

A Dialogue-based System for Sharing Development Knowledge

From: AAAI Technical Report WS-00-03. Compilation copyright © 2000, AAAI (www.aaai.org). All rights reserved.

Massoud Moussavi

The World Bank*

1818 H Street N.W.

Washington, D.C. 20433

mmoussavi@worldbank.org

Abstract

Lessons-learned systems as a technology to capture and disseminate an organization's know-how have recently been receiving considerable attention. What this means in the case of the World Bank is to capture and transfer the lessons of best and worst practices gained through economic and social development projects in developing countries. But Knowledge that increases the developing countries capability to reduce poverty cannot be passively transferred to them. Rather, they must actively acquire and absorb it within the context of their prior experience and problems. This paper describes the design and implementation of a *case-based lessons-learned* system that creates a dialogue-based environment in which policy makers and experts construct knowledge for themselves by examining cases and examples that are relevant to their situation and have direct impact on their projects and activities.

Keywords - Knowledge Sharing, Case-Based Reasoning, Lessons-Learned, Development Institutions.

Introduction

Sharing knowledge is an essential ingredient of the World Bank's Development Strategy. The reasoning behind this strategy is the fact that poor countries--and poor people--differ from rich ones not only because they have less capital but because they have less knowledge (WDR 1998). Thus, the problems of development are viewed from the perspective of knowledge rather than finance and material aid. A major component of Bank's knowledge sharing strategy is to create an on-line knowledge base that captures, stores, and transfers the lessons of best and worst practices in economic and social development projects.

But experience shows that knowledge that increases the developing countries capability to reduce poverty cannot be passively transferred to them. Rather, they must actively acquire and absorb it within the context of their prior experience and problems. We have developed a system that builds upon such a context. Using an initial set of cases taken from a recent World Development Report (WDR 1998) and in consultation with domain experts at

the World Bank, we have created a case-based lessons-learned system.

This paper describes the design and implementation of this system that was built upon our initial proposal (Moussavi 1999). A dialogue-based environment is created in which policy makers and experts construct knowledge for themselves by examining cases and examples that are relevant to their situation and have direct impact on their projects and activities. The system would help users find answers to questions such as: "How did India become the third largest producer of milk in the World?" Or "What kinds of problems does government regulation create?"

Development Assistance via Knowledge Sharing

From a fairly closed organization a few years ago, the World Bank has become a global development partner making it easier for people to find out who knows what and where the best expertise can be drawn upon, wherever it resides.

The main idea in development assistance via knowledge sharing is to put knowledge on par with money and material aid. But this is easier said than done. Knowledge does not "travel" as well as money; there are different types of knowledge: codified vs. tacit and general vs. local, and the form of knowledge involved will affect the roles of international and local knowledge institutions (Ellerman 1999). As discussed in (Hansen *et al.* 1999) there are two very different knowledge management strategies: *codification* where knowledge is codified and stored in databases and *personalization* where knowledge is shared through direct person-to-person contacts. In this paper, we deal, of course, only with knowledge that can be automated.

A fundamental question is whether development knowledge, codified or not, can be transferred. Can the

* The findings, interpretations, and conclusions expressed in this paper are those of the author and should not be attributed in any manner to the World Bank, to its affiliated organizations, to the Board of Directors or the countries they represent.

best practices of a developing country be replicated in another? A "best practice" at a company such as the Ford Motor Company is general enough that it can be replicated with big savings for the company (Wolford 1999). But when it comes to economic development it is very difficult to know *a priori* just how general a best practice is (Ellerman 1999). In fact, contrary to popular wisdom, there are times when it pays to reinvent the wheel! (Cole 1989) and often the local adaptation of a "best practice" amounts to reinventing the "best practice" in the new context. Given that there are different and sometimes conflicting viewpoints concerning development issues such as privatization and fighting corruption, Ellerman argues that development assistance needs to be based on the active learning philosophy where learners having an active role in acquiring knowledge. That is, a development institution such as the World Bank should act as an open learning institution where different theories are allowed to publicly contend and collide, and where experimentation is encouraged to see what is locally appropriate.

