

Interactive knowledge capture in the new millennium: how the Semantic Web changed everything

YOLANDA GIL

Information Sciences Institute, University of Southern California, 4676 Admiralty Way, Suite 1001, Marina del Rey, CA 90292, USA;
e-mail: gil@isi.edu

Abstract

The Semantic Web has radically changed the landscape of knowledge acquisition research. It used to be the case that a single user would edit a local knowledge base, that the user would have domain expertise to add to the system, and that the system would have a centralized knowledge base and reasoner. The world surrounding knowledge-rich systems changed drastically with the advent of the Web, and many of the original assumptions were no longer a given. Those assumptions had to be revisited and addressed in combination with new challenges that were put forward. Knowledge-rich systems today are distributed, have many users with different degrees of expertise, and integrate many shared knowledge sources of varying quality. Recent work in interactive knowledge capture includes new and exciting research on collaborative knowledge sharing, collecting knowledge from Web volunteers, and capturing knowledge provenance.

1 Introduction

For this special anniversary issue of the *Knowledge Engineering Review*, I prepared a personal perspective on recent research in the area of interactive knowledge capture. My research interest has always been human–computer collaboration, in particular how to assist people in performing complex, knowledge-rich tasks that cannot be fully delegated to a computer program. This raises many research questions: How can people provide a computer with enough knowledge about the task domain and the user’s context to provide adequate assistance? How could we make it easier to provide the knowledge needed? How do we know that the computer understands the knowledge that it is given and can make effective use of it? How can a computer be proactive in asking questions about what it does not know and is needed to improve its behavior? All these questions span a number of research areas including knowledge capture (or knowledge acquisition), intelligent user interfaces, knowledge representation and reasoning, and problem solving and decision-making, and they lie at the intersection of artificial intelligence, cognitive science, and human–computer interaction.

The landscape of this research area has radically changed with the advent of the Web, and of the Semantic Web in particular. We used to assume that there would be a single user editing the knowledge base, that he or she would have domain expertise to add to the system, that he or she would have received some training in advance, and that the system would have a centralized knowledge base and reasoner. Today, these assumptions are no longer a given. We have knowledge bases that are distributed, that are edited by groups of users with varying expertise and training, and that contain knowledge of varying quality. These very radical changes have posed new challenges, and at the same time have revived the field and expanded it well beyond the realm of artificial intelligence research.

Assumptions of Early Knowledge Systems	New Challenges in the Semantic Web	Current Research Topics
<input type="checkbox"/> One user	<input type="checkbox"/> Many, many, MANY users	<input type="checkbox"/> Harnessing volunteer contributions
<input type="checkbox"/> User must be trained	<input type="checkbox"/> No time for training	<input type="checkbox"/> Natural tutorial instruction
<input type="checkbox"/> One system	<input type="checkbox"/> Many distributed subsystems	<input type="checkbox"/> Meta-level reasoning
<input type="checkbox"/> One knowledge base	<input type="checkbox"/> Many knowledge bases	<input type="checkbox"/> OWL and semantic web languages
<input type="checkbox"/> One reasoner	<input type="checkbox"/> Many reasoners	<input type="checkbox"/> Semantic web reasoners
<input type="checkbox"/> Expert user	<input type="checkbox"/> Varying quality and coverage	<input type="checkbox"/> Provenance and trust

Figure 1 The early assumptions made in early knowledge systems were challenged by the Semantic Web, bringing about new research topics for the interactive knowledge capture community

This paper describes the evolution of the field as a personal journey, describing solely work within my research group and with collaborators. Although it focuses on my work, the journey should be representative of the changes that the field has undergone and the trends in the research emphasis of other researchers. Figure 1 gives an overarching guide to the topics discussed in the paper. We begin with a journey to the early days of knowledge acquisition.

