

Ontology Mapping - A User Survey

Sean M. Falconer¹, Natalya F. Noy², and Margaret-Anne Storey¹

¹ University of Victoria, Victoria BC V8W 2Y2, Canada
{seanf, mstorey}@uvic.ca

² Stanford University, Stanford, CA 94305, USA
noy@stanford.edu

Abstract. Ontology mapping is the key to data interoperability in the semantic web vision. Computing mappings is the first step to applications such as query rewriting, instance sharing, web-service integration, and ontology merging. This problem has received a lot of attention in recent years, but little is known about how users actually construct mappings. Several ontology-mapping tools have been developed, but which tools do users actually use? What processes are users following to discover, track, and compute mappings? How do teams coordinate when performing mappings? In this paper, we discuss the results from an on-line user survey where we gathered feedback from the community to help answer these important questions. We discuss the results from the survey and the implications they may have on the mapping research community.

1 Introduction

Ontology mapping is a complex and largely user-driven process that can benefit from tool support. In the past few years, researchers have developed many tools and techniques for creating ontology mappings. Tools include PROMPT[12], COMA++ [3], Clío [10], Chimaera [9] and OWL Lite Alignment (OLA) [1]. Much research has been spent on developing the algorithms used by these tools, and indeed the authors of [13, 15] cite more than 20 different algorithms that can be used to generate candidate mappings. However, in most cases, the mapping process cannot be fully automated [2, 4, 5] and user input is required to accept, reject, or create new mappings. Despite the necessary role users play in the mapping process, there has been little work done to understand how and why users perform mappings.

In order to design more effective tools and algorithms, we claim that a deeper understanding of the interplay between tool, user, and the process is needed. For example, who are these users that are going to use the tools? Why do they need to perform mappings and for which domains? Do they use the currently available tools and if so, how do they use them? And, do these tools meet their needs? To answer these questions, we designed a survey and gathered feedback from the ontology mapping community. To our knowledge, this survey is the first specifically designed with these goals. The information gained from this survey should be valuable to both tool and algorithm designers. For example, in part as a consequence of this survey, we believe that the biggest gains in mapping will not come from improvements in the precision and recall in matching algorithms, but rather from better tool support.

The paper is organized as follows. We first present related work in Section 2. Next, we describe our survey design in Section 3. Following this, we present the results from the survey (Section 4). In Section 5, we discuss the results and implications for tool designers. We then describe limitations of the survey in Section 6. Finally, in Section 7, we present our concluding remarks and directions for future work.

2 Related work

Specifying mappings between one or more ontologies is well recognized to be a challenging and complex process that can be made significantly easier through tool support. The typical mapping process is an iterative procedure whereby the tool presents to the user a set of candidate mappings and the user then decides to accept and reject some of those mappings. The process is repeated until the user is satisfied that the mapping is sufficiently complete.

Determining candidate mappings is a challenging algorithmic problem. Consequently much of the research to date has been expended on designing more efficient and effective algorithms for determining candidate mappings. But much of the mapping process involves a tight collaboration and coordination between the user and tool. For example, the user must decide which mappings to accept and reject, keep track of progress, and determine when enough mappings have been completed for the intended purpose. These tasks are cognitively challenging but can be made easier through an improved partnership with the tool during the mapping process. Despite the gains that can be made across the entire mapping process, little research has focused on improving the effectiveness of the user decision process. Notable exceptions include user studies with PROMPT and Chimaera [6, 11], mapping experience reports [7, 14] and our own observational user study [5].

The user study experiment conducted with PROMPT concentrated on evaluating mapping suggestions provided by the tool by having several users merge two ontologies [11]. The number of steps, suggestions followed, suggestions that were not followed, and what the resulting ontologies looked like were all recorded. Precision and recall was used to evaluate the quality of the suggestions. The experiment only involved four users, which was too small to draw generalizable conclusions. Independently, PROMPT was evaluated, along with Chimaera by Lambrix and Edberg [6] with the purpose of merging ontologies from bioinformatics. Eight users were involved in the study, four with computer science backgrounds and four with biology backgrounds. The participants were given a number of tasks to perform, a hard copy user manual, and the software's help system for support. They were also instructed to "think aloud" and an evaluator took notes during the experiment. Afterwards, the users were asked to complete a questionnaire about their experience. The tools were evaluated with the same precision and recall measurements as used in the previously described PROMPT experiment, while the user interfaces were evaluated using the REAL (Relevance, Efficiency, Attitude, and Learnability) [8] approach. Under both criteria, PROMPT outperformed Chimaera, however, the participants found learning how to merge ontologies in either tool was equally difficult. The participants found it particularly difficult to perform non-automated procedures in PROMPT, such as creating user-defined merges.