Following the philosophy of the active learning approach which is the transformation of the student into an active constructor and appropriator of knowledge, we are developing a system that would allow the users form knowledge for themselves by examining cases that are relevant to their situation. To develop such a system, two questions should be addressed: (1) How can *development knowledge*, which has to do with the content of the economic and social development projects (one's experiences and stories) be represented? And (2) What is the best mechanism to share and transfer such knowledge?

Concerning the first question, knowledge about economic and social development is intrinsically case-based. There is no single best method to approach development issues. One deals with a new situation by remembering, modifying and applying prior experience. That is, one uses case-based reasoning. Since CBR techniques provide a mechanism for representation and retrieval of experiential information (Becerra-Fernandez & Aha 1999), a case-based framework where lessons-learned and best practices are stored in terms of cases seems to be a viable method to represent development knowledge.

As for the transfer of development knowledge, the role of dialogue must be emphasized. In an active learning environment the cases and questions that are presented would be relevant to users' interests and they can actively construct and appropriate knowledge through a dialogue with the system.

General Vs. Local Knowledge: An Example

Information concerning best practices and lessons learned are usually available in two distinct forms:

- general guidelines, lessons and recommendations.

- local solutions and experiences

Examples of the general knowledge are statements such as the following: "to insure quality of products establish brand names" or "development should rely on private sector investment." But general knowledge may not suffice to determine local solutions as the following case study clearly demonstrates. It turns out that what contributed to the success of India's Dairy Production was farmer-owned cooperatives and not reliance on the private corporate sector.

India's Dairy Production (WDR 1998):

"In India in the 1950s, milk production could not keep pace with growing demand. Some milk vendors watered down the milk. Because there were many vendors, and brand names were not clearly established, vendors who did not dilute their milk could not command a premium and were squeezed out of the market. The result was an overall drop in milk quality. Enter the National Dairy Development Board, which in the early 1970s launched Operation Flood, a multifaceted program to improve the functioning of the milk market by ensuring quality. The board began by encouraging the creation of dairy cooperatives and helping them establish quality standards..."

We represent the general knowledge in terms of a set of rules in a knowledge base and the specific local solutions in terms of cases in a case library. Based on the design outline described in (Moussavi 1999), we have developed a tool in which lessons learned and case studies documented in the World Development Report (WDR 1998) and other World Bank publications are stored in terms of cases. Adaptation of a case is left to the user; the system only helps to retrieve a case that can be adapted and when possible it suggests a "best practice."

Question-based Retrieval: An Example

The retrieval of cases is closely related to the indexing method used (Leake 1996). We have opted for a hybrid of feature-based and question-based indexing. That is, each case is partially indexed with a certain number of features as well as indexed with a number of questions that it either answers or raises (Ferguson *et al.*, 1992). Feature-based indexing which identifies and stores attribute-value pairs, is most suitable when the system is used as an advisory and prescriptive tool. That is, when the user knows the problem and is looking for a solution. For example, if a World Bank task manager is working on a health project say in rural area in Nigeria he/she might want to look at how a similar problem was addressed before.

When the system is used as an educational rather than problem solving tool, the question-based indexing seems most appropriate. There are two major advantages of a

question-based indexing: it generates interest in the topic and it engages users in a dialogue. Both are crucial in the process of active learning and knowledge discovery.

To index a case, using the model of ASK systems (Ferguson *et al.* 1992; Schank *et al.* 1991), we analyze the content of each case, for the questions that it answers and also for the questions that it raises. Some of such questions concerning the example on India's Dairy Production that we discussed earlier would be as follows:

- "Why did the state intervene in the case of milk production in India?"
- "How do brand names provide quality assurance?"
- "What kinds of problems does government regulation create?"

Figure 1 shows an interaction with a user. The user is first prompted to choose a particular country and topic of interest. (In this example, we have only shown country and topic as the only indexed features. However, the case-base is indexed with more features.) Once the user selects a country and a topic, a list of relevant questions are displayed.

