

Statement of research

Semi-automatically and cooperatively built knowledge bases

Dr Philippe MARTIN

Abstract. The general goal of my research during the last ten years, and one direction of research I'd like to take in the medium term, is to find out and support *scalable* ways of extracting, entering, sharing and managing precise, normalized and well interconnected information. I have so far mainly focused on (i) the design of [intuitive, expressive and normalizing notations](#) (and automatically generated forms) for knowledge entering, querying, navigation and combination, (ii) the [use of such notations within Web documents to index any of its elements](#) (or, conversely, document the represented knowledge), (iii) the creation of a [server permitting the users to update a same very large knowledge base \(KB\) without lexical/semantic conflicts nor forcing them to agree with each other but encouraging or forcing them to keep the KB structured](#), (iv) the creation of a [multi-source ontology](#) (built by transforming WordNet into a genuine lexical ontology and then joining various top-level ontologies to it) to ease knowledge entering and sharing, and (v) various [proof-of-concept domain modellings](#) (including three courses and one application in tourism). I'd like to apply and extend this research in the context of applications and make a very substantial progress on the following tasks: (i) automatic and semi-automatic knowledge-based document element indexation and retrieval, (ii) e-learning, ubiquitous learning and semantic learning grids, (iii) building a cooperatively-built state of the art in knowledge engineering, and (iv) complementing Wikipedia for technical or original materials and hotly debated issues. Given my approach, these researches are related. With respect to my research so far, one new aspect but natural continuation will be the use and refinement of natural language understanding techniques.

Table of Contents

1. [Problems, approaches, observations and hypothesis](#)
2. [Accomplishments](#)
3. [Research goals](#)
4. [References](#)

1. Problem, approaches, observations and hypothesis

The problem. Nowadays, a person sharing information often has to repeat or summarise ideas or facts that she or other persons have written in other documents. Furthermore, publishing a document only adds to the volume of poorly structured and heavily redundant data that has to be searched for finding information.

Insufficiency and complementarity of current approaches. There are many different approaches that partially address the above problem. Document retrieval systems and question-answering systems exploit legacy data and are easy to use but the document retrieval systems can only provide a list of "possibly interesting and relevant" documents while question-answering systems currently only retrieve sentences from documents. Ideally, the semantic content of documents would be automatically interpreted and merged into a well structured semantic network, thus permitting to (i) answer queries with information which is as precise and complete as possible, (ii) display tables or hierarchies (e.g., specialization hierarchies and argumentation structures) to support information comparison, and (iii) filter information according to users' preferences/models (e.g., a user might not want to see anything written by a member of certain organisations). Given the state of the art in natural language understanding, knowledge representation and reasoning, and general knowledge bases, the above cited ideal is far away and few research projects claim to work on achieving it directly ([CYCORP](#) is a well-known exception). Instead, information providers are asked to take small steps toward knowledge representation: they are asked to store small non-contextual units of information into shared repositories (e.g., (Downes, 2001)), and/or provide metadata. However, outside very limited domains, these steps are insufficient to support an efficient retrieval and comparison of information. Indeed, (i) most current units of information (e.g., most learning objects and most pages in Wikipedia) are still informal documents, (ii) most current metadata schemes, database schemas and independently developed ontologies are still very simple and cannot be automatically merged in a precise or logically sound way since they are poorly inter-related, follow different conventions and lack precision (Euzenat & al., 2005), (iii) knowledge representation (KR) languages used for these ontologies (e.g., RDF+OWL or KIF) are too low level to permit genuine knowledge sharing (additional lexical, ontological and structural conventions are needed), and (iv) most of the user-friendly KR notations (e.g., controlled languages, topic maps, or notations in hypertext systems or argumentations systems) are either too informal or too restrictive to guide the entering of knowledge that can be re-used for inference purposes and hence for automated knowledge retrieval and comparison (Marshall & Shipman, 2003).

