User Semantic Web Applications, in: *The Semantic Web - ISWC 2003, Second International Semantic Web Conference, Sanibel Island, FL, USA, October 20-23, 2003, Proceedings*, D. Fensel, K.P. Sycara and J. Mylopoulos, eds, Lecture Notes in Computer Science, Vol. 2870, Springer, 2003, pp. 738–753. doi:10.1007/978-3-540-39718-2_47.
[57] S. Auer, S. Dietzold and T. Riechert, OntoWiki - A Tool for Social, Semantic Collaboration, in: *The Semantic Web - ISWC 2006, 5th International Semantic Web Conference, ISWC 2006, Athens, GA, USA, November 5-9, 2006, Proceedings*, I.F. Cruz, S. Decker, D. Allemang, C. Preist, D. Schwabe, P. Mika, M. Uschold and L. Aroyo, eds, Lecture Notes in Computer Science, Vol. 4273, Springer, 2006, pp. 736–749. doi:10.1007/11926078_53.
[58] B. Lee, C.S. Parr, C. Plaisant, B.B. Bederson, V.D. Veksler, W.D. Gray and C. Kotfila, TreePlus: Interactive Exploration of Networks with Enhanced Tree Layouts, *IEEE Trans. Vis. Comput. Graph.* **12**(6) (2006), 1414–1426. doi:10.1109/TVCG.2006.106.
[59] J. Davies and R. Weeks, QuizRDF: Search Technology for the Semantic Web, in: *37th Hawaii International Conference on System Sciences (HICSS-37 2004), CD-ROM / Abstracts Proceedings, 5-8 January 2004, Big Island, HI, USA*, IEEE Computer Society, 2004. doi:10.1109/HICSS.2004.1265293.
[60] G. Tummarello, C. Morbidoni and M. Nucci, Enabling Semantic Web Communities with DBin: An Overview, in: *The Semantic Web - ISWC 2006, 5th International Semantic Web Conference, ISWC 2006, Athens, GA, USA, November 5-9, 2006, Proceedings*, I.F. Cruz, S. Decker, D. Allemang, C. Preist, D. Schwabe, P. Mika, M. Uschold and L. Aroyo, eds, Lecture Notes in Computer Science, Vol. 4273, Springer, 2006, pp. 943–950. doi:10.1007/11926078_69.
[61] E. Hyvönen, E. Mäkelä, M. Salminen, A. Valo, K. Viljanen, S. Saarela, M. Junnila and S. Kettula, MuseumFinland - Finnish museums on the semantic web, *J. Web Semant.* **3**(2–3) (2005), 224–241. doi:10.1016/j.websem.2005.05.008.
[62] D. Huynh, S. Mazzocchi and D.R. Karger, Piggy Bank: Experience the Semantic Web Inside Your Web Browser, in: *The Semantic Web - ISWC 2005, 4th International Semantic Web Conference, ISWC 2005, Galway, Ireland, November 6-10, 2005, Proceedings*, Y. Gil, E. Motta, V.R. Benjamins and M.A. Musen, eds, Lecture Notes in Computer Science, Vol. 3729, Springer, 2005, pp. 413–430. doi:10.1007/11574620_31.

[63] E. Ikkala, E. Hyvönen, H. Rantala and M. Koho, Sampo-UI: A full stack JavaScript framework for developing semantic portal user interfaces, *Semantic Web–Interoperability, Usability, Applicability* (2020).

[64] N. Gibbins, S. Harris and M. Schraefel, Applying mSpace Interfaces to the Semantic Web, in: *World Wide Web Conference 2004 (01/01/04)*, 2003. https://eprints.soton.ac.uk/258639/.

[65] G. Vega-Gorgojo, M. Giese, S. Heggestøyl, A. Soylu and A. Waaler, PepeSearch: Semantic Data for the Masses, *PLOS ONE* **11**(3) (2016), 1–12. doi:10.1371/journal.pone.0151573.

