SHARING-IN-ACTION[1]

How designers can exchange insights without knowing

Ann Heylighen, Mathias Casaer, Herman Neuckermans

Katholieke Universiteit Leuven,
Dept. of Architecture, Urban Design & Planning, Kasteelpark Arenberg 1, B-3001 Leuven, Belgium
ann.heylighen/mathias.casaer/herman.neuckermans@asro.kuleuven.ac.be

ABSTRACT

In architecture, design ideas are developed as much through interaction as by individuals in isolation. This awareness inspired the development of a Dynamic Architectural Memory On-line, an interactive platform to share ideas, knowledge and insights in the form of concrete building projects among designers in different contexts and at different levels of expertise. Interaction with various user groups revealed this platform to suffer from at least two thresholds. First of all, making projects available to other platform users takes time, effort, and specific skills. Secondly, designers tend to sense a psychological threshold to share their ideas and insights with others. In trying to tackle both thresholds, this paper proposes to conceive the platform as an associative network of projects, and develops ideas about how the relationships in this network can be determined and updated by exploiting the insights implicitly available in the project documentation and user (inter)actions. In the long run, this should allow the platform to learn from *all* designers using the platform, including those who do not release information on their own projects, and to apply the lessons learned to continuously enhance its performance.

KEYWORDS

architectural design, connectionism, self-organization

1. INTRODUCTION

Recently, the awareness has been growing that in architecture, like in other design domains, design ideas are developed as much through interaction as by individuals in isolation. Motivated by this observation, we developed a Dynamic Architectural Memory On-line (DYNAMO), an interactive multimedia platform to stimulate and support the sharing of ideas, knowledge and insights in the form of concrete building projects among designers in different contexts and at different levels of expertise.

A longitudinal study of DYNAMO's interaction with various user groups during the past six years revealed two obstacles that prevent the platform from being used as originally intended. First of all, making projects available to other platform users takes time, effort, and specific skills. Secondly, designers tend to sense a psychological threshold to share their ideas and insights with others.

In trying to tackle both obstacles, we propose to conceive the platform as an associative network of projects, and to determine the relations between them by exploiting meaningful information that is implicitly embedded in their documentation and consultation. Using a combination of statistical analysis and specific database queries, patterns and subtle relationships between the different projects, indices and other contents can be identified. In the long run, this should allow the platform to learn from users' interactions with the project collection, and to apply the lessons learned to continuously enhance its performance.

In the sections which follow, the basic concepts of DYNAMO are presented, followed by their implementation and results from experiments with various user groups. Based on this foundation, new features are identified to better fulfill DYNAMO's original ambition by using patterns of usage. Subsequently, a concrete approach is advanced for implementing these new features effectively. The paper closes with a summary and plans for future research.

---

[1] After Donald Schön's concept of reflection-in-action (1983).

2. DYNAMO IN A NUTSHELL

2.1 Idea

DYNAMO is a web-based design assistant for students and professional designers in the field of architecture (Heylighen & Neuckermans 2000). DYNAMO aims to incorporate quite literally the view of cognition underlying Case-Based Design (CBD) and at the same time extrapolate it beyond the individual.

Firmly rooted in the Theory of Dynamic Memory (Schank 1982), the CBD approach propounds that people's knowledge does not only consist of abstract, generally applicable principles, but also of specific experiences, so-called cases (Riesbeck & Schank 1989, Kolodner 1993). Moreover, it claims that human memory is dynamically changing with every new experience. Several years of observing and analyzing people's remindings have nurtured the hypothesis that experiencing, understanding, remembering and learning cannot be separated from one another. Our understanding grows by trying to integrate new things with what we already know. As a result, understanding causes us to come across old experiences as we process new ones. A significant side effect of this process of understanding is that memory never behaves exactly the same way twice, since it changes as a result of its own experiences. As experiences are recalled and used, memory gets an opportunity to try out the knowledge associated with them. This allows memory to re-organize and re-define itself dynamically, in other words to learn from its experiences (Kolodner 1993).