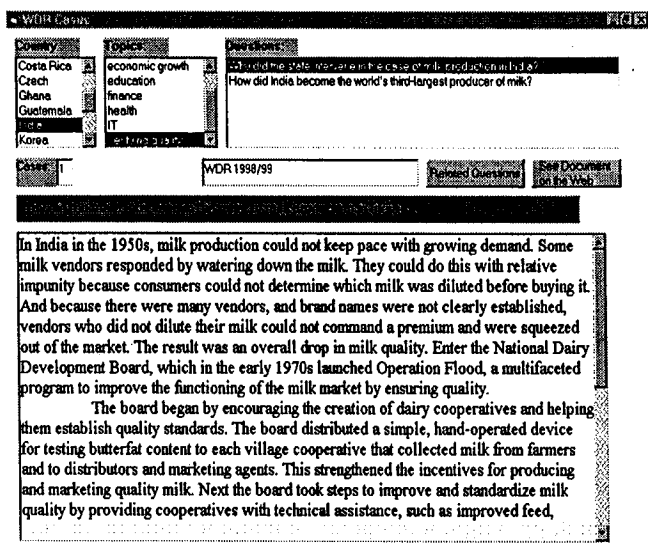


Figure 1: The User Interface

The interface allows the user to browse the case-base by selecting one or more questions. Once a case is displayed, the user can browse related cases by clicking on the "Related Questions" button and then selecting one of the related questions. This cycle continues until the user is satisfied. Adaptation of cases in this system is left up to the user as adapting development knowledge to local conditions and culture requires considerable effort (Ellerman, 1999).

Knowledge Representation

We have taken a minimal representation approach (Edelson 1993) where the primary research is concerned with representing as little as possible about the domain and yet achieve the educational goal of the system. That is a case-based knowledge-sharing tool need not know much about what it teaches. Rather it should be able to provide users with the information that they need.

To create a database of cases, we model in terms of objects and relationships between them. The objects include *cases*, *questions* and *rules*. Two broad classes of cases are identified as *theoretical cases* that explain economic and development theories and principles and *specific cases* which describe lessons learned from specific projects. Questions are also broken into two classes *general questions* (e.g., "Why regulate?") and *specific questions* (e.g., "How did Ghana implement a successful telecom reform?") As discussed earlier, *rules* are statements that prescribe a particular action or contain a recommendation. As for relationships among objects, the most important relationship is the ones between cases and questions. A case is linked to questions that it raises or answers.

Each case is represented in terms of a number of indexed features or attributes (geographical region, country, sector, topic, etc.) and a text field that contains the description of the situation, the solution and the outcome. Since there is a natural abstraction hierarchy between some of the attributes (e.g., region and country), the objects are further categorized into subclasses.

We have represented 100 cases and 70 questions taken from The World Development Report (WDR 1998) and other Bank publications in a relational database. All the cases have been indexed with a number of attributes such as region, country, sector, sub-sector, etc. Currently, about 50 cases are linked to questions. The cases are fairly small with textual parts not exceeding 500 words. Framing questions was fairly straightforward as a number of the case studies in these publications posed a major question as part of their titles.

Discussion

We have built an interactive case-based system using a "conversational CBR" framework (Aha and Breslow 1997) that allows users to navigate through a library of case studies to find an example that is similar to their situation. This is achieved by prompting users to answer questions on indexed features. Additionally, we have created cross-links between a number of cases by indexing those cases with questions that they raise or answer. The users will be able to browse the case library by selecting a question that interests them, retrieve a case and browse other related cases

by selecting one or more questions. In effect, an active and participatory learning environment is created.

The tool was developed using Microsoft Visual Basic to build the user interface and the CBR engine. We have also created a Web enabled interface to the database. The advantage of a Web interface is that it would allow for easy integration of other sources of information concerning the country or the topic.

Building a case-based system to represent development knowledge is a big challenge because it implies indexing a large number of documents. One of the main issues is how to scale up. We are currently exploring document summarization techniques (McKeown 1999; Mani and Bloedorn 1999) so that a larger number of documents can be indexed more economically. At the same time we are investigating the applications of semi-automated indexing techniques (Bruninghaus and Ashley 1999) in the area of development information. We will also research the applicability of adaptation techniques to retrieval of relevant cases. Adaptation-guided retrieval is a novel technique that assesses the adaptation requirements of cases during retrieval (Smyth and Keane, 1995).