[66] G. Vega-Gorgojo, L. Slaughter, B.M. von Zernichow, N. Nikolov and D. Roman, Linked Data Exploration With RDF Surveyor, *IEEE Access* **7** (2019), 172199–172213. doi:10.1109/ACCESS.2019.2956345.

[67] P. Haase, M. Schmidt and A. Schwarte, The Information Workbench as a Self-Service Platform for Linked Data Applications, in: *Proceedings of the Second International Workshop on Consuming Linked Data (COLD2011), Bonn, Germany, October 23, 2011*, O. Hartig, A. Harth and J.F. Sequeda, eds, CEUR Workshop Proceedings, Vol. 782, CEUR-WS.org, 2011. http://ceur-ws.org/Vol-782/HaaseEtAl_COLD2011.pdf.

[68] P. Heim, S. Hellmann, J. Lehmann, S. Lohmann and T. Stegemann, RelFinder: Revealing Relationships in RDF Knowledge Bases, in: *Semantic Multimedia, 4th International Conference on Semantic and Digital Media Technologies, SAMT 2009, Graz, Austria, December 2-4, 2009, Proceedings*, T. Chua, Y. Kompatsiaris, B. Mérialdo, W. Haas, G. Thallinger and W. Bailer, eds, Lecture Notes in Computer Science, Vol. 5887, Springer, 2009, pp. 182–187. doi:10.1007/978-3-642-10543-2_21.

[69] G. Burel, A.E. Cano and V. Lanfranchi, Ozone Browser: Augmenting the Web with Semantic Overlays, in: *Proceedings of the Fifth Workshop on Scripting and Development for the Semantic Web , co-located with 6th European Semantic Web Conference, SFSW@ESWC 2009, Heraklion, Greece, May 31, 2009*, S. Auer, C. Bizer and G.A. Grimnes, eds, CEUR Workshop Proceedings, Vol. 449, CEUR-WS.org, 2009. http://ceur-ws.org/Vol-449/Challenge1.pdf.

[70] F. Haag, S. Lohmann, S. Bold and T. Ertl, Visual SPARQL querying based on extended filter/flow graphs, in: *International Working Conference on Advanced Visual Interfaces, AVI 2014, Como, Italy, May 27-29, 2014*, P. Paolini and F. Garzotto, eds, ACM, 2014, pp. 305–312. doi:10.1145/2598153.2598185.

[71] M. Hearst, Design recommendations for hierarchical faceted search interfaces, in: *Proc. SIGIR 2006, Workshop on Faceted Search*, 2006, pp. 26–30.

[72] N. Seidel, Making web video accessible: interaction design patterns for assistive video learning environments, in: *Proceedings of the 20th European Conference on Pattern Languages of Programs, EuroPLoP 2015, Kaufbeuren, Germany, July 8-12, 2015*, C. Link and V. Eloranta, eds, ACM, 2015, pp. 17:1–17:16. doi:10.1145/2855321.2855339.

[73] N. Shadbolt and K. O'Hara, Linked Data in Government, *IEEE Internet Computing* **17**(4) (2013), 72–77–. doi:10.1109/MIC.2013.72.

[74] C. Keßler, M. d'Aquin and S. Dietze, Linked Data for Science and Education, *Semant. Web* **4**(1) (2013), 1–2–.

[75] E. Gamma, R. Helm, R. Johnson and J.M. Vlissides, *Design Patterns: Elements of Reusable Object-Oriented Software*, 1st edn, Addison-Wesley Professional, 1994. ISBN 0201633612. http://www.amazon.com/Design-Patterns-Elements-Reusable-Object-Oriented/dp/0201633612/ref=ntt_at_ep_dpi_1.

[76] M. Hassenzahl and N. Tractinsky, User experience - a research agenda, *Behav. Inf. Technol.* **25**(2) (2006), 91–97. doi:10.1080/01449290500330331.