Learning from experience can occur in different ways (Riesbeck & Schank 1989). New episodes are stored in terms of old expectations generated by previous experiences. Eventually expectations that used to work may have to be invalidated. Indices to unique experiences that were once useful will cease to do so because similar experiences have been encountered. In short, memory learns from experience by acquiring new cases, grouping similar cases, or re-indexing cases stored improperly at first.

Inspired by the cognitive model underlying CBD, DY*N*AMO is conceived as an (inter)active workhouse rather than a passive warehouse. It actively develops and is interactively developed by architects' design knowledge, by stimulating and supporting several modes of interaction (see Figure 1):

1. interaction among cases, for cases are labeled and linked to related cases by various features architects address during design;
2. interaction between (human) designer and (computer) memory, for users cannot only consult cases in DYNAMO, they can improve its content in various ways;
3. interaction among individual designers in different contexts and at different levels of expertise, for DYNAMO is meant for collective use by students in architecture schools and professionals in design firms;
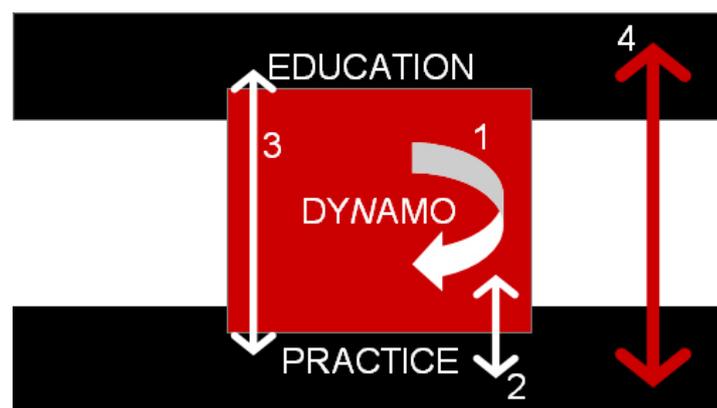4. and thus also interaction between practice and education in architectural design.

Figure 1. DYNAMO stimulates and supports several modes of interaction.

2.2 Implementation

From a technical point of view, DYNAMO can be thought of as a learning content management system. The platform is designed to support the creation, storage, use and reuse of learning content in the granular form of building projects (see http://dynamo.asro.kuleuven.ac.be). All learning content is organized by means of a dynamic metadata classification system and stored in a data repository embodied by a relational database, subdivided into four sub-databases as follows (see Figure 2):

The *cases* database labels every building project by various features: project name, architect and location, but also aspects of form and space, function, construction and context. These metadata serve as filter criteria during retrieval and as links to projects with analogous characteristics. In order to avoid confusion, a clear distinction should be made between *categories* and *values*. The term category refers to the name of an index, e.g. spatial configuration. Each category provides a place for a project to characterize itself with one or more values, e.g. cluster, linear, radial, 'plan libre'. A value thus refers to the concrete realization of a category for a certain project. It characterizes a specific project but can, and in many cases will be the same for several projects. In other words, whereas categories are chosen generally and shared by all projects, the values for these categories are assigned to each case specifically. For some categories, materials for instance, a single project can have multiple values.
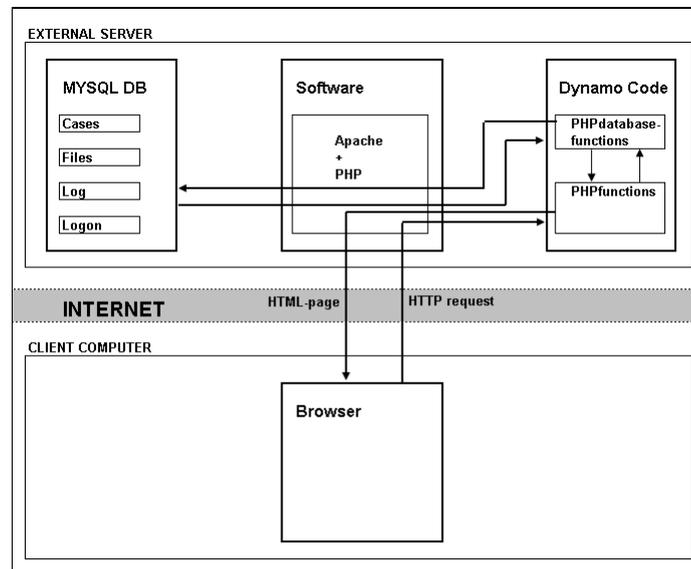

