

# Distributed (design) knowledge exchange

ANN HEYLIGHEN<sup>1</sup>, FRANCIS HEYLIGHEN<sup>2</sup>, JOHAN BOLLEN<sup>3</sup> & MATHIAS CASAER<sup>1</sup>

<sup>1</sup>*Katholieke Universiteit Leuven, Dept. ASRO, Kasteelpark Arenberg 1, B-3001 Leuven, Belgium,*

<sup>2</sup>*Vrije Universiteit Brussel, Center "Leo Apostel", Pleinlaan 2, B-1050 Brussels, Belgium*

<sup>3</sup>*Old Dominion University, Computer Science Dept., ...*

Corresponding author: ann.heylighen@asro.kuleuven.be t: +32 16 321741 f:+32 16 321984

Acknowledgements: Ann Heylighen is a Postdoctoral Research Fellow of the Fund for Scientific Research-Flanders (FWO).

**Abstract.** The domain of social intelligence design attempts to integrate three until recently separate realms—mind, society, and matter—yet lacks a coherent theoretical framework to do so. Despite the intrinsic complexity of integrating individual, social and technologically-supported intelligence, the paper proposes a relatively simple ‘connectionist’ framework for conceptualizing distributed cognitive systems. Shared information sources (documents) are represented as nodes connected by links of variable strength, which increases as the documents co-occur in the usage patterns. This learning procedure captures and exploits its users’ implicit knowledge to help them find relevant information, thus supporting an unconscious form of exchange. The principles are applied to a concrete problem domain: architects sharing design knowledge through a database of associatively connected building projects.

*Keywords: architectural design, connectionism, distributed cognition*

## 1 A general perspective on social intelligence design

Social intelligence design (SID) is one of an emerging cluster of approaches, originating in different disciplines and application domains, that attempt to bring together three until recently separate realms: *mind* (cognition, intelligence), *society* (social interaction, organizations, institutions), and *matter* (objects, tools, technologies). Related approaches are known as distributed cognition (Hutchins 1995), situated and embodied cognition (Susi & Ziemke 2003), activity theory (Hasan et al. 1998), extended mind (Clark 1995), and collective intelligence (Levy 1997, Heylighen 1999). All of them extend the cognitive processes that characterize learning, memory and intelligence outward: from the individual brain to the surrounding social and physical environment.

One dimension of extension is the collectivity or social system formed by individual minds interacting. It is a commonplace that most of our knowledge develops in interaction with

others: an individual cut-off from society (e.g. a child raised by wolves) would be virtually powerless to comprehend our complex environment—if capable of surviving at all..

More recently, the awareness has been growing that cognitive processes heavily rely on structures in the physical environment, i.e. on various forms of symbols, tools, and spatial arrangements (Kirsh & Maglio 1994). These structures, exemplified by items such as scribbled notes, traffic signs, file cabinets, and computers, are used to store, process, communicate or organize information, thus relieving our limited mental capacities. Since most of these structures have been consciously designed for that purpose, we may call them *information technologies*.

However, this should not mislead us to immediately think in terms of digital computers or electronic networks. When shepherds put pebbles in a bag to count their sheep or drop them along the path to remember where they passed, they are using a cognitive artefact just as much as if they were calculating with a palmtop or determining their location with GPS. In both cases, they manipulate their physical environment so that it can hold information, thus augmenting their interior capacity. Similarly, architects who design a building so that exits or bathrooms are easy to find create an information structure intended to minimize the cognitive load on its users. For an even more vivid example of the need for external structure to support cognition, imagine a household in which all individual items (cooking utensils, clothes, cleaning products, cans with food, books, pens, etc.) are distributed completely at random over the different cupboards, drawers and shelves: since our brain would never be able to memorize all their locations, these items would effectively become irretrievable and thus useless.

The social and material dimensions of extended cognition are inextricably linked. Exchanging information with others necessarily entails a shared physical medium. In the simplest case, this is just the air that carries the sounds of our speech. But as soon as the information becomes too complex to be easily remembered, we begin using more structured media, such as paper, computer discs or spatial arrangements, which can accurately register and shape the information being shared so as to facilitate reuse. In the ‘random distribution’ example above, even if after years of diligent training we had come to memorize the locations of all household items, we would never be able to communicate this knowledge to another person. A ‘common sense arrangement’ (e.g. clothes in the wardrobe, cleaning products under the kitchen sink), on the other hand, merely needs a few general indications to be used by others.