[77] M. van Welie and G.C. van der Veer, Pattern Languages in Interaction Design, in: *Human-Computer Interaction INTERACT '03: IFIP TC13 International Conference on Human-Computer Interaction, 1st-5th September 2003, Zurich, Switzerland*, M. Rauterberg, M. Menozzi and J. Wesson, eds, IOS Press, 2003.

[78] N.B. Harrison, The Language of Shepherding: A Pattern Language for Shepherds and Sheep, *Pattern Languages of Program Design* **5** (1999), 507–530.

[79] U. Government, Regulated Food and Feed Products for Great Britain, 2021.

[80] O.E. Group, Portal de datos bibliográficos de la Biblioteca Nacional de España, 2021.

[81] D. for Environment Food & Rural Affairs, Hydrology Data Explorer, 2021.

[82] H. Rantala, I. Jokipii, M. Koho, E. Ikkala, J. Tuominen and E. Hyvönen, Building a Linked Open Data Portal of War Victims in Finland 1914-1922, in: *Proceedings of the Digital Humanities in the Nordic Countries 5th Conference, Riga, Latvia, October 21-23, 2020*, S. Reinsone, I. Skadina, A. Baklane and J. Daugavietis, eds, CEUR Workshop Proceedings, Vol. 2612, CEUR-WS.org, 2020, pp. 310–317. http://ceur-ws.org/Vol-2612/short17.pdf.

[83] P. Hassanzadeh, E. Hyvönen, E. Ikkala, J. Tuominen, S. Thomas, A. Wessman and V. Rohiola, FindSampo Platform for Reporting and Studying Archaeological Finds Using Citizen Science, in: *Proceedings of the Third Workshop on Humanities in the Semantic Web (WHiSe 2020) co-located with 15th Extended Semantic Web Conference (ESWC 2020), Heraklion, Greece, June 2, 2020 (online)*, A. Adamou, E. Daga and A. Meroño-Peñuela, eds, CEUR Workshop Proceedings, Vol. 2695, CEUR-WS.org, 2020, pp. 33–40. http://ceur-ws.org/Vol-2695/paper4.pdf.

[84] P. Leskinen and E. Hyvönen, Linked Open Data Service about Historical Finnish Academic People in 1640-1899, in: *Proceedings of the Digital Humanities in the Nordic Countries 5th Conference, Riga, Latvia, October 21-23, 2020*, S. Reinsone, I. Skadina, A. Baklane and J. Daugavietis, eds, CEUR Workshop Proceedings, Vol. 2612, CEUR-WS.org, 2020, pp. 284–292. http://ceur-ws.org/Vol-2612/short14.pdf.

[85] M. Warren and P.J. Hayes, Bounding Ambiguity: Experiences with an Image Annotation System, in: *Proceedings of the 1st Workshop on Subjectivity, Ambiguity and Disagreement in Crowdsourcing, and Short Paper Proceedings of the 1st Workshop on Disentangling the Relation Between Crowdsourcing and Bias Management (SAD 2018 and CrowdBias 2018) co-located the 6th AAAI Conference on Human Computation and Crowdsourcing (HCOMP 2018), Zürich, Switzerland, July 5, 2018*, L. Aroyo, A. Dumitrache, P. Paritosh, A.J. Quinn,

C. Welty, A. Checco, G. Demartini, U. Gadiraju and C. Sarasua, eds, CEUR Workshop Proceedings, Vol. 2276, CEUR-WS.org, 2018, pp. 41–54. http://ceur-ws.org/Vol-2276/paper5.pdf.

[86] U. Government, HM Land Registry Open Data, 2021.

[87] STLab, ODP Reactor, 2021.

[88] S. Castano, A. Ferrara and S. Montanelli, inWalk: Interactive and Thematic Walks inside the Web of Data, in: *Proceedings of the 17th International Conference on Extending Database Technology, EDBT 2014, Athens, Greece, March 24-28, 2014*, S. Amer-Yahia, V. Christophides, A. Kementsietsidis, M.N. Garofalakis, S. Idreos and V. Leroy, eds, OpenProceedings.org, 2014, pp. 628–631. doi:10.5441/002/edbt.2014.60.