Reed and Lenat [14] reported on their experiences with manually mapping the CyC ontology to other ontologies over a period of 15 years. The process relied on trained ontologists collaborating with domain experts. Over time, interactive clarification-dialog-based tools were developed to help ease the mapping procedure. The authors believed that the major barrier to the adoption of ontology mapping is the heavy reliance on someone setting up the source schemas and access protocols. They also stated that better tools were needed in order to allow domain experts to perform mappings rather than relying on ontology experts.

Lomax and McCray [7] described their experiences with mapping the Gene Ontology (GO) to the National Library's Unified Medical Language System (UMLS). The authors used a combination of methods to perform the mapping, starting with a preliminary exploration of both ontologies looking for overlap and then using an automated system to map 25% of the GO terms to UMLS terms. Following this, one of the authors visited the UMLS team for a month to work with the team in an attempt to complete the mapping. While many problems surfaced during this time, these were eventually addressed, debated and resolved. Through a combination of automated techniques, analysis and collaboration from the UMLS and GO teams, and mapping verification by humans, the GO ontology was fully mapped to UMLS.

In our previous work, we presented results from a user study where we observed teams participating in a "think-aloud" ontology mapping session with two different tools [5]. The goal of this study was to gain a deeper understanding of the user needs and how they could be met more effectively through tool support. Although the results of the study were informative, we were left with questions that required feedback from the ontology mapping community. For example, users had trouble remembering what mappings they had created or verified while working with the tools. Also, the participants in the study were frustrated by not knowing how much of the mapping task they had already completed and what was left to be completed. The participants also had difficulty learning and working with the tools, which reinforces the findings reported in the Lambix and Edberg study. The participants in our study were not typical ontology mapping tool users and indeed were trained to use the tools before the study. Hence, we are interested in discovering if the problems encountered by our test users are also experienced by ontology mapping tool users with pragmatic and pressing needs for ontology mappings.

Despite some preliminary work in this area of understanding mapping tool users, we believe there is a lack of knowledge about the tools currently used, the users themselves, and the problems faced during the mapping process. Hence, we designed the survey that is presented in the next section of this paper to gain more insight into these questions.

3 The survey

We conducted a user survey to gain a basic understanding of why people perform ontology mappings, the tools they use, problems they experience, and what processes, if any, they follow.

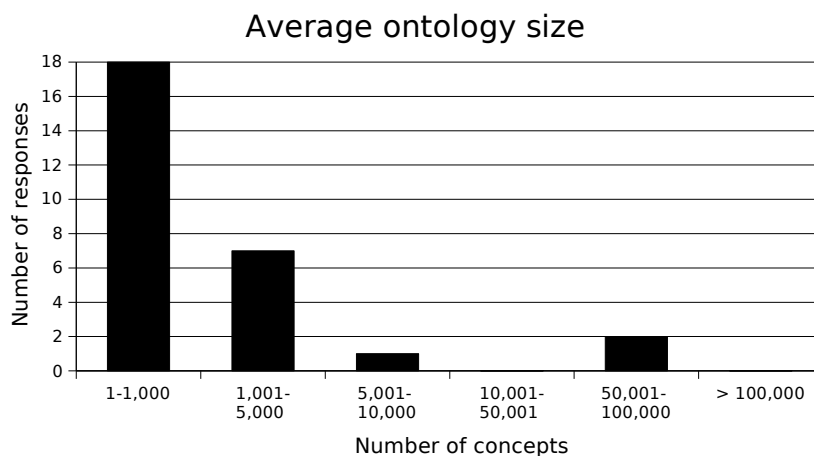


Fig. 1: Size of ontologies being used.

3.1 Participants

The survey was made available online for a two week period to both industry and academia via semantic web related mailing lists and news groups. Twenty-eight people responded to the survey.