2 Early work: the knowledge acquisition bottleneck

Our early research investigated knowledge acquisition in the context of a user teaching a computer how to perform the aspects of the task that he or she wanted to delegate. We assumed the system had a knowledge base that the user was modifying and extending. This was dubbed as ‘the knowledge acquisition bottleneck’ in the 1980s, and there were a number of knowledge acquisition tools developed to address it. Each tool was developed for a specific type of problem-solving task. An early example was SALT, developed by expert system pioneer John McDermott and colleagues to acquire knowledge for constraint satisfaction problem solvers (Marcus & McDermott, 1989).

Our research demonstrated that this proliferation of knowledge acquisition tools was unnecessary and that a single general-purpose knowledge acquisition system with an underlying generic problem solver was sufficient, as long as the problem-solving method at hand could be represented in the problem solver’s language and the knowledge acquisition system could reason about it with respect to the knowledge that needed to be acquired. In Gil and Melz (1996), we showed that all the questions that SALT was designed to ask from users were automatically generated by our EXPECT knowledge acquisition system, and that additional questions were created by EXPECT to account for gaps in the knowledge base that SALT did not address. This meant that there was no need to develop a different system for each style of reasoning, rather that one single knowledge acquisition system with more explicit representations was sufficient.

Another issue that we investigated is maintaining a knowledge base consistent through many individual changes (Gil & Tallis, 1997; Tallis & Gil, 1999). We found that the more the knowledge acquisition system can reason about different kinds of gaps in its knowledge, the better it can hypothesize how those gaps could be filled and therefore be more proactive and cause less burden to the user. This was reported in Blythe *et al.* (2001), the last paper written about the EXPECT project and one that received a Best Paper award at the 2001 Intelligent User Interfaces conference.

Through this work, we found that we needed a methodology to assess the quality of knowledge acquisition systems, and were surprised there was not one for this area of research. We designed a methodology for user studies that drew from and extended methods from the literature on evaluating programmers, tutoring systems, and user interfaces (Tallis *et al.*, 2001). We have used this methodology over the years to evaluate a number of knowledge acquisition systems in a variety of application domains, with users ranging from biologists to military officers to project assistants (Kim & Gil, 1999, 2000, 2001; Clark *et al.*, 2001).

Thanks to these user evaluations and experiments, we were able to obtain important quantitative findings. We learned that users were able to complete complex tasks only using the assistance of a

knowledge acquisition system, that users enter knowledge much faster, and that knowledge acquisition systems help regardless of the level of programming knowledge or sophistication of the user.

3 From a single user to millions of knowledge contributors

Unlike our earlier assumption that a single user would be teaching, there are now vast amounts of people on the Web interested or willing to provide knowledge to computers, whether it be deep domain expertise or personal preferences. Myriads of Web users contribute structured content that ranges from simple labels or tags to complex mashups.

Our research in this area focused on how to acquire common knowledge from volunteer contributors. A broad repository of common knowledge would allow knowledge systems to be less brittle. We investigated how to keep volunteers engaged in contributing, and developed techniques to validate the knowledge acquired (Chklovski & Gil, 2005b). An interesting aspect of this work is to accommodate natural language statements, which are very easy for users to provide, while constraining them to be in a form that makes their processing possible. We also designed user interfaces that facilitated the collection of specific types of knowledge from users, including process knowledge and argumentation structures (Chklovski & Gil, 2005a, 2005c). By 2006, our knowledge repository contained more than 700 000 statements concerning common knowledge, becoming one of the top 100 largest OWL documents on the Web (Ebiquity, 2006).

It became clear very quickly that there are very few universally agreed-upon truths and that accommodating different views is important in an open knowledge system such as the Web. We are currently working on novel frameworks to create knowledge resources that incorporate alternative views within a community of contributors (Vrandečić *et al.*, 2010).

4 No training required

When teaching a computer a complex task, it is important to receive some training ahead of time concerning the abilities and limitations of the system, the methods to test that it has understood correctly, and most of all the use of a language to communicate knowledge. Long periods of training get in the way of doing the teaching and are generally not acceptable to users. In addition, users much prefer to communicate knowledge in a manner that is natural to them.