[89] A. Katifori, C. Halatsis, G. Lepouras, C. Vassilakis and E.G. Giannopoulou, Ontology visualization methods - a survey, *ACM Comput. Surv.* **39**(4) (2007), 10. doi:10.1145/1287620.1287621.

[90] F. Gandon, A survey of the first 20 years of research on semantic Web and linked data, *Ingénierie des Systèmes d Inf.* **23**(3–4) (2018), 11–38. doi:10.3166/isi.23.3-4.11-38.

Table 2
Summary of the thematic analysis results of the pattern mining interviews.

#	Problem	Solution
AUTHORING		
1	Allow inexperienced users to create and edit Linked Data.	Provide features to create and edit data in the similar and familiar ways present in other web applications, and transform it in the background in RDF triples. Present clear boxes to users, so that the correct information is entered in the correct place.
2	—	Provide a feature to download the data present in search results, single resources, graphs, or tables.
3	Aid users when choosing what to input, while creating or editing data.	Provide recommendations of valid options for each text area available. Filter the recommendations to the most relevant according to the resource and domain, or by the previously used by the user.
VISUALISING		
4	Visualising all the triples associated with a resource can be too much information at once for the users.	Define for a resource the set of the most important properties and related resources to be displayed to the user. Divide them in different sections and order according to importance. Provide the remaining properties and related resources in closed sections for the user to consult if needed.
5	Mapping the data to visual information that can be easily understandable by the users.	Use primitive intuitive formats of visualisation (geographic data into maps, time data into calendars or timelines) and combine those different types of visualisations in one component.
6	Information is presented to the user scattered over several pages or sections, and thus can be hard to fully comprehend.	Aggregate sub-graphs of the data that contain the complete information for a use case or user task and create a visualisation component for that complex data aggregation. Know when information should be presented together and when information should be in another page navigable to by a link.
7	Finding a way to make the user understand the underlying structure of the dataset.	If not too dense, show a graph visualisation with the different classes in the dataset and the relations between them. Not easy if the application is not constrained to a specific domain or dataset, because an overview could mean an overview of the web of data.
8	—	Allow the user to choose from different visualisations of the displayed data (table, list, graph, map, calendar, picture, video).
9	—	Always have raw data available in addition to other kinds of visualisations. Have the option for users to be able to dig deeper into whatever direction they want.
10	Users have difficulty interpreting the labels of the resources or properties being displayed.	The definition of something, a resource or a property, should be available, for example when you hover the name.
BROWSING		
11	Due to the fact that everything is linked, it's difficult to provide effective navigation that is intuitive and keeps track of where the users come from. Otherwise, they can feel lost while browsing the data.	—
12	Allow the users to browse the data.	Provide links to the internal and external resources and properties.
SEARCHING		
13	—	Provide users with a SPARQL query editor in addition to the simple keyword-based search approach. Allow the user to visualise the corresponding SPARQL query of the search task performed.

Table 3

Visualising patlets - part 1.