Figure 2. Schematic overview of DYNAMO's implementation. As we write, DYNAMO represents 536 projects, 7032 files, 1911 keywords (index values), 713 registered users, and 85126 logs.

The *files* database contains all information about the files documenting the projects (file name, author, source, file type, etc). Projects can be documented with a combination of various media, ranging from sketches and drawings, over digital models and animation, to pictures, video and text. The advantages of this combination are manifold. Compared to written data, visual and spatial representations better fit architects' designerly way of thinking and working (Cross 1982). Furthermore, mixing multiple modes of representation provides users with a richer learning experience, which results in a better understanding of the project.

The log base keeps a log of all user interactions with the platform. For every user (inter)action, the log base stores parameters such as the identity of the user (user name), the location of the page that registered the action, the query string used during the action (which typically contains variables such as the project and/or file id), the date and time the action was performed, the client ip address, and the type of action performed (check file, delete file, delete project, delete value, logon, recheck file, save file, search, view file, view project).

The *logon* database takes care of user privileges and administration. The current version of DYNAMO distinguishes between three types of users. DYNAMO *users* have limited privileges: they can consult the platform, feed it with new projects, project features and documentation, and create new categories. *Monitors*

have extended privileges in that they can approve, alter or delete user-added materials. *Administrators* have access to all DYNAMO features, including user and monitor administration.

2.3 Use

DYNAMO is being used and developed by various groups and individuals in architecture schools, and provides an emerging pool of knowledge for design practice. As such, the platform holds great potential to establish a learning community and narrow the existing gap between practitioners and student designers in architecture. The question whether this potential is effectively exploited, motivated a series of experiments to assess the added value of the platform in supporting architectural practice and education.

This longitudinal study has been reported more fully elsewhere (Heylighen et al. 2004) and only the findings relevant in the context of this paper are described here. To give a little background: the study monitored and analyzed various user groups (students and practitioners, both novice and expert designers) while using the platform in solving concrete design tasks. To this end, a variety of data was drawn on, including think aloud protocols, surveys and log files.

Protocol analysis revealed some interesting differences between novice and expert designers. While novices tend to scan DYNAMO for inspiration on interesting concepts, experts rather try to project their own concepts onto the design projects available. In a comparative study of two designers, the novice under consideration did not consult DYNAMO because he was looking for anything in particular, but to get inspired by concepts in the project collection. By contrast, the expert explicitly looked for concepts related to his design ideas. His search was more structured in that he tried to match his own concepts with those available in DYNAMO. This seems to suggest that novices and experts make different associations between projects: novices seem to rely, primarily, on superficial similarities between the design task at hand and the projects in DYNAMO, assuming that the latter will provide them with useful concepts for their task; experts, for their part, seem to draw more insightful parallels between projects that may look quite different at first sight.

The study also identified a number of obstacles that prevented DYNAMO from being used as originally intended. Worth mentioning here are the obstacles encountered in establishing a community of users who exchange and learn from each other's insights and experiences. In principle, we expected users to consult DY*N*AMO's collection, but also to (inter)actively participate in feeding it, be it by adding, supplementing, connecting or commenting on projects. Yet, apart from students submitting an obligatory project analysis, hardly any user so far exploited the opportunities for (inter)action. Perhaps a good term to characterize this use is 'free-ridership' (Agre 2003): users who do not participate in contributing to developing DYNAMO's content will nonetheless benefit from it. One reason for this phenomenon may relate to the submission procedure, which users seemed to find rather complex at first. In general, making material available to others takes time, effort, and sometimes specific skills (Van House 2003). Especially the articulation of appropriate metadata, which are key to identifying inter-project relationships, is far from trivial a task. Moreover, making material available seems inconsistent with the habits and priorities in architectural design practice, which tends to be highly secretive in nature. Professional architects turned out to be highly skeptic about sharing information on their own projects with other platform users, anxious to give away the secret of their success to other design firms. If practitioners are not willing to share their knowledge, insights, and expertise across firm boundaries, this may become a heavy burden on DYNAMO's future.