Some projections argue that 90 million end users without having programming expertise will be interested in creating applications of reasonable complexity such as spreadsheets, websites, and computer game content by 2012 (Scaffidi *et al.*, 2005). Their ‘programs’ will be shared with others and extended or adapted to other problems. The focus of our research in this area is learning procedure descriptions from users through tutorial instruction. An extensive survey of this area revealed that people find it very natural to teach procedures by giving general instructions but that typically give poor instruction. We used descriptions of data analysis workflows built by biologists and shared in a social website, and developed a knowledge acquisition system that can take such descriptions in natural language and create workflows for the user (Groth & Gil, 2009a, 2009b).

We are currently developing an approach to learn procedures from human tutorial instruction given in natural language, where the system uses heuristics to handle ambiguities and other imperfections in the instruction (Gil *et al.*, 2011).

5 Functioning with limited knowledge

Our knowledge acquisition systems enabled unsophisticated users to enter knowledge about complex processes in expert-level domains. Alas, although the knowledge was mostly correct, the missing bits did not allow the system to reason appropriately about the task. Perfect knowledge cannot be a requirement. Knowledge systems must be able to operate despite containing varying quality and coverage. The Web is full of broken links, which was traditionally considered unacceptable in hypertext circles. Similarly, knowledge systems cannot assume that the knowledge will

be consistent and complete and that it can always be subject to such assumptions typically made for reasoning in first-order logic.

We turned to research on human learning to understand the principles and cognitive approaches that humans take when acquiring knowledge incrementally through a lesson (Kim & Gil, 2002, 2003; Gil & Kim, 2002). We studied this from the student's perspective as well as from the tutor's perspective regarding the organization of lessons and the expectations of the student's ability to reason with limited knowledge.

In order to provide a computer with the ability to understand what it knows and what knowledge it is missing, we are investigating meta-reasoning capabilities that reason over the contents of a knowledge base and detect gaps and inconsistencies. At any given point, the system should be aware of its knowledge gaps and decide whether they affect its ability to do specific kinds of problem solving. We have developed a framework to keep track of multiple alternative models of a procedure that evolve over time as additional information is provided (Kim & Gil 2007; Gil *et al.*, 2010a; Kim *et al.*, 2010). These techniques are used in our work on learning procedures from natural instruction mentioned above (Gil *et al.*, 2011).

6 Distributed problem solving

In contrast with early knowledge-based systems, today both knowledge and reasoning tend to be physically and organizationally distributed on the Web. Often, the knowledge sources needed for a task reside in distributed subsystems. Reasoning functions may also be distributed, for example, as provided by third-party services. Knowledge development may also be distributed, formed by assembling knowledge bases (or ontologies) that are separately developed.

Our research is looking at distributed knowledge sources and distributed reasoning in multi-institutional scientific collaborations. Some institutions may provide access to their data as a repository or as a remote service, while others may contribute software and tools to process data. Scientists in collaboration develop applications that retrieve, process, characterize, and store data. Our work focuses on computational workflows to represent declaratively the application software components and the dataflow among them (Gil, 2006). We have developed semantic representations of workflows that can be used to reason about the data and the processing steps to automatically generate workflows as well as metadata descriptions of new data products that result from the computations (Gil, 2009). Our algorithms to reason about semantic workflows assume that data repositories and software repositories are distributed and managed separately from the workflow system (Gil *et al.*, 2010b). This is a major departure from today's planning and reasoning algorithms, which assume that all the pertinent axioms are accessible to a centralized reasoner. Our algorithms call out to the external repositories with well-defined requests, which cause data-specific or component-specific reasoning at the remote sites.

A major component of this research is a collaboration with grid computing researchers to combine workflow reasoning with workflow execution planning (Deelman *et al.*, 2003, 2005; Maechling *et al.*, 2005; Gil *et al.*, 2007a, 2007b; Kumar *et al.*, 2009). We have applied these ideas to improve compiler optimization (Hall *et al.*, 2008), resource-bound computations (Kumar *et al.*, 2010), and scientific application performance (Woollard *et al.*, 2008).