Patterns	Problem
TEMPLATE	When resources have a long list of properties and related resources, define a set of the most important ones and display them to the user in the resource's page. In addition, define an outline dictating how the properties and related resources must be displayed to the user, they should be divided in different sections according to the different kinds of information presented. Provide the remainder of the properties and resources in a collapsible section that the users can consult when needed.
MULTI-VISUALISATION	When resources or sets of resources have different types of data associated, provide the user with different types of visualisations for that data: tables, lists, maps, graphs, calendars, graphics, timelines. This will provide the users with the visualisation that best suits each type of data and a variety of user needs and tasks. In addition to whatever set of visualisations chosen, always provide a list or table to ensure accessibility of the data.
SHOW DEFINITIONS	When the resources' names are not understandable by humans or are unknown terms to not domain experts, provide definitions for all resources and related properties of your data. Provide the definitions for either internal vocabularies and external vocabularies, when possible.
GRAPH VISUALISATION	Only use graph visualisations for subsets of the data where the edges of the graph all have the same or similar meaning, like family trees and friend networks. If the entire graph is not understandable without zooming in or out, the subset might be too big to be visualised in a graph. Provide clear labels for the nodes and edges and minimise overlapping.
SEARCH RESULTS VISUALISATION	Define the items that constitute the types of search results for your data. Identify the most important information for each type. Present the search result items in a list in differentiable sections, each with the identified information. Allow the user to change the order in which the results are presented, based on the associated data, on numeric fields, or on how close or far they are to a given location. The option to customise the order to the results will allow the users to more quickly find what they are looking for, and better compare the results according to certain properties.
RAW DATA	When an application is targeted at users with some knowledge of Linked Data technologies, provide the presented data in raw format. This will allow the users to consult this information if needed, enabling them to explore the data in whichever direction and to validate their interpretations of it.

Table 4
Visualising patlets - part 2.

Patterns	Problem
RAW DATA	When an application is targeted at users with some knowledge of Linked Data technologies, provide the presented data in raw format. This will allow the users to consult this information if needed, enabling them to explore the data in whichever direction and to validate their interpretations of it.
COMPLETE INFORMATION	When the information needed for a user task is a set of interconnected resources, previously aggregate those small subsets of data that form a complete piece of information and present it to the user.
VISUAL INFORMATION	Map the presented data to visual information according to their type, using familiar visualisation techniques. For numeric and textual data use lists and tables, for geographical data use maps, for temporal data use calendars and timelines, for statistical data use graphics. If your data represents concepts that are naturally related to themselves in a network format, like family trees or friend networks, use graph visualisations.
COLLECTION OVERVIEW	Provide a high-level overview of the data collection to the users, when the application is used by people with little to none experience on the domain. This will lead to a better understanding of the data, and consequently less time spent on tasks. When the application is not restricted to a domain, providing an understandable overview of the collection structure can be difficult. Opt for an alternative like resources' hierarchical trees instead.
HIERARCHICAL TREE	When the data collection is composed of a large number of classes and subclasses, in each resource's page provides its hierarchical tree. This will help the user situate the resource in the collection and gather more understanding of its overall structure.
KEYWORD QUERIES	When an application wants to provide all users with a familiar and easy to use approach to search the data collection, allow the users to perform search tasks using a keyword based approach. Due to the fact that keyword based queries are the most common search method, they should be provided in addition to more alternative and novel approaches, in order to be accessible to all users.

Table 5
Searching, browsing patlets.

Pattern	Problem
KEYWORD QUERIES	When an application wants to provide all users with a familiar and easy to use approach to search the data collection, allow the users to perform search tasks using a keyword based approach. Due to the fact that keyword based queries are the most common search method, they should be provided in addition to more alternative and novel approaches, in order to be accessible to all users.
SPARQL QUERIES	When an application provides its users with search features and already implements keyword queries or a similar simple search approach, provide a SPARQL query editor. This will allow the more experienced users to perform complex search tasks without any restriction that may be imposed by the more traditional approaches. It should also be possible to view the corresponding SPARQL query for search tasks constructed with a keyword approach and additional advanced options.
ADVANCED SEARCH	When an application provides its users with search features and already implements keyword queries or a similar simple search approach, provide a set of advanced search options. These options should be composed of various filters on different properties and characteristics of the data, tailored to the application domain. Make sure to match the filter to the corresponding type of data, using checkboxes to filter the classes of the results, dropdown lists when the field has a fixed set of options, sliders for filters on numeric fields, calendars to select dates or time intervals, and maps with area selection features to filter resources with geographical data.
LINK BROWSING	When an application wants to provide users with means to browse the data, allow users to traverse it following links and visiting the pages of internal and external resources. When designing a resource's page include the names of any related resources as hyperlinks to the corresponding pages.
BROWSE PATH	When browsing through Linked Data following links or with a similar approach, the possible paths between resources are enormous, and users may lose track of where they started and the path that was taken. Keep track of the users' path through the data and present it in a way that lets the user return to previous states or previously visited resources.
FACETED BROWSING	More traditional and familiar browsing approaches require an initial search query. When an application is targeted at users with little to no domain knowledge this can pose a challenge to them. Additionally to other browsing approaches, provide the user with a faceted browsing feature, based on the resources' classes. It should allow the user to filter the full list of resources in the application, according to their classes and the classes of the related resources. This approach will allow the users to start browsing without any previous knowledge and at the same time will contribute to their understanding of the structure of the collection.