3.2 Survey design

We designed the questions, organization, and presentation of the survey with the help of several researchers with human-computer interaction and ontology-mapping experience. The survey consisted of **multiple-choice** and **open-ended** questions (Table 1). Ten of the questions were completely open-ended, three were multiple-choice, and five were multiple choice answers that also allowed an open-ended response.

The questions can be classified into three categories. Questions 1 through 5 were **user context questions**, questions 6 through 8 and 17 were **tool questions**, and questions 9 through 16 were **process questions**. The user context questions were designed to gather data relating to the use cases for ontology mapping. The tool-related questions were designed to explore our questions regarding which tools people are using and whether the users find the tools useful. The process-related questions investigated the difficulties with performing mappings, whether users are working in team environments, and what processes people are following for coordinating their mapping efforts. Table 1 lists the specific questions that were asked.

4 Survey results

4.1 User context questions

The first question asked participants what domains of ontologies they work with. We had participants from a variety of domains, the most popular being biomedical, media, information-system design, business, and travel. Several respondents worked with

Table 1: Survey questions

No.	Question	Response
<i>User context questions</i>		
1	What are the domains of ontologies that you work with?	open-ended
2	What is the average size of the ontologies you work with?	multiple-choice
3	What type of ontologies do you work with?	combined
4	How often do you create/edit ontology mappings?	open-ended
5	What are these mappings used for?	combined
<i>Tool questions</i>		
6	What mapping tools have you used?	combined
7	Of the tools you've used, what do you find most useful and why?	open-ended
8	Of the tools you've used, what do you find to be deficient about these tools?	open-ended
17	If you were to design your perfect mapping tool, what features will it have?	open-ended
<i>Process questions</i>		
9	Do you find it difficult to create an ontology mapping?	multiple-choice
10	If you answered "Yes" to question 9, please explain why.	open-ended
11	What process do you use when performing mappings?	combined
12	How do you remember which mappings you have created/verified versus mappings that are left to create/verify?	combined
13	When do you consider that the mapping is complete?	open-ended
14	Do you experience problems while performing ontology mappings? If so, please explain.	open-ended
15	How many people participate in the creation process of the mappings you perform?	multiple-choice
16	If the number of participants is greater than one, please briefly explain how the work is coordinated.	open-ended

different domains of ontologies, either for research purposes or based on their current project.

Figures 1 and 2 show the distribution of results for the questions on the size of the ontologies and the ontology languages used. Most participants worked with ontologies with less than 1000 concepts and the primary ontology languages are OWL and RDFS.

In the next user context question, we looked at how often participants create and edit ontology mappings. The responses varied, the most popular being that mappings were created either per-application or as often as their work dictated that a mapping was required. Two of the respondents indicated they performed this operation often, but for research and tool testing purposes.

The final user context question asked why the participants created mappings. Participants were allowed to select from a list of responses (as many that applied) and also provide a free-form response in an "Other" field (see Figure 3). The use cases listed were "Instance data sharing", "Ontology merging", "Query translation", "Web service integration", and "Other". Two of the "Other" responses stated that they create mappings purely for research purposes, while one was for automatic SQL generation.

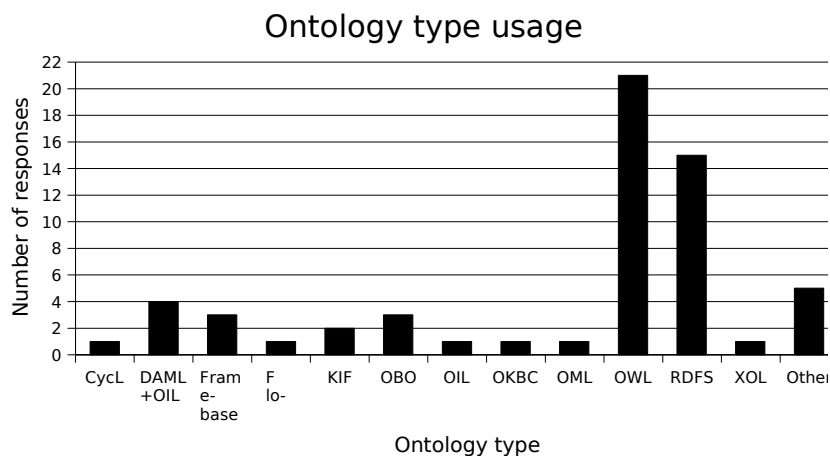


Fig. 2: Ontology language usage.