This brings us to the perspective of *distributed cognition*: to understand complex information processes, we must consider the distributed organization constituted by different individuals with different forms of knowledge and experience, the social network that links them together, and the artefacts, media or information technologies that support their individual thought and interindividual communication. A classic example of this approach is the work of Hutchins (1995), who used ethnographic methods to study a Navy warship as a distributed cognitive system.

The distributed cognition approach, and by extension the domain of SID, as yet offers little more than a heterogeneous collection of ideas, observations, and case studies. It lacks a solid foundation for building models of concrete systems and processes (Heylighen et al. 2004b, Susi & Ziemke 2001). Despite the intrinsic complexity of the problem of integrating individual, social and artefact-supported intelligence, we believe that a few well-chosen

concepts and methods can help make the issue much more tractable, offering a relatively simple paradigm to conceptualize a distributed cognitive system. The remainder of this paper first briefly sketches our ‘connectionist’ conceptual framework, and then describes in more detail its application to a concrete problem domain: the dynamic sharing of design knowledge among multiple architects by means of a database of associatively connected building projects. The paper closes by situating our approach in the context of related research and outlining directions for future work.

## **2 A connectionist perspective on distributed cognition**

In essence connectionism models cognitive systems as networks of nodes connected by (typically weighted and directed) associations or links. This static representation is complemented by dynamic rules that govern the short-term interaction between the nodes, and its long-term effect on the links. This perspective is inspired by the organization of the brain, where neurons act as nodes, synapses as links, and interaction occurs by the transmission of electrical activation from neuron to neuron, with a strength proportional to the ‘conductivity’ of the connecting synapse. This conductivity is adjusted according to the principle that successfully used links grow stronger, while unsuccessfully used ones weaken, thus continuously improving the network’s overall efficiency. This can be formalized by the Hebbian rule, stating that the increase in strength of a link is proportional to the strength of the co-activation (product of activations) of the nodes it connects (Heylighen & Bollen 2002). Each experience of use thus leaves a trace in the linking pattern. The network therefore functions as a self-organizing, dynamic memory, which becomes an increasingly reliable guide for dealing successfully with situations similar to those experienced in the past.

While this approach grossly simplifies the actual processes in the brain, research on artificial neural networks has shown that it allows to successfully model most fundamental cognitive processes (McLeod et al. 1998). The advantage of the relative abstractness and simplicity of this connectionist representation is that we can easily extend it to systems that do not consist of actual brain tissue. In particular, we propose to extend it to distributed cognitive systems that consist of human individuals and physical information sources (‘documents’), linked by communication channels.

For the sake of simplicity, this paper focuses on the mechanisms of distributed memory, learning and information retrieval, assuming that more complex processing either takes place in individuals or emerges from the collective dynamics of the network (Heylighen et al. 2004b). We further assume that the documents act as a shared memory, through which individuals can exchange information, access, and contribute to, the whole of their collective knowledge (Heylighen 1999). The network of linked documents thus constitutes a medium for indirect cognitive collaboration between its users. (Note that the connectionist perspective can also fruitfully model the direct communication between individuals (Van Overwalle & Heylighen, 2006).

Given the present information explosion, documents can easily number in the millions. As the household example illustrates, making this memory usable requires an intuitive ordering that makes sense to all users. While an individual might easily memorize that the

cleaning products are stored in the wardrobe, and the trousers under the kitchen sink, most people would look for the trousers in the wardrobe, and vice-versa. But for an ever growing collection of novel and abstract data, there are no such conventional location schemes. We need to develop a *collective mental map* (Heylighen 1999), that indicates the relative locations of all documents, and constantly adapts to changes in information content or user preferences.