Finding relevant components in distributed architectures is a problem for users. We have developed matchers that are given a natural language description of a task by a user and find relevant agents available to automate that task (Chalupsky *et al.*, 2001; Gil & Ramachandran, 2001; Gil & Ratnakar, 2008). This capability is crucial for the approaches mentioned above to enable users to create and reuse workflows (Gil *et al.*, 2009), as well as to specify procedures through tutorial instruction (Groth & Gil, 2009a, 2009b; Gil *et al.*, 2011).

7 Provenance and trust

New knowledge is often obtained from open sources on the Web, whether they be Web pages, people, mashups, or services. These sources have varying quality and coverage, and may be

incomplete and contradict one another. The information is often of unknown origins and there is often no prior history with many of the sources that may be used to assess their reputation.

One focus of this research is how to determine whether to trust a particular piece of information provided by a source. Most existing trust systems are centered around the notion of entity trust, that is, whether or not to trust a given entity (Artz & Gil, 2007). We defined *content trust* as a finer-grained model than entity trust in that it assigns trust to pieces of information rather than the overall entity. We identified major factors that affect trust in information sources (Gil & Artz, 2007). We also developed a mathematical model and algorithms to derive source reputation from individual trust ratings (Gil & Ratnakar, 2002a; Gil & Artz, 2006).

Another aspect of this research is on tracking the provenance of processed information, which is synthesized by processing inputs from many individual sources. Representing the provenance of information is key to being able to track what sources were used and how they were combined (Gil & Ratnakar, 2002a). We investigated provenance in knowledge bases, where in order to be trusted each axiom must be tied to the source used to create it (Gil & Ratnakar, 2002b). Provenance of knowledge base axioms is even more critical when the knowledge base is to be developed and used by a community, as shown by our study of the most widely adopted biomedical ontology to date (Bada *et al.*, 2004). We also investigated provenance in workflow systems, and augmented existing provenance models that capture execution details with process and data semantics (Kim *et al.*, 2008). As described earlier, our semantic workflow representations can be used to reason about metadata properties of data sources as well as new data products as they are processed by the workflow steps, resulting in detailed provenance records that support semantic queries (Kim *et al.*, 2006).

Recent research aims at creating a better understanding of provenance requirements, and at fostering shared representations and standards for provenance of Web resources (Gil, 2010; Moreau *et al.*, 2010).

8 Conclusions

The influence of the Semantic Web on many research areas in artificial intelligence is palpable. This article discussed my views on how it has influenced the area of interactive knowledge capture and pointed to a number of current research topics in the field. Overall, the sheer overwhelming demand from all netizens of a more useful Web creates ample opportunities for increasing the challenges and potential for broader impact of our research in interactive knowledge capture.

Acknowledgments

I would like to thank all my collaborators over the years for their inspiration and contributions to the research described here. I am also thankful for the funding provided for this work by grants from a number of sources, including the Air Force Office of Scientific Research, the Air Force Research Laboratories, the Defense Advanced Projects Agency, the Intelligence Advanced Research Projects Activity, the National Institutes of Health, and the National Science Foundation. For this article, I gratefully acknowledge support from the US Defense Advanced Research Projects Agency under grant number HR0011-07-C-0060.

