MEDINFO 2004
M. Fieschi et al. (Eds)
Amsterdam: IOS Press


# Knowledge Management in Pediatric Pain: Mapping On-Line Expert Discussions to Medical Literature

**Syed Sibte Raza Abidi[a], Allen Finley[b], Evangelos Milios[a], Michael Shepherd[a], David Zitner[b]**

[a]*Health Informatics Lab, Faculty of Computer Science, Dalhousie University, Halifax, Canada*
[b]*Health Informatics Lab ,Faculty of Medicine, Dalhousie University, Halifax, Canada*

**Abstract**

*Clinical decision-making can be vastly improved with the availability of the right medical knowledge at the right time. This concept paper presents a knowledge management research program to (a) identify, capture and organize the tacit knowledge inherent within on-line problem-solving discussions between pediatric pain practitioners; (b) establish linkages between topic-specific pediatric pain discussions and corresponding published medical literature on children's pain available at PubMed—i.e. linking tacit expert knowledge to explicit medical literature; and (c) make these knowledge resources available to pediatric pain practitioners via the WWW for timely access to various modalities of clinical knowledge.*

***Keywords***

Knowledge Management, Pediatric Pain, Tacit Knowledge, Discussion Lists, Medical Literature

## Introduction

Children's pain is a significant burden to society that causes suffering not only for the child but also for the family, and interferes with the child's schooling and social development. In the past 15 years there has been a wider recognition of the under-treatment of children's pain and the simultaneous need for more knowledge in the area of pediatric pain. Yet, the reality is that the vast majority of children's pain is managed by family physicians, general pediatricians, or nurses, who have generally had little or no training in the area, or by parents. To complicate matters further, many pediatric pain problems are relatively rare and not well represented within the published medical literature. Hence, physicians, nurses, and other professionals in everyday practice often have to rely on other practitioners' clinical experience to make sound decisions pertaining to pediatric pain care.

Clinical decision-making can be vastly improved with the availability of the right medical knowledge at the right time. The Institute of Medicine reports that errors increase when clinicians do not have timely access to clinical evidence [1]; in fact, only about half of clinical treatments are supported by high-quality evidence [2]. Having said that, it can be argued that healthcare practitioners do not necessarily have timely access to various modalities of clinical knowledge—i.e. published medical literature, clinical experiences of experts and peers, structured abstracts and remote consultations with domain experts. Hence, there is an imminent need to provide access for healthcare professionals to medical knowledge resources so that they become effective users of best evidence throughout their careers.

This concept paper presents a unique Knowledge Management (KM) approach to the timely provision of clinical knowledge to healthcare practitioners. Our approach is to supplement best evidence vis-à-vis published medical literature with related tacit knowledge of healthcare experts. Tacit knowledge—i.e., knowledge that reflects practical clinical experiences—will be derived from archived online discussions (collected via an email-based discussion list) between a community of healthcare practitioners, who share their clinical experiences and research findings to address a particular clinical problem/case. We present a KM research program to (a) identify, capture and organize the tacit knowledge inherent within on-line problem-solving discussions between pediatric pain practitioners; (b) establish linkages between topic-specific pediatric pain discussions and corresponding published medical literature on research and clinical management of children's pain available at PubMed—i.e. linking tacit expert knowledge to explicit medical literature; and (c) make these knowledge resources available to pediatric pain practitioners via the WWW for timely access to various modalities of clinical knowledge.

## Pediatric Pain Knowledge Sources

The *Pediatric Pain Discussion List* (PPDL) is a preeminent international source of clinical discussion on pediatric pain, initiated by Finley in 1993 to foster informal email-based communication between pediatric pain professionals [3]. At present there are over 700 subscribers including clinicians, researchers, academics and editors of major current texts in pediatric pain. Currently, the PPDL archive contains more than 10,000 messages covering various pediatric pain topics. The message content includes clinical problems, research problems or proposals, successful experiences, suggestions and administrative aspects of children's pain management and prevention. The message content is subjected to a critical, if informal, "peer-review" by over 700 experts.