Additional observations underlying my approach. Even if the semantic content of documents was automatically interpreted and merged, people would still need to learn simple KR notations to visualise, navigate and refine parts of the semantic network. Knowledge representations, their sources and their documentation should be well connected or, whenever possible, merged (Fensel, 1998); this was well acknowledged by the knowledge acquisition community (Fensel, 1998) but is only fully addressed in CGKAT (Martin, 1995) and, as much as HTML permits, in WebKB-1 (Martin & Eklund, 1999). My numerous experiments with CGKAT and WebKB-1 consistently showed that even for a single user with the possibility to re-use a large lexical ontology without performance penalties, using only documents to store knowledge (and hence loading these documents into the tools when performing queries) is not a scalable avenue: in many cases, a large scale knowledge base server (KBS) is required to facilitate, cross-check and guide knowledge modelling and retrieval. The use of a large shared KBS is even more crucial for knowledge sharing since otherwise an information provider cannot simply add one concept or statement "at the right place" and cannot be guided in her knowledge modelling by a KBS server with a unique large ontology: she has to find relevant ontologies on the Web, understand them, choose between them and combine them to create a new ontology, which is difficult, time-consuming and adds to the number of ontologies that other people or inference engines will later have to choose from. A normalized way to represent and insert a statement into a knowledge base (KB) is a requirement for keeping it organised and hence for its scalability. This implies the absence of redundancies and a unique place to add a precise category or statement (or a set of equivalent places). For simple modellings, this can be achieved (Dromey, 2006). In any case, forcing knowledge providers to relate a new object (category or formal/informal statement) to its already stored generalizations and specializations lead or encourage these persons to keep the KB organised (e.g., it encourages them not to add redundant objects even when the redundancies cannot be detected automatically). Euzenat (2005) recommends the exploitation and update of a shared KBS by ontology aligning agents. Only two KBSs seem to have special protocols to support cooperation between people: Co4 (Euzenat, 1996) and WebKB-2 (Martin & Eklund, 2001). Co4 is not used anymore. Its approach was based on peer reviewing, the result of which was a hierarchy of KBs, the uppermost ones containing the most consensual knowledge and the lowermost ones being the private KBs of the contributing users. This approach was intuitive but generated separate KBs just for storing how consensual each piece of information was.

My research approach. As opposed to Co4 or document based approaches, the approach of WebKB-2 consists of enabling to associate any kind of information about any object with this object and then, if a user provides filtering criteria, filter the information presented during navigation or in answer to a query. (For example, a user may not want to see statements from people or kinds of people that have authored or backed-up ideas or arguments that were un-interesting or invalid from the viewpoint of this user). Hence, separate *physical modules* (KBs or documents) are not needed for storage, they can be generated, including those used as sources. Each object (category, relation or statement) must have a recorded creator and/or source module. Thus, a KB can be cooperatively updated by many users with editing protocols that lead information providers to explicit relationships between inconsistent or partially redundant statements via specialization relations (e.g., "example") or "corrective relations" (e.g., "corrective_restriction", as in the following Formalized-English statement: `any bird is agent of a flight'(John) has for corrective_restriction `most healthy French birds are able to be agent of a flight' (Joe)). The statements do not have to be formal and can be arbitrarily large or small (in order to allow for incremental refinements) but have to be connected by relations. WebKB-2 does not currently detect inconsistencies or partial redundancies between completely informal statements but in the future it could use heuristics to perform such a detection and hence suggest (instead of enforce) the use of corrective relations.

This shared KB approach is more efficient and scalable than the approach of Co4 since (i) it maintains a minimal organisation of the knowledge, (ii) similar peer-reviewing protocols can be used but only one KB is to be managed by these protocols, and (iii) it permits the use of much more complex knowledge evaluation algorithms than predefined peer-reviewing protocols - see (Martin, 2005) for a template algorithm based on votes and argumentation relations on statements; in the future, users will be allowed to personalise this template to evaluate or filter knowledge. Other advantages of a large shared KBS over private KBs or documents have been noted or hinted at in the previous paragraph (for example, a feature of my approach is to use the content of the KB to cross-check and guide knowledge entering, for example via automatically generated forms).