3. ROOM FOR IMPROVEMENT

DYNAMO, whose first implementation dates back to 1998 (Segers 1998), has outgrown the stage of prototype and is ready to enter the next phase, providing a more solid implementation of the concepts described above. As the platform is being used more extensively, questions arise about the frequency and usefulness of consultation by the different types of users. As mentioned above, various experiments have been set up to evaluate the platform's influence on design, but also the effectiveness of the tool through analysis of the actual usage. Besides general insights regarding design and design tool support, these evaluations have generated invaluable feedback to help identify new, more specialized ways to fulfill the original goal of the platform.

Originally, DYNAMO's goal was to embrace and profit from several kinds of interaction between designs and designers in different contexts and at different levels of expertise to support the actual design process. As the platform grows, this interaction becomes even more crucial if users are to see the wood for the trees. Indeed, on the level of data management, the continuously growing data pool demands proportional efforts from the DYNAMO monitors and administrators to keep the information provided of an acceptable quality and trustworthiness. Furthermore DYNAMO users are confronted with ever larger selections of projects, calling for more sophisticated search and navigation facilities. To suit the profile of a collaborative learning environment, DYNAMO should be able to make suggestions based on previous user and monitor experiences.

So far, we have mentioned, interaction remains a quite limited phenomenon, in the sense that most DYNAMO users act as 'free-riders'. In a slightly different sense, however, there is in fact a lot of interaction—between user and platform and between different projects—which currently remains implicit, yet holds great potential to improve DYNAMO's performance. Since most of the specific user (inter)actions are logged into a separate database, a lot of information on the way the tool is actually used is implicitly available. Currently, this log base is completely hidden for the users, because it has little value as such. Yet when processed by methods to extract meaningful parameters out of it, these log data may help steer DYNAMO further towards its initial goal. Data mining techniques can be used to discover currently hidden facts about projects and their contents. Using a combination of statistical analysis and specific database queries, patterns and subtle relationships between the different projects, indices and other contents can be identified. This would allow DYNAMO in the long run to learn from its experiences with user (inter)actions and to apply what it has learned by providing more relevant information to an individual user, based on what other users have consulted before.

The information extracted from the log base could be translated into simple, practical features that are easy and intuitive to understand, especially for inexperienced users. Currently, a single project in DYNAMO is accessible in three different ways. First, it can be accessed through a fixed, ordered or random list. Different orderings are available, but regardless of what other users did before, the list will always be the same under the same circumstances. Second, users may specify a keyword or an index value to find all projects that contain that value. Third, projects can be opened through the web of indices between them. Clicking on one of the index values presented in the project details causes to appear all projects with the same value for that index. The latter method differs from the second in that the user does not explicitly searches for an actual keyword but rather for similarity between the different projects.

These ways of accessing projects show some intrinsic shortcomings, which will only get worse as the volume of the collection increases. Some of the problems derive from the categorization through freeform keywords. Whereas the user can access documents through the definition of those keywords, the translation of the concept of interest into the correct word(s) is not always straightforward, especially not for inexperienced users. Furthermore defining meaningful metadata for project categorization demands great effort and special expertise, and therefore will always show room for improvement. As DYNAMO's information partly originates from free-rider users, the issue gets even more problematic.

The added value of exploiting the log data on user (inter)actions is that DYNAMO would be able to present a set of projects the user did not ask for explicitly, but only implicitly by using one of the three access methods outlined above. Since inexperienced users have no overview of all relevant information stored in the platform, such approach would offer two major advantages from the perspective of a learning experience. First, DYNAMO would be able to exploit the expertise and insights used by more experienced users while navigating through the projects, to make suggestions to users who are not familiar with the contents. This would enable newcomers to benefit from decisions that have led to successful results in the past. Second, we can think of a scenario where DYNAMO could be able to derive and propose alternative ways to look for projects, by analysing the way these projects have been accessed by a certain group of users.