In practice, PPDL represents a unique approach to knowledge creation and translation that both supports and supplements the published literature by providing rapid expert consultation and access to a wider community of pediatric pain practitioners. This form of communication reflects and promotes the articulation of

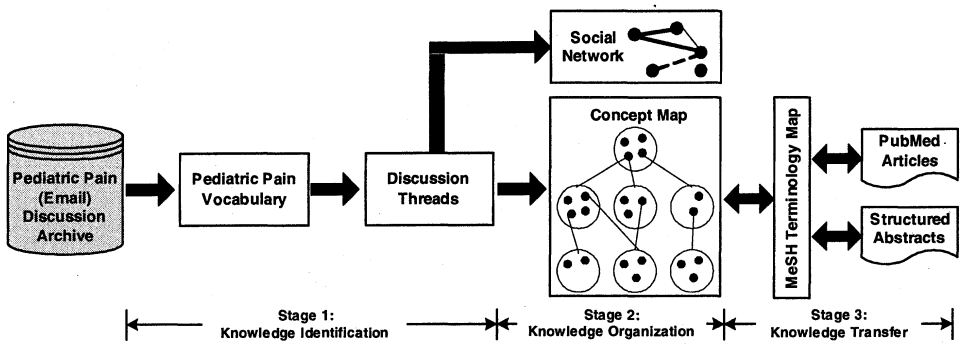


*Figure 1 - Schematic view of the KM methodology for linking tacit knowledge with explicit knowledge*

problem-specific tacit knowledge and clinical experience of pediatric pain specialists, which might not otherwise be shared or made explicit via published medical literature. The workflow of PPDL is as follows: (1) a practitioner initiates discussion on a specific pediatric pain issue; (2) multiple pediatric pain practitioners discuss the issue—a debate involving experts ensues during which they share experiences, relate theory to practice and collaboratively suggest a solution; (3) expert's tacit knowledge explicated during the discussion is disseminated to various pediatric pain practitioners. We believe that such discussions capture the accumulated experience and clinical judgment of many expert contributors—akin to a compendium of "wisdom"—that needs to be captured, organized and utilized to impact the delivery of healthcare [4].

Medical knowledge can be differentiated along the lines of *explicit* and *tacit knowledge*. Explicit knowledge is about *how things should work* and can be described as canonical knowledge presented as clinical guidelines, studies and methods as reported in published medical literature. Tacit knowledge is the innate knowledge of medical experts; it embodies their experiential know-how, skills, and intuitive judgment—*what really works* and *how to make it work*. Even though not "evidence-based" by the traditional definition, this tacit knowledge is necessary, and is recognized as valid, valuable, and an essential part of professional practice [5]. Based on the above classification of medical knowledge, at one end of the spectrum we have PPDL as representative of the cumulative tacit knowledge of a vibrant *community of practice*—i.e. a group of knowledge workers informally sharing their experience, expertise and commitment to a common topic/area—dealing with pediatric-pain management. At the other end of the spectrum there is published medical literature reporting clinical research findings that is representative of explicit knowledge. Indeed, medical decision making can be significantly improved if medical practitioners are able to leverage both tacit knowledge of medical experts (if available) and published explicit knowledge. Yet, at present there are no mechanisms whereby these two vital modalities of medical knowledge can co-exist, inter-relate with each other and be concurrently presented to medical practitioners as all-encompassing best evidence to impact clinical decision-making.

### The Case for Mapping On-Line Expert Discussions to Medical Literature

It is our contention that the potential linking of tacit knowledge (i.e. PPDL) to explicit knowledge (i.e. published literature) will have a significant impact on healthcare practitioners whilst dealing with atypical and ill-represented/researched medical conditions. The importance of such a linkage is further amplified by the fact that clinicians need to provide advice in many situations where strong evidence is lacking. Even when studies have been completed, clinicians may be unaware of the pertinent research. In the absence of clinical algorithms based on scientific evidence, clinicians then need to rely on intuitive approaches that, although omnipresent in practice, are difficult to describe and teach. It is in these situations that tacit knowledge of healthcare experts—reflecting what solution will work, why it will work, and how to make it work—can be useful to practitioners if made available as a clinical decision support resource. Vice versa, clinicians at times need to verify or seek the opinion of their peers with respect to the veracity, practicality and impact (under different operational conditions) of research findings presented in medical literature. In this situation they will be able to refer to the clinical experiences of the respective community of practice.

## Pediatric Pain Knowledge Management Methodology

The broad objective of the KM research program is to translate the PPDL into a pediatric pain knowledge resource—i.e. to *identify*, *extract*, *organize* and *share* the tacit knowledge of pediatric pain practitioners inherent within PPDL discussions. Our methodology involves three KM stages (see figure 1).

### Stage 1: Knowledge Identification

This stage involves: (a) the extraction of a *controlled pediatric pain vocabulary*, including emerging pediatric pain terms, identified within the pediatric pain discussions; and (b) the identification of problem-mediated and topic-specific *discussion threads* comprising a continuous sequence of individual email messages.