Furthermore, the WebKB-2 approach ("gathering and storing any information about any of the objects of interest of the KB") also means that each WebKB-2 server should (ideally) periodically check *related servers* (more general servers, competing servers and slightly more specialized servers), import the information relevant to its domain and, for the other information, store pointers to those servers. Via knowledge replication and, if necessary, query redirections, it would not matter much which server a Web user attempts to query or update first. Within a peer-to-peer network such replications mechanisms can be supported or enforced and it is then not detrimental for each user to have her own WebKB-2 server. Automatically or semi-automatically importing knowledge from large or similar KBs may be difficult but is much simpler than importing knowledge from semi-independently developed small KBs since they provide less information to exploit.

Hypothesis underlying my research approach. One rather safe hypothesis is that different, internally consistent, ontologies do not have to be modified to be integrated into a unique consistent semantic network. Some precision (e.g., contexts or assumptions) that were implicit in the source ontologies may have to be made explicit in order to solve semantic conflicts and the relations between categories may have to become indirect rather than direct. However, all these relations can still hold, and hence the source ontologies can be regenerated and newer versions of these ontologies can be rather easily integrated. A stronger hypothesis is that the categories of these different ontologies can always be strongly

inter-related into a well-organised semantic network, especially via specialization relations of one sort or another, in order to genuinely permit knowledge sharing and re-use. For example, categories from a "4D ontology" (where a temporal aspect is associated to each referred object) cannot be said to be "subtype" of their counterparts in a "3D" ontology (where a temporal aspect does not have to be associated to each object but can be later specified using temporal contexts within statements) but a "3Dto4Dsubtype" relation type can be defined and used to strongly interconnect the two ontologies and allow logical inferences similar to those that can be made using "subtype" relations. An hypothesis that generalises the two previously cited ones is that semantic conflicts can always be solved by adding more precision or making explicit how they boil down to mere "preferences", and hence that solving conflicts increases the organisation of the KB. A related hypothesis is that solving conflicts (adding precision) can be done incrementally (when a new piece of information is added) and that people do not have to meet nor compromise to merge their ontologies.

From a sociological viewpoint, a related hypothesis is that a sufficient number of persons will take the time to be precise and learn the notations or conventions to do that. This will not be a problem once the approach becomes popular since (i) this seems less difficult than learning musical notations, programming languages or XML, (ii) the social success of Wikipedia and many open-source projects has shown that many persons are willing to contribute, (iii) the approach would solve many problems of Wikipedia or other repositories (e.g., their need for a committee to decide what should or should not be kept in the repository and hence for certain users to make content selections for other users; most wikis also give extensive deletion rights to any user). Until this approach becomes popular, semi-automatic knowledge extraction and very simple textual or graphic notations should be proposed. To conclude, the strongest hypothesis is that, for each domain a KB is dedicated to, the contributors to the KB can be sufficiently guided to keep the KB well organised; this implies that the initial knowledge of the KB, the provided protocols, notations, knowledge entering forms, semi-automatic knowledge extraction procedures, and normalisation procedures constitute a sufficient guide. Providing a good basis for all these guiding elements is difficult and central to my research. However, the social success of Wikipedia shows that even a very weakly organised repository can be seen by many persons as interesting.

2. Accomplishments

During my PhD, I designed [CGKAT](#) ([Martin, 1995](#); [Martin & Alpay, 1996](#)), a tool helping knowledge engineers to represent, link and search information within KBs and documents. To achieve this, I integrated the structured document editor Thot, the Conceptual Graph workbench Cogito, models of the KADS library, my extension of WordNet 1.5 with some top-level ontologies, and the Unix shell. Because of its re-use of Thot, CGKAT remains the only tool that fully integrates and combines KB management with document management.

During my postdoc in 1997 and then as a research fellow at Griffith University, I designed [WebKB-1](#) ([Martin & Eklund, 1999, 2000](#)), a KB server enabling knowledge engineers to load or execute knowledge representations/queries stored in Web documents and enabling them index any part of any Web document (furthermore, as in CGKAT, associating a query to an hyperlink permits the generation of a virtual document). By the end of 1997, WebKB-1 began to propose user-friendly notations for representing or indexing parts of Web documents, and a language of commands to retrieve or combine knowledge in a conceptual, structural or lexical way. Nowadays, many tools - generally exploiting an XML-based language - have similar purposes but the notations and commands they propose are most often very cumbersome, restricted in scripting and knowledge representation/querying capabilities, and often cannot be used within informal textual documents. In Ontobroker (Fensel, 1998) and nowadays [Semantic MediaWiki](#), some relations about the "object that a document is about" can be hidden within hyperlinks in this document; this approach is user-friendly for casual readers of the pages but quite un-friendly and restrictive for knowledge providers since only simple relations about the object of the page can be represented (hence, for example, the semantic content of a table cannot be represented within one document). To sum up, WebKB-1 integrates KB management with document management as much as current Web browsers permit.