References

- Artz, D. & Gil, Y. 2007. A survey of trust in computer science and the Semantic Web. *Journal of Web Semantics* 5(2).
- Bada, M., Stevens, R., Goble, C., Gil, Y., Ashburner, M., Blake, J. A., Cherry, J. M., Harris, M. & Lewis, S. 2004. A short study on the success of the gene ontology. *Journal of Web Semantics* 1(2).
- Blythe, J., Kim, J., Ramachandran, S. & Gil, Y. 2001. An integrated environment for knowledge acquisition. In *Proceedings of the ACM International Conference on Intelligent User Interfaces (IUI-01)*, Santa Fe, New Mexico.
- Chalupsky, H., Gil, Y., Knoblock, C. A., Lerman, K., Oh, J., Pynadath, D. V., Russ, T. A. & Tambe, M. 2001. Electric elves: applying agent technology to support human organizations. In *Proceedings of*

- the *Thirteenth Annual Conference of Innovative Applications of Artificial Intelligence (IAAI-2001)*, Seattle, WA.
- Chklovski, T. & Gil, Y. 2005a. User interfaces with semi-formal representations: a study of designing argumentation structures. In *Proceedings of the ACM International Conference on Intelligent User Interfaces (IUI-2005)*, San Diego, CA.
- Chklovski, T. & Gil, Y. 2005b. An analysis of knowledge collected from volunteer contributors. In *Proceedings of the Twentieth National Conference on Artificial Intelligence (AAAI-05)*, Pittsburgh, PA.
- Chklovski, T. & Gil, Y. 2005c. Improving the design of intelligent acquisition interfaces for collecting world knowledge from web contributors. In *Proceedings of the Third International Conference on Knowledge Capture (KCAP-05)*, Banff, Alberta.
- Clark, P., Thompson, J., Barker, K., Porter, B., Chaudhi, V., Rodriguez, A., Thomere, J., Mishra, S., Gil, Y., Hayes, P. & Reichherzer, T. 2001. Knowledge entry as graphical assembly of components. In *Proceedings of the First International Conference on Knowledge Capture (K-CAP'01)*, Victoria, BC.
- Deelman, E., Blythe, J., Gil, Y., Kesselman, C., Mehta, G., Vahi, K., Blackburn, K., Lazzarini, A., Arbree, A., Cavanaugh, R. & Koranda, S. 2003. Mapping abstract workflows onto grid environments. *Journal of Grid Computing* **1**(1).
- Deelman, E., Singh, G., Su, M.-H., Blythe, J., Gil, Y., Kesselman, C., Mehta, G., Vahi, K., Berriman, G. B., Good, J., Laity, A., Jacob, J. C. & Katz, D. S. 2005. Pegasus: a framework for mapping complex scientific workflows onto distributed systems. *Scientific Programming* **13**(3).
- Ebiquity. 2006. *Big OWL Documents on the Semantic Web*. http://ebiquity.umbc.edu/blogger/?page_id=452
- Gil, Y. 2006. Workflow composition. In *Workflows for e-Science*, Gannon, D., Deelman, E., Shields, M. & Taylor, I. (eds). Springer Verlag.
- Gil, Y. 2009. From data to knowledge to discoveries: scientific workflows and artificial intelligence. *Scientific Programming* **17**(3).
- Gil, Y. (ed.). 2010. *Report of the W3C Provenance Incubator Group, World Wide Web Consortium*. <http://www.w3.org/2005/Incubator/prov/wiki/>
- Gil, Y. & Artz, D. 2006. Towards content trust of web resources. In *Proceedings of the Fifteenth International World Wide Web Conference, (WWW-06)*, Edinburgh, Scotland.