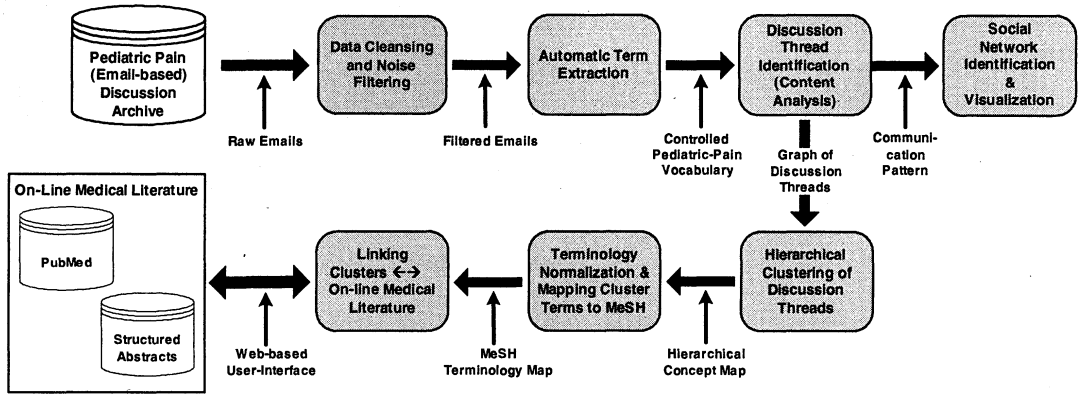


*Figure 2 - Schematic view of the research plan, highlighting the various research modules*

**Stage 2: Knowledge Organization**

This stage involves (a) the formation of a higher-level abstraction of the discussion threads in terms of a hierarchical *concept map* that represents clusters of contextually and semantically similar discussion threads. Users will be able to search the concept map via a graphical web-based user interface, such that clicking on a pediatric pain topic on the concept map will retrieve the most pertinent discussion threads related to that topic; and (b) the identification of communities of practice—in terms of *social networks*—from the interaction patterns between PPDL contributors during various discussions. The emergent *social network* will depict the mutual interactions between practitioners.

**Stage 3: Knowledge Transfer**

This final stage involves establishing linkages between the pediatric pain concept map (i.e. tacit expert knowledge) with online medical literature (i.e. explicit published knowledge) at PubMed. This will allow healthcare practitioners to simultaneously access related clinical knowledge from both these knowledge resources. In practice, salient pediatric-pain terms within a topic-specific discussion thread will be used as the search criteria for automated queries to retrieve relevant published literature from PubMed. Working in reverse, users accessing medical literature at PubMed will be able to link to discussion threads (within the concept map) related to the contents of the retrieved literature.

Since medical literature at PubMed is indexed based on MeSH (Medical Subject Headings), there is a need to design a MeSH terminology map to transform the controlled pediatric pain terminology to MeSH compliant terms.

Finally, a web-based user-interface will be designed for medical practitioners to access the various pediatric pain knowledge resources—i.e. PPDL, concept map, social network and online medical literature via the concept map.

## Knowledge Management Methods: An Overview

The proposed KM research program raises interesting questions about how to extract usable tacit medical knowledge from informal online discussions; how to link the extracted, and subsequently organized, knowledge to published best evidence; how the availability of both tacit and explicit knowledge impact the delivery of healthcare; and how this affects the sources of knowledge and their usage by healthcare practitioners. The forthcoming discussion presents a research plan (illustrated in figure 2) that stipulates a sequence of research modules that suggest methods to address the abovementioned KM issues. At present, these modules are under development.

**Module 1: Cleaning of the PPDL archive**

Email messages do not follow an *a priori* structure, vocabulary, discourse, and context. There is a need to develop computational methods to pre-process the PPDL archive in order to remove potential noise (i.e. announcements, personal communications or system messages) using sophisticated content analysis techniques that leverage machine learning algorithms. Our approach is to look beyond the subject line of the email—i.e. to use both the body text and the subject—to determine its relevance to pediatric pain discussions. We plan to use a feature vector approach, where a feature vector will describe each email along a number of dimensions, such as message length, density of non-medical terms, density of medical terms, existence of particular words and frequency of medical terms. An inductive analysis of the feature vector may determine the appropriateness of each email message [6].