From July 2000 to December 2003, as a senior research fellow at DSTC, I designed [WebKB-2](#) ([Martin & Eklund, 2001](#); [Martin, 2003a](#)), a KB server that can not only manage a very large KB but also allows people or software agents to store and tightly interconnect their knowledge into it *without* having to discuss and agree on terminology or beliefs. To that end, I built a Web-accessible multi-source KB management system above an object-oriented database and designed special [editing protocols to encourage knowledge interlinking and keep the KB consistent](#) (as far as the inference engine can tell). To initialise the default KB (the one "proposed" to people), I [transformed WordNet 1.7 into a genuine lexical ontology, corrected it, and extended it with several top-level ontologies](#) ([Martin, 2003b](#)). The result of this merge (which, unlike most other merging efforts, modifies the source ontologies only if some inconsistencies within them are detected) and the on-going integration of the DOLCE and SUMO top-level ontologies, has been named the [MSO](#) (Multi Source Ontology). It has been [voted "a material to work on" by the Standard Upper Ontology Working Group](#), and can be accessed and extended by any Web user or software agent via WebKB-2.

For WebKB-1 and WebKB-2, I designed and continue to refine [notations that are more intuitive, expressive and normalizing](#) that currently existing notations ([Martin, 2002](#)). One of them is adapted to the case of simple relations between categories and permits to represent a large volume of knowledge in a structured way and a small amount of space (which is important for browsing a large KB). Two others, named Frame-CG and Formalized English, are derived from the Conceptual Graph Linear Form and not only extend it but improve on the qualities that made its success: its intuitiveness

and "knowledge normalization" effect. That is, people are better led to follow [good "lexical, structural and ontological principles"](#) that I collected and refined (Martin, 2000). Thus, these notations ease the visualization, handling, entering and sharing of knowledge. WebKB-2 can also export some of its knowledge in some other notations such as RDF+OWL, CGIF and KIF. For WebKB-2, I also designed "cascading knowledge entering forms": such forms are automatically generated from definitions and general statements in the KB and they can be combined to guide knowledge entering on any object once its type has been selected. For an accommodation broker (namely, Wotif), I [represented some information about accommodations on the Sunshine Coast \(Australia\) and completed the accommodation retrieval predefined interfaces that I designed \(e.g., using Google maps\) with an access to the above mentioned generated forms to support unforeseen kinds of queries.](#)

Since February 2004, I extend WebKB-2 to permit it to exploit some external inference engines and, most importantly, support collaboratively-built *well-organized semi-formal* repositories. For example, I designed a user-friendly notation providing all the necessary constructs to engage in or represent "structured discussions" (i.e., semantically organized collections of arguments and counter-arguments for various statements) and designed an [algorithm to evaluate the popularity and originality of each contribution and contributor](#) based on the structure of the discussions and votes on each statement (Martin & al., 2006). Such repositories could potentially be used for precision-oriented corporate memories and as enhanced versions or complements to the current on-line yellow-pages and auction sites. However, I focused on showing how they could help researchers, lecturers and students by supporting them in cooperatively building a semantically structured "state of the art" or a "learning object repository" in their domains. Although these two kinds of repositories are generally distinguished, the same semantic network can, and ideally should, be used for both research and learning/teaching. I organised a [workshop about such repositories and the use of structured discussions to build them.](#)

As a starting point for a semantically structured "state of the art" in knowledge engineering and to permit the comparison of knowledge management tools, I have begun an ontology of such tools, even though an important part of this ontology is currently focused on Conceptual Graph tools (Martin & Eboueya, 2007). As part of my on-line teaching of Workflow Management and, the last semester of 2006, as part of my ["Griffith E-Learning" research grant project](#), I represented a good part of the lecture materials for three different courses into semantic networks within WebKB-2, asked the students to contribute to parts of these semantic networks (e.g., as a replacement for an informal "learning journal" and hence as a better way to train and evaluate their critical thinking) and asked them to fill surveys about the approach. Many of them recognized the help that the centralisation and categorisation of pieces of information scattered all across the lecture materials provided them in accessing and understanding these information. In (Martin & al., 2006, 2008) I analysed the various interests and applications of my work for e-learning and the design/sharing/ retrieval of learning objects.