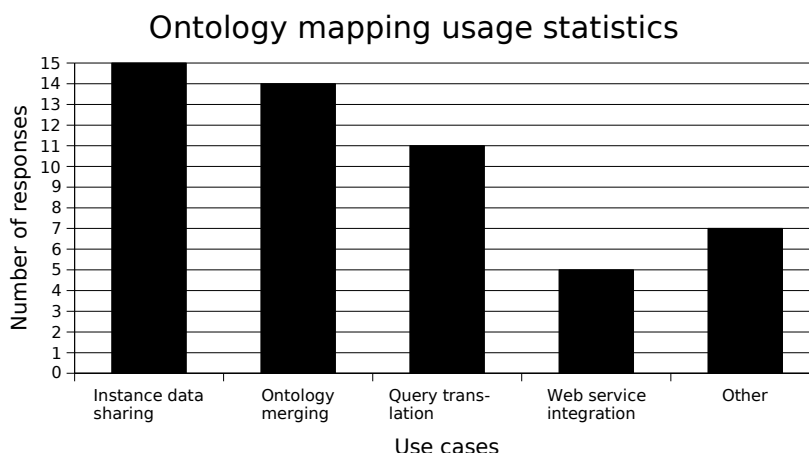


Fig. 3: Ontology mapping use cases.

4.2 Tool questions

The first tool-related question asked which tools users had tried. Respondents could choose from seven tools: Chimaera, COMA++, FOAM, MoA Shell, OLA, PROMPT, and QOM. They could also list any other tools in the “Other” field. Each of the listed tools was used by at most one to three participants with the exception of MoA Shell, which none of the respondents used. As shown in Figure 4, no tool was particularly dominant. The bulk of the feedback came in the “Other” category, which had 17 participant responses.

Other tools included Protégé, Internet Business Logic, AUTOMS, Crosi, WSMT with Ontostudio, X-SOM, OMAP, Falcon-AO, HMatch, and Snoggle. Each of these tools were used by only one participant, except X-SOM, which had been used by two.

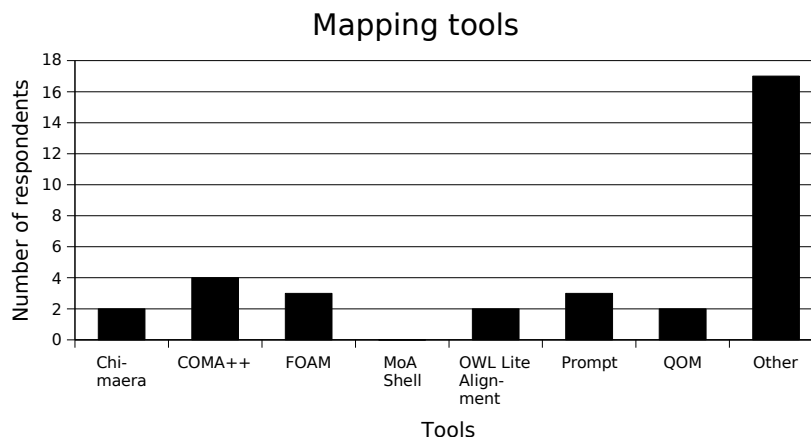


Fig. 4: Ontology mapping tools.

Two participants indicated that they use a custom built solution, while one indicated that they use a completely manual process.

We asked which tools and features participants found most useful and what deficiencies they found with the tools. Both Crosi and COMA++ were found to be useful because they integrate a large variety of similarity algorithms and are available online. One user indicated that they like tools to provide simple suggestions and automatic help, while another user had a contrasting view, stating that they like statistically-based tools because others require too much designer opinion. Protégé was highlighted as being good for manual creation of mappings as it makes it easy to create ontologies. Several participants pointed out that many tools are too general and are built without domain-specific mechanisms. One of the custom built solutions was indicated to be similar to PROMPT, but was built to take advantage of domain knowledge, specifically term normalization algorithms and synonyms for their domain of interest. The requirement for the tools to incorporate domain-specific analysis and features was a common theme in response to several questions in the survey. Another common theme was the lack of visual displays or easy to use tools. Specifically, one participant indicated that PROMPT's interface was too complicated to give to a domain expert to do the mapping. One user criticized specific tools for their lack of documentation, for being buggy, and for not working as described. Other interesting observations were the lack of precision and recall for mappings in real world ontologies and that the tools do not allow for expressive enough mappings (e.g. some tools only support 1-1 mappings).