- Gil, Y. & Artz, D. 2007. Towards content trust of web resources. *Journal of Web Semantics* **5**(4).
- Gil, Y. & Kim, J. 2002. Interactive knowledge acquisition tools: a tutoring perspective. In *Proceedings of the 24th Annual Meeting of the Cognitive Science Society (CogSci)*, Fairfax, VA.
- Gil, Y. & Melz, E. 1996. Explicit representations of problem-solving strategies to support knowledge acquisition. In *Proceedings of the Thirteen National Conference on Artificial Intelligence (AAAI-96)*, Portland, OR.
- Gil, Y. & Ramachandran, S. 2001. PHOSPHORUS: a task-based agent matchmaker. In *Proceedings of the Fifth International Conference on Autonomous Agents (Agents'01)*, Montreal, Canada.
- Gil, Y. & Ratnakar, V. 2002a. Trusting information sources one citizen at a time. In *Proceedings of the First International Semantic Web Conference (ISWC)*, Sardinia, Italy.
- Gil, Y. & Ratnakar, V. 2002b. TRELIS: an interactive tool for capturing information analysis and decision making. In *Proceedings of the 13th International Conference on Knowledge Engineering and Knowledge Management (EKAW)*, Sigüenza, Spain.
- Gil, Y. & Ratnakar, V. 2008. Automating to-do lists for users: interpretation of to-dos for selecting and tasking agents. In *Proceedings of the Twenty-Third Conference of the Association for the Advancement of Artificial Intelligence (AAAI-08)*, Chicago, IL.
- Gil, Y. & Tallis, M. 1997. A script-based approach to modifying knowledge bases. In *Proceedings of the Fourteenth National Conference on Artificial Intelligence (AAAI-97)*, Providence, RI.
- Gil, Y., Ratnakar, V. & Fritz, C. 2011. TellMe: learning procedures from tutorial instruction. In *Proceedings of the ACM International Conference on Intelligent User Interfaces (IUI-11)*, Palo Alto, CA.
- Gil, Y., Ratnakar, V., Deelman, E., Mehta, G. & Kim, J. 2007a. Wings for Pegasus: creating large-scale scientific applications using semantic representations of computational workflows. In *Proceedings of the 19th Annual Conference on Innovative Applications of Artificial Intelligence (IAAI-07)*, Vancouver, British Columbia, Canada.
- Gil, Y., Deelman, E., Ellisman, M., Fahringer, T., Fox, G., Gannon, D., Goble, C., Livny, M., Moreau, L. & Myers, J. 2007b. Examining the challenges of scientific workflows. *IEEE Computer* **40**(12), 24–32.
- Gil, Y., Kim, J., Florez, G., Ratnakar, V. & Gonzalez Calero, P. A. 2009. Workflow matching using semantic metadata. In *Proceedings of the Fifth International Conference on Knowledge Capture (K-CAP)*, Redondo Beach, CA.
- Gil, Y., Kim, J. & Spraragen, M. 2010a. Principles for interactive acquisition and validation of workflows. *Journal of Experimental and Theoretical Artificial Intelligence* **22**(2).
- Gil, Y., Gonzalez-Calero, P. A., Kim, J., Moody, J. & Ratnakar, V. 2010b. Automatic generation of computational workflows from workflow templates using distributed data and component catalogs. *Journal of Experimental and Theoretical Artificial Intelligence* (In Press).