The basic advantage of the connectionist model is that it allows having such a mental map self-organize. Simply represent each document by a node, and derive the strength of the link between two nodes from the degree to which they are mutually relevant or associated. To retrieve a document, start from the already known document that is most associated with it, and then move ‘upstream’, each time selecting the link that seems most strongly associated with the target. This is how people navigate through hypertext networks (Bollen 2001). The effectiveness of the strategy depends on how well the actual hyperlinks reflect the users’ intuitive associations. And this depends on how well the collective mental map represented by the linking patterns captures the collective knowledge of its users. The Hebbian learning rule allows the network to assimilate the knowledge implicit in how the network is being used (Heylighen & Bollen 2002). We simply assume that two documents A and B are *co-activated* if they are consulted by the same user within a relatively short time interval. This indicates that B is likely to be relevant for A, and that it is worth strengthening the link  $A \rightarrow B$ . The next time a user consults A, B will be easier to reach, thus facilitating the search for relevant documents. In that way, the network continuously adapts to better reflect the users’ intuitive expectations collectively.

Co-activation, in the sense of being consulted (nearly) simultaneously, can be generalized to *functional co-occurrence*: appearing in the same category of usage, e.g. being cited in the same paper, bought by the same customer, or present in the same library section, having the same author or the same keyword in the title, sharing specific features such as document type or subject... Some of these data are already being used to find related documents, e.g. co-citation to locate scientific papers on the same subject, or co-purchase to recommend relevant books to customers of web bookshops. In our generalized connectionist model, all these co-occurrence data are potentially useful to indicate associations between documents, i.e. to efficiently structure the collective mental map. Combining all data available should improve the quality of the associations, especially in the beginning stages when the network has not yet had time to ‘learn’ from the way it is being used.

The only problem is to determine the relative importance of the different contributions: is having the same author a better or worse indication of mutual relevance than being bought by the same customer, or having the same keyword? A possible solution would be to consider the degree of association implied by each type of co-occurrence as a variable, and then to calculate the correlation coefficients between all variables over the whole collection of documents. Applying the statistical technique of factor analysis to the matrix of cross-correlations, we can then compute a ‘principal component’, a linear combination of all variables that maximally co-varies with each of them. The load of a variable on this ‘principal’ dimension can then be interpreted as an indication of its importance for the overall calculation of associative strength. Co-occurrence variables that correlate poorly with the others have low loads, implying that they are unreliable indicators of relevance.

In conclusion, our connectionist approach is able to intuitively structure any—even random—collection of documents, first by aggregating the various features that two documents share (co-occurrence) to determine their initial degree of association. 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 community of users. This makes it much easier for the users to find the documents they are interested in and globally orders the collection, which allows various more complex forms of use e.g. categorizing documents by means of cluster analysis, finding the most relevant documents with particular features, producing tailor-made, context-dependent recommendations for individual users, or discovering global trends (Heylighen & Bollen 2002). The method relies on an implicit, technology-mediated collaboration between the different users, which does not demand any effort from them except the one they have to invest anyway to use the network purely for their individual sake.

### **3 A concrete application**

Having sketched our connectionist framework to conceptualize distributed cognitive systems, this section switches attention to a concrete application in architecture. In this domain, design ideas are developed as much through interaction as by individuals in isolation. This observation inspired the development of a Dynamic Architectural Memory On-line (DYNAMO), an interactive platform to share ideas, knowledge and insights in the form of concrete building projects among a multitude of architects.

Interaction with professional and student architects revealed DYNAMO to suffer from two thresholds: 1) making projects available to other platform users takes time, effort, and specific skills; 2) architects tend to sense a psychological threshold to share their ideas with others. This section shows how the framework outlined above, and in particular the notion of co-occurrence can help addressing both thresholds. Equipping DYNAMO with a self-organizing mental map reflecting the relative locations of all projects, and constantly adapting to changes in content or user preferences, would enable the platform to benefit from its users' insights without any extra effort or even awareness on their part.

The subsections which follow present the basic concepts of DYNAMO, followed by their implementation and results from experiments with various user groups. Based on this foundation, we point out the power of the connectionist framework to help the platform effectively fulfill its original ambition.

#### **3.1 DYNAMO in a nutshell**

DYNAMO (<http://dynamo.asro.kuleuven.be>) is a web-based design assistant for student and professional architects. It aims to incorporate quite literally the view of cognition underlying Case-Based Design (CBD) and at the same time to extrapolate it beyond the individual (Heylighen & Neuckermans 2000).