3. Research goals

The general goal of my research during the last ten years, and the research I'd like to do in the medium term, is to find out and support *scalable* ways of extracting, entering, sharing and managing precise, normalized and well interconnected knowledge representations. This implies pursuing my work on ontologies, notations, normalization, protocols, representation of various domains, and experiments in collaborative KB building.

This also implies starting to re-use and refine semantic based methods for natural language understanding (the most ancient and well-known being Roger Schank's conceptual dependency theory). Indeed, scalability requires the use of automatic (or at least semi-automatic) knowledge extraction techniques, and this is a natural extension of my works on the normalised representation of expressive knowledge and on ontology merging. For example, WebKB-2 should be able to give a formal representation of the simple sentence "governments should enforce animal rights", insert the category "enforcement_of_animal_rights_by_governments" into the process hierarchy of WebKB-2, and associate the normalized sentence "the enforcement_of_animal_rights_by_governments should happen" to this category. Although many sentences will not fully be automatically understandable, normalizing and indexing many of their components (and especially their main processes) will go a long way toward permitting the classification and comparison of ideas. Such an NLU will require the integration of ontologies such as OntoWordNet, Extended WordNet and FrameNet into WebKB-2; both the process and the result of this integration will be of interest to many research communities besides the NLU one.

Knowledge extraction via NLU permits to index small document elements (DEs), not just whole documents, and hence permits question answering rather than document retrieval. Actually, retrieving DEs via the querying and navigation of an organised KB offers more possibilities than current question answering approaches (direct retrieval of sentences in documents). For example, no document retrieval or current question answering system can answer questions such as "What are the characteristics of the various theories and implemented parsers related to [Functional Dependency Grammar](#) and how do these theories and parsers respectively compare to each other?", nor do they allow people to correct or complement the given information. On the other hand, Wang & al. (2003) present techniques automatically indexing documents via an existing domain ontology. Such techniques could take advantage of (simple parts of) the ontologies that can be cooperatively built via WebKB-2 servers. Conversely, connections from categories to documents would increase the use of these ontologies and be starting points for manual or semi-automatic extensions of these ontologies or more precise indexations/representations of some documents. Hence, a relatively easy but fruitful extension of WebKB-2 is to add a module re-using some of these techniques. Regarding knowledge extraction via NLU, I'll essentially focus on

extracting the most important information for knowledge sharing and NLU itself: specializationOf and partOf relations between categories, especially process categories (other relations will mainly come from the integration of existing ontologies).

Here are more details below about three projects on which I'd like to make a very substantial progress within three years.