4. A PROMISING APPROACH

Today, the World Wide Web offers a wide range of professional and commercial digital libraries, which utilize various techniques to improve information display through analysis of user interaction. One well known example is the use of data mining techniques by on-line stores to recommend alternative products to

customers, based on their current selection and what other customers with similar interests have bought. Most of these techniques, however, are implemented as purely practical solutions to achieve straightforward goals, and have no underlying theory to explain how users navigate through a web of linked documents. An interesting exception to this rule is the connectionist approach proposed by Francis Heylighen and Johan Bollen (2002). Not only is the approach characterized by a unique, coherent paradigm inspired by the working of the brain; this paradigm perfectly chimes with the cognitive model underlying DYNAMO.

The essence of connectionism is to model cognitive systems as networks of nodes connected by (typically weighed and directed) associations or links. Drawing on the analogy between the organization of the brain— where neurons play the role of nodes and synapses the role of links—and the dynamic organization of a document network, the paradigm assumes that documents develop stronger associations as they are more frequently co-activated or, more in general, as they functionally co-occur, c. q. appear in the same category of usage. Co-activation decreases exponentially with the time between consultations, while its strength is proportional to the user's interest in the document. The latter may be evaluated explicitly or inferred implicitly from user actions or the duration of the consultation.

The major advantage of the connectionist model is that it enables document networks to self-organize. Simply represent each document by a node, and then derive the initial strength of the link between two nodes from the degree to which they are mutually associated (Bollen 2001), by aggregating various features that the documents share (co-occurrence). This network of associations is then continuously improved, updated and fine-tuned by letting it learn from the actual way it is used by its user community.

In the light of this connectionist approach, DYNAMO can be thought of as an associative network in which the meaning of a single project is defined by the whole of its associations with other projects rather than by its independent content. In other words, projects can be considered as the nodes in the network and the strength of the links between them as the degree to which they are mutually relevant or associated. At this point, these links are determined by the projects' characteristics, i.e. by the metadata explicitly specified by users. As a result, functional co-occurrence among projects is necessarily limited to having the same value for one or more categories, such as being designed by the same architect, using the same material, being built in the same location or period, having the same function, and the like. Moreover, once these values have been specified, the degree of association between two projects stays the same, and may alter only when a user explicitly creates a new category.

Following the ideas outlined above, however, this network may be considerably enriched by exploiting the various modes of interaction DYNAMO aims to stimulate and support. A first possibility would be to take into account not only those features of a project that are explicitly specified as category values, but also features that are contained in the files documenting that project. For instance, while a user submitting the Bus station in Leuven (Belgium) designed by Manuel de Solà-Morales may forget (or fail) to label it explicitly as making use of brick, a journal article documenting this project is likely to mention this material, as it plays an instrumental role in realizing the architect's design concept. (Along the same line, one could even think of distilling numerical data out of digital models, or using image recognition techniques to derive visual features from pictures or drawings.)

In addition to features extracted from files that document projects, DYNAMO could exploit information on how users actually interact with projects, as recorded by the log base. By analyzing the records of each project consultation, the degree of inter-project association could be computed on the basis of relations such as being consulted by the same user or being consulted within the same time slot (Heylighen & Bollen 2002). Also, if any documents related to the projects are requested, one could derive the user's relative measure of interest in each project by counting the number of requested files or the total time spent consulting the project, and this interest gives a direct indication of the degree of "activation" to be accorded in a connectionist model (Heylighen & Bollen 2002, Heylighen 1999).