In the final tool-related question, we asked the respondents to describe which features the perfect mapping tool would have. In the presentation of the survey, this question came at the end, but here we categorize it as a tool question. There were several interesting themes that came up in the responses.

The first theme was that six of the 19 responses related to the desire for *better and easier to use tools*. Specifically, participants stated that they needed better interfaces, graphical cognitive support, improved user interactivity, and facilities for explaining manual mappings. Users highlighted a large number of *desired features for*

the algorithms for generating candidate mappings: powerful and pluggable matching algorithms, recognition of re-occurring patterns of representing information, identification of not only simple correspondences but also of complex ones, and extending beyond mere word-pair associations and semantics. Four of the responses stated the requirement for perfect precision of recall for the mapping algorithms. Three participants also indicated that they want better facilities for *testing mappings and support for more expressive mappings*. The final interesting theme was *collaborative support*. Most of the respondents indicated that they work on their mappings in teams (see *process questions*). Most available tools do not support this type of team development scenario.

4.3 Process questions

The first process-related question asked whether the participants found the creation of an ontology mapping difficult. 30.8% replied “No” to this question, while 69.2% said “Yes”. The follow-up question to those participants that answered “Yes”, asked participants to explain why they found the process difficult. Ten of the 21 responses discussed semantic issues, such as the process being too subjective or ambiguous. One participant pointed out that the *“semantics of the underlying ontologies are not usually well defined. Without a solid understanding of the semantics, it is almost impossible to perform the mapping correctly.”* Respondents also discussed a lack of domain expertise for performing mappings, and that *“[y]ou have to get into the brains of the original developers of the ontologies being mapped.”* Participants also stated that tools are not flexible enough for application-specific needs, resulting in the manual creation of mappings, which is neither scalable nor efficient. One participant indicated that the OWL primitives for expressing mappings are poor and that users are faced with making difficult decisions when two related concepts *“almost but not exactly match.”* Three participants also indicated that problems with creating mappings resulted from poorly designed and documented tools.

We next asked participants what process they followed when performing mappings. Available responses were “Tackle the easy matches first” (37.0%), “Focus on a familiar area of the ontology” (51.9%), “Rely on the tool’s ordering of mappings” (14.8%), and “Other” (22.2%). Responses for the “Other” category included performing an automated matching up front and then a debugging step, while two of the responses indicated that they first applied lexical, then structural, and finally semantic methods.

In the next question, we asked how the participants remembered mappings they had created or verified. Most respondents chose from the provided answers, “The tool supports this” (37.0%) and “Paper and pencil” (55.6%), while 22.2% filled out the “Other” option. In the “Other” responses, one user indicated that the tool they use supports this, which works well when mappings are done in a single pass, but extra help is needed for multiple passes. Another respondent indicated that they use their own codes to report the mappings they create, which is similar to tracking the information by paper and pencil. Finally, one respondent indicated that they did not follow any good process. It is interesting that the majority resort to tracking this information manually by paper and pencil. Similar types of changes exist in software development and most IDEs and source control systems handle the tracking of this data.

We then asked when the participants considered the mapping to be complete. Ten of the 25 responses indicated that they used some form of testing (automated or manual) to verify that the mapping was completed to their satisfaction. For many respondents, this testing meant determining whether the mapping supports whatever application they were working on. Five responded that they knew the mapping to be complete when all concepts had been mapped. However, this implies either a perfect mapping, or that they knew when all reasonable concepts had been mapped. Interestingly, three participants responded that they never knew when the mapping was complete. Only one respondent indicated that they relied on tool support for determining whether the mapping was complete, although one participant stated that they must trust the system when mapping large ontologies because verification by hand is too slow.

We also asked participants about the types of problems they experienced while performing mappings. Similar issues outlined in previous questions came up again. Specifically, one respondent stated that *“most ontology tools are difficult for business users to understand.”* Testing the mapping was also a popular theme along with issues related to the problem that people model conceptualizations differently.