Rooted in the Theory of Dynamic Memory (Schank 1982), the CBD approach propounds that people's knowledge does not only entail abstract, generally applicable principles, but also 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 people's reminders 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, and 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, DYNAMO dynamically develops and is dynamically developed by architects' design knowledge by stimulating and supporting several modes of interaction:

- interaction among building projects, for projects are labeled and linked to related projects by various features architects address during design;
- interaction between (human) designer and (computer) memory, for users cannot only consult projects in DYNAMO, they can improve its content in various ways;
- 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;
- and thus also interaction between practice and education in architectural.

### **3.2 Implementation**

Technically, DYNAMO can be thought of as a learning content management system. It is designed to support the creation, storage, use and reuse of learning content in the granular form of building projects. All learning content is organized by a dynamic metadata classification system and stored in a data repository embodied by a relational database, subdivided into four sub-databases:

The *cases* database labels each building project by 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 features. 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, etc. For some categories, materials for instance, a single project can have multiple values.

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, e.g. sketches, drawings, pictures, digital models, animation, video and text.

For every user interaction with the platform, the *log* base stores parameters such as the user's identity (user name), the location of the page that registered the action, the query string used during the action (typically containing variables the the project or file id), the date and time the action was performed, the client ip address, and the type of action performed (e.g. save file, view project, search).

The *logon* database takes care of user privileges and administration. DYNAMO distinguishes between three types of users. *Users* 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.

As we write, DYNAMO represents 552 projects, 7439 files, 1911 keywords (index values), 1010 registered users, and 85126 logs.

### **3.3 Use**

DYNAMO is being used by various groups and individuals in architecture schools, and provides an emerging knowledge pool for design practice. As such, it holds great potential to become a distributed cognitive system for student and professional architects. To what extent this potential is effectively exploited was evaluated through a series of experiments in architectural practice and education.

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

Protocol analysis revealed 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 projects in the platform. In a comparative analysis, the novice under consideration consulted DYNAMO not because he was looking for anything in particular, but to get inspired by concepts in other projects. 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 in the project collection. This suggests that novice and expert designers make different associations between projects: novices tend to rely, primarily, on superficial similarities (e.g. the function) between their design task and DYNAMO's projects, hoping that the latter provide useful concepts for this task; experts

seem to draw more insightful parallels (e.g. the underlying concept or spatial organization) between projects that may look quite different at first sight.

The study also revealed several obstacles that hamper DYNAMO in establishing a community of users who exchange insights and experiences. In principle, we expected users to consult the project collection, but also to (inter)actively participate in feeding it, by adding, supplementing, or connecting projects. Yet, apart from students submitting an obligatory project analysis, hardly any user so far exploited the opportunities for (inter)action. This pattern of use is called ‘free-ridership’ (Agre 2003): users who do not participate in contributing to developing DYNAMO’s content nonetheless benefit from it.

One reason for this ‘free-ridership’ may relate to the submission procedure, which users seemed to find rather complex at first. 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 practice. Professional architects turned out to be highly skeptic about sharing information on their projects with other users, anxious to give away ‘the secret of their success’. If practitioners are not willing to share their knowledge and insights, this may become a heavy burden on DYNAMO’s future.

### **3.4 Towards self-organization**

In the light of the connectionist model, however, this ‘free-ridership’ does not need to be as problematic as first meets the eye. DYNAMO can be thought of as a collective mental map 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 nodes in an associative network and the strength of the links between them as the degree to which they are mutually relevant or associated.

Originally, these links are determined by the projects’ features, i.e. the metadata explicitly specified by users. As a result, functional co-occurrence among projects is limited to having the same value for one or more categories, e.g. being designed by the same architect, using the same material, being built in the same location or period, having the same function etc. Moreover, once these values have been specified, the degree of association between two projects stays the same, and may alter only when a new category is created.

This collective mental map may be considerably enriched by exploiting the various modes of interaction DYNAMO stimulates and supports. So far, interaction remains a quite limited phenomenon, in the sense that most users of the platform act as ‘free-riders’. In a different sense, however, there is in fact a lot of interaction—across different projects and between user and platform—that currently remains implicit, yet holds great potential to determine the degree of association between projects in a more sophisticated and dynamic way.