- **Cooperatively-built state of the art in knowledge engineering.** There is a need for methods to compare tools and techniques in knowledge engineering. An example of this need is the interest that [Michael Denny's "Ontology editor survey"](#) attracted despite the fact that its very coarse level of details made it quite unusable for its main goal: tool comparison. In the ["CG tools" page on Wikipedia](#) I categorised seven well-known tools from the Conceptual Graph (CG) community according to 160 criteria organised by specialization relations and grouped into 6 sections and tables. I also [began to complete the KB of WebKB-2 for representing this categorisation as well as other information related to research in CGs, and more generally, knowledge engineering](#); some preliminary results have been published in ([Martin & Eboueya, 2007](#)). I wish to continue this work at least up to a point where the content of the existing KB is (i) sufficiently large and interesting for researchers in knowledge engineering to be motivated to complete (at least to "advertise" their tools), and (ii) a sufficient guide to enable *scalable additions* (of tools, features, objections, references to articles or people, etc.). In other words, it will then be easy for the researchers to know the "right places" where they should insert new categories or statements for the KB to remain well organised. The generation of the kinds of tables described in ([Martin, 2005](#)) to ease the comparison of objects (products, techniques, etc.) will also have to be refined and implemented. This goal is ambitious but achievable with three years and will be an important and well recognised research landmark.
- **Complement to Wikipedia for technical or original materials and hotly debated issues.** The goals and approach of Wikipedia restrict it to be poorly structured, of a general "encyclopaedic nature" and "consensual" (when edit wars end satisfactorily). These restrictions are not and will not be lessened by the existence of sister-projects (Wiktionary, Wikiquote, Wikibooks, etc.), the [proposals for many more sister-projects](#), nor the potential use of [Semantic MediaWiki](#) in the future. The centralised and structured approach of WebKB-2 (along with replication mechanisms between KBs) permits to go beyond these restrictions. Hence, WebKB-2 could be used as a complement for Wikipedia. As a first step, using fine-grained argumentation structures, I'll represent the debates of many Wikipedia pages marked with having a debated content and, at the bottom these pages, will invite people to explore and complement the structures in WebKB-2. (I have begun the representation of some of these debates).
- **Application of my collaboration-support techniques to semantic grids and learning grids.** Semantic grids combine Grid Computing techniques and Semantic Web techniques, two promising but not-yet-mature kinds of techniques for sharing data, knowledge, services and CPU resources. These techniques will provide the most suited platforms for collaboration or e-learning, at least within restricted communities. However, retrieving, comparing, entering and sharing/re-using precise and well-structured knowledge on these grids remains as difficult as on the Web or current Semantic Web. In Section 1, I showed how my research approach reduces these problems in the context of a large shared KB server but could also be applied to a peer-to-peer network and hence to a semantic grid. I'll work on this application/adaptation in collaboration with the above mentioned PhD student and her main supervisor, A.Pr. Michel Eboueya. Currently, the users allowed to access the KB of a WebKB-2 server can access any information in it; as part of my future research on knowledge sharing protocols, I'll allow a user to specify rules to hide some of her information until she receives some other information about the reader. This feature would be useful for many on-line services such as match-making services, employment services or auction sites.

Finally, WebKB-2 is planned to be re-used by two Australian research teams: (i) the [RMIT](#) (Australia) as part of an [Agent-based Web Service Composition project for NineSigma Pty Ltd](#), and (ii) the [Software Quality Institute](#) (Griffith Uni) for sharing, querying, discussing and documenting software specifications.

4. References

- Downes S. (2001). *Learning Objects: Resources For Distance Education Worldwide*. International Review of Research in Open and Distance Learning, Vol. 2, No.1, Oct. 1st 2001.
- Dromey G. R. (2006). *Scaleable Formalization of Imperfect Knowledge*. Proceedings of AWCVS-2006, Asian Working Conference on Verified Software, 29-31 October 2006, Macao SAR, China. www.iist.unu.edu/www/workshop/AWCVS2006/
- Euzenat J. (1996). *Corporate memory through cooperative creation of knowledge bases and hyper-documents*. Proceedings of 10th KAW, (36)1-18, Banff, Canada, Nov. 1996.
- Euzenat J., Stuckenschmidt H. & Yatskevich M. (2005). *Introduction to the Ontology Alignment Evaluation 2005* Proceedings of K-Cap 2005 (pp. 61-71), workshop on Integrating ontology, Banff, Canada, 2005.