The final two questions dealt with whether participants worked in teams and what sort of process they followed for coordinating their efforts. 53.9% indicated that only 1 or 2 people were involved in the mapping process, 42.3% worked in groups of 3 to 5 people, and finally, 3.9% worked with 6 to 10 people. Based on results from the team process question, we were able to determine that of the 53.9% working in teams of 1 to 2 people, 53.8% of these actually work in a team of 2, which means that only 24.0% work completely on their own. It is interesting that these users felt that there could be an automatic algorithm that they would trust completely.

The team-coordination processes ranged from weekly meetings to collaborating through wiki's to coordinating through CVS. 18 of the 20 respondents relied on non-software solutions for managing the team or a combination of *ad hoc* communication strategies like CVS, wiki's and e-mail along with meetings. Smaller teams typically had one team lead and one implementer, and coordinated with face to face meetings. Participants also indicated that they worked as a group or partitioned the ontologies and then performed a group check to validate the mapping. Some teams used domain expertise for reviewing the composed mappings or during the mapping process for input. One respondent indicated that they use a “brainstorming” team process for coordinating the mapping effort.

5 Discussion of results

We found it surprising how many tools had been tried by our respondents. There has clearly been a large effort from the research community to develop so many tools, yet there does not appear to be a dominant tool that is a benchmark for mapping tool design. This may be due in part to the variety of user needs. Some respondents highlighted that they had domain-specific needs or that existing tools do not support sophisticated enough mappings.

Most of the problems, deficiencies, and issues with ontology mapping uncovered by the survey can be classified into one of two categories: fundamental issues with

language and semantics, and tool-specific issues. Fundamental issues, such as different model conceptualizations and language ambiguity, are difficult, if not impossible problems to solve. It is interesting that some of the responses to the “perfect mapping tool” question were that the tool would have 100% precision and recall or full natural language understanding. While a perfect, fully-automated solution would be ideal, it is probably not realistic for any but the most simple, straight-forward mappings. As the survey also highlighted, many of the problems that people face in ontology mapping, are difficult for even a team of human experts to resolve.

It is also interesting that these users felt that there could be an automatic algorithm that they would trust completely. Issues of trust also came up in our previously mentioned user study. Although our users stated that they liked PROMPT’s explanation facility, they were also often confused when it made a suggestion that was obviously wrong. Incorrectly generated candidate mappings would sometimes lead to our users ignoring the suggestions and switching to a completely manual process.

Tool specific issues such as better user interfaces, graphical support, better testing facilities (data translation based on mappings), interactivity, algorithm explanation capabilities, and so forth are all problems that ontology-mapping tool developers can help with. As discussed in the results, one respondent indicated that PROMPT was too complex to give to their end-users. This sentiment was also echoed by a non-computer science participant in our user study. Mapping is a complex process, it is important that we do not further burden our users with learning a difficult to use software suite. Instead we must support them via the software. In agreement with Bernstein *et al.* [2], we believe that at this point the biggest productivity gains in mapping tasks will come from better cognitive support rather than from an improvement of precision and recall in matching algorithms.

The responses to the process-related questions brought up several interesting issues that tool designers and researchers may also need to address. First, it is noteworthy that many of the participants relied on paper and pencil to remember the mappings they perform. One individual even noted that some tools work for a first pass, but then they “forget” the steps previously performed during the second pass. Tool support should be able to address this issue. Second, it appears that most users work in small teams but tools currently lack support for team communication and collaboration, as well as for partitioning the mapping process into manageable chunks that can be tackled by individuals on the team. Many teams work together to validate the prepared mappings. Again, tool support could help with team work. Metadata annotations (perhaps visualized via color-coding) could be used to record who composed the mapping and why they made certain decisions. In addition, the ontology mapping community could borrow ideas from the Web 2.0 social networking community³. E.g., tools could support voting on mappings, commenting on and annotating mappings, and associating instance data with a conceptualization. There has been some experimentation with community-driven ontology mapping [16], but tool support is currently limited.