**Module 2: Automatic term extraction from PPDL archive**

In order to organize the PPDL into coherent discussion threads it is important to have a controlled vocabulary that will help identify the content-based linkages between various messages. We propose to automatically identify and extract a controlled pediatric-pain vocabulary—comprising both atomic and composite terms—from the PPDL discussions. The advantages of working with a data-driven controlled vocabulary as opposed to a standard vocabulary such as MeSH are identified as follows: (a) we anticipate the presence of composite concepts—concepts created from two or more atomic concepts—that are mandated by the pediatric pain community, and may be unique to pediatric pain discussions only; (b) the presence of specific terms can be easily traced to the context for meaning elicitation; (c) statistical information can be maintained by keeping the terms in their

original form; and (d) we can avoid the normalization of existing terms.

Our approach is to use NLP techniques to extract specialized terms from technical documents [6]. For instance, the use of 'light-weight' unification-based grammars can lead to the identification of salient noun-phrase terms, which are good candidates for terms in a controlled vocabulary. The choice of unification-based grammars is appropriate because they combine syntactic and semantic features in describing sentence structures, and in our case semantic information is necessary for identifying content-bearing phrases, i.e., terms.

**Module 3: Discussion thread identification**

In our approach, the granularity of a conceptual unit is not an individual email message, but a discussion thread comprising multiple email messages. Message threading involves the grouping of individual E-mail messages into a topic-specific discussion thread that maintains the temporal order in which the messages are created. We believe that message threading will not only group various contextually-similar messages into a coherent discussion thread, but will additionally lead to the reduction of information overload presented by a collection of heterogeneous messages. We argue that identification of discussion threads should be based on the following criteria: (a) subject heading of the message; (b) message content; (c) time-stamps of the message; (d) message discourse, such as reply to an earlier email; and (e) author information associated with the message.

Our message threading approach presupposes the PPDL archive as a *graph* of discussion threads, where nodes are individual messages and directed links exist between consecutive messages on the same thread. The problem here is to associate a message to the most relevant discussion thread based on the abovementioned criteria. We suggest the use of network-based knowledge representation formalisms for determining and representing discussion threads. We believe that graph analysis and partitioning techniques will be highly appropriate for this purpose as this will not only depict the relational schema between the messages but also allow for dynamic growth as new discussion threads originate. The graph representation will make explicit the threading connections between messages, in preparation for the subsequent content-based clustering of the discussion threads.

An interesting problem that will also be addressed is *topic detection and tracking*—i.e., to identify the spawning of a new discussion thread off existing discussion threads. We anticipate that long threads may change focus, and if so then they need to be segmented into sub-thread units, where segmentation will be based on the discussion content and the discontinuities in the terms used in the earlier part of a long thread.

**Module 4: Organization of discussion threads as a hierarchical concept map**

We propose that the organization of the emergent discussion threads is best represented in terms of a hierarchical structure—i.e. a *concept map*—that depicts the conceptual relationships between the various discussion threads [7]. At the highest-level the concept map will represent clusters of contextually-similar discussion threads. Functionally, users will be able to traverse the entire pediatric pain archive by selecting a topic and then following the links to sub-topics, whilst reading the discussion threads associated with the topic.

Our approach is to inductively derive a concept map from the totality of discussion thread and sub-thread units based on the full textual content of the constituent messages, using the controlled vocabulary derived in module 2. The machine learning literature comprises several methods for hierarchical clustering, including classical clustering algorithms, Principal Direction Divisive Partitioning, and Self Organizing Maps. We argue that the choice of the clustering method should be based on *a priori* defined criteria, including computational efficiency considering the size of the PPDL archive, quality (cohesion), visualization (in terms of distinct boundaries), population size of the clusters and the levels of the hierarchical organization of the clusters in terms of sub-topics.

**Module 5: Organization of discussion threads as a social network**

We propose to generate a social network of PPDL contributors based on their mutual interactions during a discussion—the act of replying to a message is viewed as an embodiment of social interactions [8]. We believe that a social network can assist healthcare practitioners, following a particular discussion, to identify the following: central contributors, expertise profiling and collaboration over a period, topic and institutions. We believe that simple graph algorithms are best suited to create a dynamically evolving social network and probabilistic methods can be used to extract sub-groups that exhibit higher levels of collaboration. Healthcare practitioners should be able to interact with social networks, such that clicking on an edge (representing interaction between two contributors) will retrieve the discussion thread and clicking on a node (representing a contributor) will list all messages posted by a contributor.