- Euzenat J. (2005). *Alignment infrastructure for ontology mediation and other applications*. Proceedings of 1st ICSOC international workshop on Mediation in semantic web services, Amsterdam (NL), pp. 81-95, 2005
- Fensel D., Decker S., Erdmann M. & Studer M. (1998). *Ontobroker: Or How to Enable Intelligent Access to the WWW*. Proceedings of KAW'98 (11th Knowledge Acquisition Workshop), pp. 8-23, Banff, Canada, 1998.
- Marshall C.C. & Shipman F.M. (2003). *Which Semantic Web?*. Proceedings of ACM Hypertext 2003, pp. 57-66.
- Martin P. (1995). *Knowledge Acquisition Using Documents, Conceptual Graphs and a Semantically Structured Dictionary*. Proceedings of KAW 1995 (9th International Knowledge Acquisition for Knowledge-Based Systems Workshop), pp. 1-19, Banff, Canada, February 26 - March 2, 1995.
- Martin P. & Alpay L. (1996). *Conceptual Structures and Structured Documents*. Proceedings of ICCS 1996, 4th Int. Conf. on Conceptual Structures (Springer, LNAI 1115, pp. 145-159), Sydney, Australia, August 19-22, 1996.
- Martin P. & Eklund P. (1999). *Embedding Knowledge in Web Documents*. Proceedings of WWW8 (pp. 324-341, 8th Int. World Wide Web Conf., Toronto, Canada, May 11-14, 1999) and special issue of "Computer Networks, Int. Journal of Computer and Telecommunications Networking", Vol. 31 (11-16), pp. 1403-1419, February 1999.
- Martin P. (2000). *Conventions and Notations for Knowledge Representation and Retrieval*. Proceedings of ICCS 2000, 8th Int. Conf. on Conceptual Structures (Springer, LNAI 1867, pp. 41-54), Darmstadt, Germany, August 14-18, 2000.
- Martin P. & Eklund P. (2000). *Knowledge Indexation and Retrieval and the Word Wide Web*. *IEEE Intelligent Systems*, special issue "Knowledge Management and Knowledge Distribution over the Internet", pp. 18-25, May/June 2000.
- Martin P. & Eklund P. (2001). *Large-scale cooperatively-built heterogeneous KBs*. Proceedings of ICCS 2001, 9th Int. Conf. on Conceptual Structures (Springer, LNAI 2120, pp. 231-244), California, USA, July 30 to August 3, 2001.
- Martin P. (2002). *Knowledge representation in CGLF, CGIF, KIF, Frame-CG and Formalized-English*. Proceedings of ICCS 2002, 10th Int. Conf. on Conceptual Structures (Springer, LNAI 2393, pp. 77-91), Bulgaria, July 15-19, 2002.
- Martin P. (2003a). *Knowledge Representation, Sharing and Retrieval on the Web*. Chapter 12 of a book titled "Web Intelligence" (Springer; editors: N. Zhong, J. Liu, Y. Yao; pp. 263-297; ISBN 3-540-44384-3; Web Intelligence Consortium's book), January 2003.
- Martin P. (2003b). *Correction and Extension of WordNet 1.7*. Proceedings of ICCS 2003, 11th International Conference on Conceptual Structures (Springer, LNAI 2746, pp. 160-173), Dresden, Germany, July 21-25, 2003.
- Martin P., Blumenstein M. & Deer P. (2005). *Toward cooperatively-built knowledge repositories*. Proceedings of ICCS 2005, 13th Int. Conf. on Conceptual Structures (Springer, LNAI 3596, pp. 411-424), Germany, July 18-22, 2005.
- Martin P., Eboueya M., Blumenstein M. & Deer P. (2006). *A Network of Semantically Structured Wikipedia to Bind Information*. Proceedings of E-learn 2006 (pp. 1684-1702), *AACE Conference* on E-learning in Corporate, Government, Healthcare, & Higher Education, Honolulu, Hawaii, October 13-17, 2006.
- Martin Ph. & Eboueya M. (2007). *Sharing and Comparing Information about Knowledge Engineering*. Published in the *WSEAS Transactions on Information Science and Applications*, Issue 5, Volume 4 (pp. 1089-1096; ISSN: 1790-0832), May 2007.
- Martin Ph. & Eboueya M. (2008). *For the ultimate accessibility and re-usability*. Book chapter to be published early 2008 in the *Handbook of Research on Learning Design and Learning Objects: Issues, Applications and Technologies*.
- Trombert-Pavot B., Rodrigues J.M., Rogers J.E., Baud R., van der Haring E., Rassinoux A.M., Abrial V., Clavel L. & Idir H. (2000). *GALEN: a third generation terminology tool to support a multipurpose national coding system for surgical procedures*. *International Journal of Medical Informatics*, Vol. 58-59 (pp. 71-85), September 1st 2000.
- Wang B. B., McKay R. I., Abbass H. A. & Barlow M. (2003). *A comparative study for domain ontology guided feature extraction*. In Proceedings of ACSC-2003, 26th Australian Computer Science Conference, pages 69-78. Australian Computer Society, 2003.