- Groth, P. & Gil, Y. 2009a. A scientific workflow construction command line. In *Proceedings of the ACM International Conference on Intelligent User Interfaces (IUI-09)*, Sanibel, FL.
- Groth, P. & Gil, Y. 2009b. Scaffolding instructions to learn procedures from users. In *AAAI Spring Symposium on Agents that Learn from Human Teachers*, Stanford, CA.
- Hall, M., Gil, Y. & Lucas, R. 2008. Self-configuring applications for heterogeneous systems: program composition and optimization using cognitive techniques. In *Proceedings of the Institute of Electrical and Electronics Engineers, Special Issue on Edge Computing*.
- Kim, J. & Gil, Y. 1999. Deriving expectations to guide knowledge base creation. In *Proceedings of the Fourteenth National Conference on Artificial Intelligence (AAAI-99)*, Orlando, FL.
- Kim, J. & Gil, Y. 2000. Acquiring problem-solving knowledge from end users: putting interdependency models to the test. In *Proceedings of the Fifteenth National Conference on Artificial Intelligence (AAAI-2000)*, Austin, TX.
- Kim, J. & Gil, Y. 2001. Knowledge analysis on process models. In *Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence (IJCAI-2001)*, Seattle, WA.
- Kim, J. & Gil, Y. 2002. Deriving acquisition principles from tutoring principles. In *Proceedings of the Intelligent Tutoring Systems Conference (ITS)*, Biarritz, France.
- Kim, J. & Gil, Y. 2003. Proactive acquisition from tutoring and learning principles. In *Proceedings of the 11th International Conference on Artificial Intelligence in Education (AI-ED)*, Sydney, Australia.
- Kim, J. & Gil, Y. 2007. Incorporating tutoring principles into interactive knowledge acquisition. *International Journal of Human-Computer Studies* **65**(10).
- Kim, J., Gil, Y. & Ratnakar, V. 2006. Semantic metadata generation for large scientific workflows. In *Proceedings of the Fifth International Semantic Web Conference (ISWC-06)*, Athens, GA.
- Kim, J., Myers, K., Gervasio, M. & Gil, Y. 2010. Goal-directed meta-control for integrated procedure learning. In *Metareasoning: Thinking about Thinking*, Cox, M. T. & Raja, A. (eds). MIT Press.
- Kim, J., Deelman, E., Gil, Y., Mehta, G. & Ratnakar, V. 2008. Provenance trails in the wings/pegasus workflow system. *Concurrency and Computation: Practice and Experience, Special Issue on the First Provenance Challenge* **20**(5).
- Kumar, V., Kurc, T., Ratnakar, V., Kim, J., Mehta, G., Vahi, K., Nelson, Y. L., Sadayappan, P., Deelman, E., Gil, Y., Hall, M. & Saltz, J. 2010. Parameterized specification, configuration, and execution of data-intensive scientific workflows. *Cluster Computing Journal* **13**.
- Kurc, T., Hastings, S., Kumar, V., Langella, S., Sharma, A., Pan, T., Oster, S., Ervin, D., Permar, J., Narayanan, S., Gil, Y., Deelman, E., Hall, M. & Saltz, J. 2009. High performance computing and grid computing for integrative biomedical research. *Journal of High Performance Computing Applications*.
- Maechling, P., Chalupsky, H., Dougherty, M., Deelman, E., Gil, Y., Gullapalli, S., Gupta, V., Kesselman, C., Kim, J., Mehta, G., Mendenhall, B., Russ, T., Singh, G., Spraragen, M., Staples, G. & Vahi, K. 2005. Simplifying construction of complex workflows for non-expert users of the southern california earthquake center community modelling environment. *ACM SIGMOD Record, Special Issue on Scientific Workflows* **34**(3).
- Marcus, S. & McDermott, J. P. 1989. SALT: a knowledge acquisition language for propose-and-revise systems. *Artificial Intelligence* **39**(1), 1–37.
- Moreau, L., Clifford, B., Freire, J., Futrelle, J., Gil, Y., Groth, P., Kwasnikowska, N., Miles, S., Missier, P., Myers, J., Plale, B., Simmhan, Y., Stephan, E. & Van den Bussche, J. 2010. The open provenance model core specification (v1.1). Future Generation Computer Systems, July 2010.
- Scaffidi, C., Shaw, M. & Myers, B. 2005. Estimating the numbers of end users and end user programmers. In *Proceedings of the 2005 IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC 2005)*, Dallas, TX.
- Tallis, M. & Gil, Y. 1999. Designing scripts to guide users in modifying knowledge-based systems. In *Proceedings of the Fourteenth National Conference on Artificial Intelligence (AAAI-99)*, Orlando, FL.
- Tallis, M., Kim, J. & Gil, Y. 2001. User studies of knowledge acquisition tools: methodology and lessons learned. *Journal of Experimental and Theoretical AI* **13**, 359–378.
- Vrandečić, D., Ratnakar, V., Krötzsch, M. & Gil, Y. 2010. *Shortipedia: Aggregating and Curating Semantic Web Data*. Third Prize in the Semantic Web Challenge, 2010. <http://challenge.semanticweb.org/>
- Woollard, D., Medvidovic, N., Gil, Y. & Mattmann, C. 2008. Scientific software as workflows: from discovery to distribution. *IEEE Software, Special Issue on Developing Scientific Software*.