³ <http://del.icio.us>, <http://www.flickr.com>

6 Limitations of the survey

There are of course limitations to this study, the first and foremost being the sampling size and population. Although we feel that 28 respondents gave us a wide variety of interesting and useful responses, with such a small sample it is possible that our responses are biased. Also, as we solicited participants from mailing lists, most of which were academically oriented, our sample may be biased towards researchers in the area rather than a balance between those working in research and industry. Finding and recruiting users from our target population was also an issue, because it is difficult to know how to best reach ontology mapping tool users.

As with any on-line survey, the wording of some of our questions may have potentially been confusing to some participants. For example, we asked “If you were to design your perfect mapping tool, what features will it have?”. We stated that some respondents indicated full natural language understanding and 100% precision and recall. Perhaps if the question had been worded differently to solicit feedback on a “realistic ideal mapping tool” the responses would have been different.

7 Conclusion and future work

In this paper, we discussed the design and results from our on-line user survey. The open-ended responses gave us valuable information about the types of problems users are experiencing, what features they’d like to see improved, some insight into their mapping and team process, and which tools are being used by the community.

Several issues raised by our participants indicate that their problems could be solved by improved tools and this would lead to better mapping results as well as greater adoption of the various mapping tools. This area of research has seen little activity to date. In the future, we plan to continue gathering feedback from the user community and carry out observational user studies. Our goal is to develop a comprehensive theoretical framework for cognitive support in ontology mapping. It is our hope that this will help guide the design and evaluation of future mapping tools such that the user’s role in the mapping process is fully supported.

Acknowledgments

This work was supported by the National Center for Biomedical Ontology, under roadmap-initiative grant U54 HG004028 from the National Institutes of Health. We wish to thank Philip Bernstein for his helpful suggestions about where to post our survey requests. We also wish to thank Chris Bennett and Jody Ryall for their editing suggestions.

References

1. Ontology alignment with *ola*. In *Evaluation of Ontology-based Tools*.
2. P. A. Bernstein and S. Melnik. Model management 2.0: Manipulating richer mappings. In *SIGMOD’07*, June 2007.

3. H.-H. Do. *Schema Matching and Mapping-based Data Integration*. PhD thesis, Department of Computer Science, Universität Leipzig, 2006.
4. S. M. Falconer, N. F. Noy, and M.-A. Storey. Towards understanding the needs of cognitive support for ontology mapping. In *Ontology Matching Workshop*, 2006.
5. S. M. Falconer and M.-A. Storey. Cognitive support for human-guided mapping systems. Technical Report DCS-318-IR, University of Victoria, Victoria, BC, Canada, 2007.
6. P. Lambrix and A. Edberg. Evaluation of ontology merging tools in bioinformatics. In *Proceedings Pacific Symposium on Biocomputing*, pages 589–600, Kauai, Hawaii, USA, 2003.
7. J. Lomax and A. T. McCray. Mapping the gene ontology into the unified medical language system: Research papers. *Comp. Funct. Genomics*, 5(4):354–361, 2004.
8. J. Löwgren. *Human-computer Interaction. What every system developer should know*. Studentlitteratur, Lund, 1993.
9. D. L. McGuinness, R. Fikes, J. Rice, and S. Wilder. The chimaera ontology environment. In *Proceedings of AAAI*, pages 1123–1124, 2000.
10. R. J. Miller, L. M. Haas, and M. A. Hernández. Schema mapping as query discovery. In *VLDB 2000*, pages 77–88, 2000.
11. N. F. Noy and M. A. Musen. Evaluating ontology-mapping tools: Requirements and experience. In *Proceedings of OntoWeb-SIG3 Workshop*, pages 1–14, 2002.
12. N. F. Noy and M. A. Musen. The PROMPT suite: Interactive tools for ontology merging and mapping. *International Journal of Human-Computer Studies*, 59(6):983–1024, 2003.
13. E. Rahm and P. A. Bernstein. A survey of approaches to automatic schema matching. *VLDB Journal: Very Large Data Bases*, 10(4):334–350, 2001.
14. S. L. Reed and D. B. Lenat. Mapping ontologies into cyc. In *Proceedings of AAAI*, 2002.
15. P. Shvaiko and J. Euzenat. A survey of schema-based matching approaches. *Journal on Data Semantics*, pages 146–171.
16. A. V. Zhdanova. Towards a community-driven ontology matching. In *K-CAP '05*, pages 221–222, 2005.