Added together, these various forms of association between projects form an overall, fine-grained association matrix. This can be put to use to guide DYNAMO users in various ways (Heylighen & Bollen 2002). To start with, it would be possible to append to each project a list of the projects that are most strongly associated with it, in decreasing order of association. In this way, a user who discovers a relevant project will immediately get a pointer to the projects that are most likely to be relevant as well, even though they may be labeled incompletely or differently. Another application is the clustering of projects. Submitting the association matrix to a clustering algorithm should allow DYNAMO to automatically create categories of projects, even when those categories have not been formally recognized yet. Finally, by keeping track of all the projects that a particular user has consulted, together with an estimate of the degree of interest for each of

these projects, we can define an "activation vector" representing that user's present interest profile. Multiplying this vector recurrently with the association matrix implements the connectionist process of "spreading activation", where new projects are activated proportionally to the degree that they are associated with the complete interest profile of the user, rather than just the last consulted project. In this way, DYNAMO can produce at any time a tailor-made recommendation for the user, which is updated with each further consultation activity.

These applications, if successfully implemented, would considerably help DYNAMO in tackling the obstacles it is currently facing. First of all, the extraction of relations between projects from documents and log files would free users at least in part from the cumbersome task of articulating appropriate metadata (Heylighen & Bollen 2002), while considerably enriching the content of the platform. Secondly, and more importantly, equipping DYNAMO with a mechanism for self-organization would allow benefiting from the expertise of *all* designers using the platform, including those who do not release information on their own projects. Indeed, such map would capture and exploit the insights that expert designers rely on while consulting projects without any extra effort or even awareness on their part, thereby making these insights readily available to other DYNAMO users. Each user, or rather each usage would interpret and transform the previous organization of DYNAMO's collective mental map, not in its entirety but piecemeal. Each usage would produce additional associations in the network of projects through the (often tacit) insights that users rely on to guide project selection and consultation. As pointed out above, the usage of experts seems to differ from that of novices, not because they are applying different rules or strategies, but simply because they see other, less obvious relationships across projects. Enabling DYNAMO to exchange this kind of insights effortlessly and unnoticed, i.e. without bothering the users to explicitly formulate their insights, would greatly enhance the potential of DYNAMO as a collective platform to exchange knowledge and insights in a field that is usually highly secretive.

5. SUMMARY AND FUTURE WORK

Supporting the exchange of design ideas and insights through interaction on a collective platform turns out to suffer from two major thresholds: a physical threshold caused by the effort and time needed to feed the platform, and a psychological threshold to share ideas and insights with others. In order to overcome these thresholds, this paper has proposed to consider the platform as an associative network of projects, and developed some ideas to determine its associations by extracting and exploiting the insights already available in the log data. This should allow the platform to learn from its previous experiences and progressively increase its ability to satisfactorily serve its users and eventually reach its original goal.

We are the first to admit that further evidence is needed for the value of the scenario proposed in this paper. Therefore, we are planning first of all to define a working strategy to implement the association matrix. A further step will be to analyse the results and determine whether they can be useful for an implementation visible to the user. After that, the interface of DYNAMO should be redesigned to visually integrate the results achieved through the data mining process.

In taking these future steps, there are plenty of insights from other research we can learn from, especially in the domain of recommender systems and data mining of web usage data. Collaborative filtering, for instance, produces recommendations by comparing users to other individuals in a user community and deducing possible interesting items from that comparison (Konstan 1997, Wasfi 1999). The application of collaborative filtering algorithms implies that a metric of user similarity be defined to position individual users within the context of a user community. Moreover, the principle of collaborative filtering can be inversed to generate item similarities based on usage patterns. Rather than deriving the similarity of users from the items they download, the similarity of items can be assessed on the basis of the users downloading them (Heylighen 1999). Research on collaborative filtering thus may be quite relevant to the implementation of the idea of a self-organizing DYNAMO.

Although this implementation still lies largely in the future, we are already aware of some important challenges to be dealt with. As mentioned before, most of the data needed to build the associative network are readily available in the log base. However, usage logs can be quite noisy (Pitkow 1997). For instance, when a user leaves the computer for some time, the session is interrupted and a false impression of the

244

consultation behaviour is generated. These and other types of noise will require the application of specific techniques to filter out the meaningful data patterns.