**Module 6: Mapping controlled pediatric pain vocabulary to MeSH**

Literature available at PubMed is indexed based on MeSH vocabulary standards. For efficient literature, we argue that there is a need to map our controlled pediatric pain vocabulary to MeSH terms [9]. Our approach is to establish automatic conceptual equivalence, at a terminological level, between our controlled vocabulary and MeSH terms [10]. We propose to develop a *MeSH terminology map* that applies terminology normalization techniques. We anticipate that individual terms in our controlled vocabulary may appear in different places in MeSH which might disorient the mapping exercise. Hence, we suggest the use of a set of terms in a particular message to establish the localization of the matching term in the MeSH hierarchical structure.

**Module 7: Linking topic-specific clusters in the concept map to literature at PubMed**

PubMed features multiple mechanisms to retrieve articles given a search query. We intend to leverage PubMed's set of E-Utilities and write additional PERL scripts to pass automated queries generated from the discussion thread contents to the PubMed search engine. We plan to initially retrieve the abstracts of the

relevant articles from PubMed, with the provision of full-article retrieval.

**Module 8: Web-based user interface design**

To ensure that healthcare practitioners find it convenient to access and interact with the knowledge resources created, a web-based user-interface need to be designed. Users will be able to access the PPDL, search for a discussion on a topic, browse the concept map, view any discussion thread, locate experts on the social network and link to PubMed for published literature from the concept map (and vice versa).

## Concluding Remarks and Future Outlook

From a KM perspective, there is a premium in capturing the "nuggets" of experiential knowledge articulated during online discussions and translating it to a medical knowledge resource. In this paper, we have presented the conceptual design of the proposed KM program that is under development. This unique KM program leverages a vibrant social network of pediatric pain experts to create tacit knowledge in new and interesting ways. Knowledge translation is anticipated as follows: Healthcare professionals can (a) refer to peer moderated discussions on various important topics about pediatric-pain management; (b) refer to summaries of past discussions for just-in-time clinical evidence; (c) refer to the published medical literature related to a discussion topic and vice versa; (d) use the controlled vocabulary of pediatric pain-terms; and (e) relate to a network of pediatric pain experts. The proposed KM program provides a unique tacit-explicit knowledge gateway for clinicians to support evidence-based clinical decision-making in pediatric pain, a critical evolution in clinical practice [11].

## References


[1] Institute of Medicine. Crossing the quality chasm: A new health system for the 21st century. National Academy Press: Washington, DC, 2001.

[2] Sim I, Gorman P, Greenes RA, Haynes RB, Kaplan B, Lehmann H, Tang PC. Clinical decision support systems for the practice of evidence-based medicine. *JAMIA* 2001; 8: 527-534.

[3] Finley GA, McGrath PJ. The PEDIATRIC-PAIN mailing list – An electronic interdisciplinary team. 3rd International Symposium on Pediatric Pain, Philadelphia, June 7, 1994.

[4] Abidi SSR. Knowledge management in healthcare: towards "knowledge-driven" decision-support services, *Int J Med Informatics* 2001; 63: 5-18.

[5] Patel BL, Arocha JF, Kaufman DR. Expertise and tacit knowledge in medicine. In Sternberg RJ & Horvath JA (Eds.) *Tacit Knowledge in Professional Practice—Researcher and Practitioner Perspectives.* Lawrence Erlbaum Associates, Inc.: Mahwah, NJ, 1999.

[6] Milios E, Zhang Y, He B, Dong L. Automatic Term Extraction and Document Similarity in Special Text Corpora', *Proc. 6th Conf. of the Pacific Assoc. for Computational Linguistics,* Halifax, Aug. 22-25, 2003.

[7] Chen H, Houston AL, Sewell RR, Schatz BR. Internet browsing and searching: User evaluations of category maps and concept space techniques. *J Am Soc Information Sci* 1998; 49(7): 582-603.

[8] Wasserman S, Faust, K. *Social Network Analysis.* Cambridge, MA: Cambridge University Press, 1994

[9] Lowe HJ, Bartnett GO. Understanding and using MeSH vocabulary to perform literature searches. *JAMA* 1994; 271(14): 1103-8.

[10] Abidi SSR, Manickam S. Leveraging XML-based electronic medical records to extract experiential clinical knowledge: An automated approach to generate clinical cases for medical case-based reasoning systems. *Int J Med Informatics 2002*; 68: 187-203.

[11] Jadad AR, Haynes RB, Hunt D, Browman GP. The Internet and evidence-based decision-making: A needed synergy for efficient knowledge management in health care. *Journal Canadian Medical Assoc.*, 2000; 162: 362-5.



**Address for correspondence**

Dr. Syed Sibte Raza Abidi
Health Informatics Lab
Faculty of Computer Science, Dalhousie University
Halifax, NS, B3H 1W5, Canada.
Email: sraza@cs.dal.ca