Table 6
Authoring patlets.

Pattern	Problem
MULTIPLE INPUT BOXES	When an application wants to provide users with no knowledge of Linked Data technologies with methods that allow them to compose valid and semantically correct data, provide clear text areas for the user to input each part of the triple or triple set. The text areas should be provided for each component of the triples, subject, predicate, and object, and only accept valid options.
INPUT RECOMMENDATIONS	When an application provides inexperienced users with features to author data and implements an approach like multiple input boxes, provide a list of recommendations composed of only valid options for each one of the input boxes. A limited number of recommendations should be presented to the user in a dropdown list and it should be filtered as the user types in data. The initial limited list of recommendations can be composed of the most relevant resources or properties for the triple at hand, or the previously used ones by the user. A full list of the valid options should also be provided when needed.
INPUT EXAMPLES	When an application provides authoring features to users with no technical experience with RDF or Linked Data, they may encounter some difficulties in understanding what information should be inserted. Provide examples of similar resources or triples when a user edits or creates data. These examples can be composed of illustrative values tailored to the domain and to the resource or triple to be created or edited.
DOWNLOAD DATA	When an application's data is meant to be reused by other developers, researchers, or users, provide a download data button whenever a set of data is presented to the user. This download button can be added in scenarios like the following: the resources' pages, the search result lists, graphs, maps, tables, calendars, or other sets of related resources. If an application would like to give the users the possibility to download the entirety of the data collection, provide a visually different download button available in every page of the application.

Table 7
Pattern classification results.

UI Layer	User Interaction Patterns for Linked Data
1. The Conceptual Model	0.00%
2. Task Flow	0.00%
3. Organisational Model	10.00%
4. Viewing and Navigation	45.00%
5. Editing and Manipulation	5.00%
6. User Assistance	20.00%
7. Layout	20.00%
8. Style	0.00%
9. Text	0.00%

Table 8
Pattern classification results.

Pattern	Total	Positive	Neutral	Negative	Score
MULTI-VISUALISATION	20	20	0	0	20
SHOW DEFINITIONS	20	20	0	0	20
LINK BROWSING	20	20	0	0	20
KEYWORD QUERIES	20	19	1	0	19
ADVANCED SEARCH	20	19	0	1	18
FACETED BROWSING	20	18	2	0	18
SEARCH RESULTS VISUALISATION	20	18	1	1	17
BROWSE PATH	20	18	1	1	17
VISUAL INFORMATION	20	18	0	2	16
HIERARCHICAL TREE	20	16	4	0	16
SPARQL QUERIES	19	17	1	1	16
TEMPLATE	20	16	3	1	15
GRAPH VISUALISATION	20	17	1	2	15
COMPLETE INFORMATION	20	17	1	2	15
INPUT RECOMMENDATIONS	19	16	1	2	14
COLLECTION OVERVIEW	19	13	6	0	13
RAW DATA	20	15	1	4	11
DOWNLOAD DATA	19	14	2	3	11
MULTIPLE INPUT BOXES	19	14	0	5	9
INPUT EXAMPLES	19	12	2	5	7