Summary

The paper illustrates our approach in constructing a knowledge-sharing tool to capture and represent economic and social development knowledge. The system is based on a conversational model that anticipates users' questions concerning development issues. Using this tool, policy makers and experts form knowledge by examining cases and examples that are relevant to their situation and have a direct impact on their projects and activities. The system provides the following benefits:

- A unified framework to represent the lessons learned and best practices concerning economic and social development.
- Active and participatory learning occurs because the examples and questions are couched in terms of what makes sense from the user's viewpoint and interests.
- Knowledge is not transferred to users through a sequence of lectures. Rather, it is acquired by them at their own speed through a series of structured questions.
- Provides a better user interface so that the users do not need to have extensive prior knowledge about the domain in order to find what they want in the knowledge base.
- A deeper level of knowledge is obtained because of the cross-links between similar cases. A system that links related experiences would have significantly higher operational value.

Acknowledgements

I would like to thank the AAAI reviewers for their helpful comments and feedback.

References

- Aha, D. W. and Breslow, L.A. (1997). Refining Conversational Case Libraries. In *Proceedings of the Second International Conference on Case-Based reasoning*, pages 267-278. Providence, RI: Springer-Verlag.
- Becerra-Fernandez, I., and Aha, D.W. (1999). Case-based problem solving for knowledge management systems. In *Proceedings of the Twelfth Annual Florida Artificial Intelligence Research Symposium*. Menlo Park: AAAI.
- Bruninghaus, S. and Ashley, K.D. (1999). Bootstrapping Case Base Development with annotated Case Summaries. In *Proceedings of the Third International conference on Case-Based reasoning (ICCBR-99)*. Lecture Notes in Artificial Intelligence 1650. Althoff, K.-D., Bergmann, R., and Branting, L.K. (ed.) pp. 59-73. Springer:Berlin.
- Cole, R. (1989). *Strategies for Learning*. Berkley: University of California Press.
- Edelson, D.C. (1993) Learning from Stories: Indexing and Reminding in a Socratic Case-Based Teaching System for Elementary School Biology. Dissertation Thesis. Northwestern University.
- Ellerman, D. (1999). Global Institutions: Transforming International Development Agencies into Learning Organizations. *The Academy of Management: Executive*, V. 13 (1, Feb. 99), pages 25-35.
- Ferguson, W., Bareiss, R., Birnbaum, L., and Osgood, R. (1992). ASK Systems: An Approach to the Realization of Story-based Teachers. Evanston: The Institute for the Learning Sciences.
- Hansen, M.T., Nohria, N., and Tirenay, T. (1999). What's Your Strategy For Managing Knowledge. *Harvard Business Review*, March-April 1999, pages 106-116.
- Leake, D. (1996). CBR in Context: The Present and Future. In D. Leake, editor, *Case-Based Reasoning: Experiences, Lessons, & Future Directions*, pages 3-30. Menlo Park, CA: AAAI Press.
- Mani, I., Bloedorn, E. (1999) Summarizing Similarities and Differences Among Related Documents. *Information Retrieval* 1(1): 35-67.
- McKeown, K. (1999) Toward Multi-document summarization by Reformulation: progress. In

Proceedings of the National Conference on Artificial Intelligence (AAAI-99). Menlo Park: AAAI.

Moussavi, M. (1999). A Case-Based Approach to Knowledge Management. In Aha, D.W., Becerra-Fernandez, I., Maurer, F., & Muñoz-Avila, H. (Eds.) *Exploring Synergies of Knowledge Management and Case-Based Reasoning: Papers from the AAAI Workshop* (Technical Report WS-99-10). Menlo Park, CA: AAAI Press.

Schank, R., Ferguson, W., Birnbaum, L., Barger, J., and Greising, M. (1991). ASK-TOM: An experimental interface for video case libraries. In *Proceedings of the Thirteenth Conference of the Cognitive Science society*. Northvale, NJ: Erlbaum.

Smyth, B., and Keane, M.T. (1995). Experiments on adaptation-guided retrieval in case-based design. In *Proceedings of the First International Conference on case-Based Reasoning*, pages 313-324. Sesimbra, Portugal: Springer-Verlag.

WDR (1998). *Knowledge for Development: The World Bank World Development Report*. The World Bank.

Wolford, D. (1999). *Best Practice Replication at FORD*. Knowledge Management Review. September/October 1999.