A second challenge in implementing the proposed scenario derives from the fact that the algorithm generates the best results if used frequently and in a recursive way. When a feedback loop is created the mining process will further fine-tune itself to enhance its performance. This means that the usefulness of the approach cannot be evaluated before it has more or less been integrated into the interface and tested over a significant period of time.

Moreover, a key challenge will be to determine what log data are relevant to take into account and trigger a new recursion. Indeed, tracing the 'interesting' information in the log base, i.e. insightful associations between projects as made by expert designers, in a system with 713 users is far from trivial a task. We do not want the algorithm to take into account every single usage, but only those interactions that can help enrich the content of the case base.

Awaiting the implementation of this scenario, our purpose in presenting these ideas is to point out the advantage of using analogies with cognitive processes, as both DYNAMO and the co-activation paradigm do quite literally. In our view, it is precisely these common roots in the functioning of the human brain that makes the paradigm highly compatible with DYNAMO (and vice versa) and the scenario outlined in this paper worth pursuing.

ACKNOWLEDGEMENTS

REFERENCES

Agre, P., 2003. Information and Institutional Change: The Case of Digital Libraries. *Digital Library Use. Social Practice in Design and Evaluation.* The MIT Press, Cambridge, USA, pp. 219-240.

Bollen, J., 2001. *A Cognitive Model of Adaptive Web Design and Navigation: A Shared Knowledge Perspective.* PhD Thesis, Vrije Universiteit Brussel.

Cross, N., 1982. 'Designerly ways of knowing', *Design Studies,* Vol. 3, No. 4, pp. 221-227.

Heylighen, A. and Neuckermans, H., 2000. DY*N*AMO: Dynamic Architectural Memory On-line. *In Educational Technology and Society*, Vol. 3, No. 2, pp. 86-95.

Heylighen, A. et al. 2004, ICT revisited—from information & communication to integrating curricula? *ITcon*, Vol. 9, Special Issue Digital Media Libraries, pp. 101-120.

Heylighen, F., 1999. 'Collective Intelligence and its Implementation on the Web: algorithms to develop a collective mental map', *Computational and Mathematical Theory of Organizations* 5(3), p. 253-280.

Heylighen, F. and Bollen, J., 2002. Hebbian Algorithms for a Digital Library Recommendation System. *Proceedings 2002 International Conference on Parallel Processing Workshops* (IEEE Computer Society Press).

Kolodner, J.L., 1993. *Case-Based Reasoning.* Morgan Kaufman, San Mateo, USA.

Konstan, J. A., Miller, B. N., Maltz, D., Herlocker, J. L., Gordon, L. R. and Riedl, J., 1997. 'Grouplens: applying collaborative filtering to usenet news', *Communications of the ACM*, 40(3), pp. 77-87.

Pitkow, J., 1997. 'In search of reliable usage data on the www', *Proceedings of the Sixth International WWW Conference*, Santa Clara, California, April 7-11, 1997.

Riesbeck, C.K. and Schank, R.C., 1989. *Inside Case-Based Reasoning.* Lawrence Erlbaum, Hillsdale, USA.

Schank, R.C., 1982. *Dynamic Memory—A Theory of Reminding and Learning in Computers and People.* Cambridge University Press, Cambridge, USA.

Schön, D., 1983. *The reflective practitioner: how professionals think in action.* Basic Books, New York, USA.

Segers, R., 1998. *Prototype van een interactieve case-base voor architectuur, geïllustreerd met woningen op helling.* Master's thesis, K.U.Leuven, Dept. ASRO, Leuven, Belgium

Van House, N. A., 2003. Digital Libraries and Collaborative Knowledge Instruction. *In …Digital Library Use. SocialPractice in Design and Evaluation.* The MIT Press, Cambridge, USA, pp. 271-296.

Wasfi, A. M. A., 1999. 'Collecting user access patterns for building user profiles and collaborative filtering', *Proceedings of the 4th international conference on Intelligent user interfaces*, ACM Press, pp. 57-64.