One possibility would be to determine this degree based not only on project features that are explicitly specified as category values, but also on features contained in the project documentation. For instance, while the user submitting the project ‘s Hertogenmolens in Aarschot (Belgium) by Noa Architects may forget (or fail) to label it explicitly as making use of corten steel, a journal article documenting this project is likely to mention this material, as it plays an instrumental role in realizing the architects’ design concept.



In addition, DYNAMO could exploit information on how users interact with projects, as each user (inter)action is recorded by the log base. By analyzing the records of each project consultation, the degree of inter-project association could be computed based on relations like being consulted by the same user or within the same time slot (Heylighen & Bollen 2002). Also, one could derive the user's relative measure of interest in each project from the number of requested files or total time spent on the project, and this interest reflects the degree of 'activation' to be accorded in a connectionist model (Heylighen & Bollen 2002, Heylighen 1999).

Added together, these various forms of inter-project association determine an overall, fine-grained association matrix, which can guide DYNAMO users in various ways (Heylighen & Bollen 2002). One could append to each project a list of 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 be pointed to the projects that are most likely to be relevant as well, but may be labeled incompletely or differently. Another application is project clustering. Submitting the association matrix to a clustering algorithm allows DYNAMO to automatically create categories of projects, even when those categories have not been formally recognized yet. Finally, by keeping track of all projects a particular user has consulted, together with an estimate of the degree of interest for each, 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', whereby new projects are activated proportionally to the degree that they are associated with the user's interest profile, rather than just the last consulted project. In this way, DYNAMO can give users at any time a tailor-made recommendation, which is updated with each further consultation activity.

These applications would considerably help DYNAMO in tackling the obstacles it is currently facing. First, extracting inter-project relations from documents and log files would free users at least in part from articulating appropriate metadata (Heylighen & Bollen 2002), while considerably enriching the platform's content. As DYNAMO grows, this becomes even more important if users are to see the wood for the trees. The continuously growing data pool demands proportional efforts from monitors and administrators to keep the information provided of an acceptable quality and trustworthiness. Furthermore, users are confronted with ever larger project selections, calling for more sophisticated search and navigation facilities.

Secondly, equipping DYNAMO with a self-organizing mental map would allow benefiting from the expertise of *all* architects who use the platform, including those who do not release information on their projects. 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 users. Each user, or rather each usage would interpret and transform the organization of DYNAMO's collective mental map, not in its entirety but piecemeal. Each usage would produce additional associations in the project network through the (often tacit) insights that users rely on in project selection and consultation. As pointed out above, the usage of experts differs from that of novices, not because they apply different rules or strategies, but 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 them, would be a considerable step towards establishing a distributed cognitive system for architecture.

## 4 Related work

In making this step a reality, we can draw on plenty of insights from other research. In particular, the production of self-organizing mental maps strongly relates to techniques applied in recommender systems and data mining of web usage data.

Research on recommender systems can be split into two categories according to how the relationship between individual users and their context in a user community is regarded: personalization and collaborative filtering. The former category aims to provide users with interface customizations related to individual preferences. It tends to stress the needs and preferences of individual users, and to isolate them from top-down design verdicts imposed by system engineers and other users, even when WWW interfaces are involved (Rucker & Polanco 1997). Its bias is against communal standards and for individual needs. Personalization systems have been designed to make use of log and usage data (Takano & Winograd 1998, Harvey et al. 1998, Mobasher et al. 2001).

Collaborative filtering on the other hand aims to produce recommendations by comparing users to other individuals in a user community and deducing possibly interesting items from that (Wasfi 1999, Konstan 1997). Since users who are similar may like similar items, items that one user preferred are recommended to other, similar users (Herlocker et al. 1999). Applying collaborative filtering algorithms therefore implies defining a metric of user similarities to position individual users within a contextual community of other users. Individual preferences are addressed on the backdrop of a communal, semantic model (Mock & Vemuri 1997), simulating the social processes involved in retrieval (Harvey et al. 1998), such as word-of-mouth (Shardanand & Maes 1995), human recommendation (Twidale et al. 1997) or distributed problem solving (Bouthors & Dedieu 1999). The literature on collaborative filtering is thus quite relevant to the mentioned principles of self-organizing mental maps and their application.

In addition, the principle of collaborative filtering can be transposed to generate item similarities based on usage patterns. Rather than the similarity of users being derived from the items they downloaded, the similarity of items can be derived from the users who downloaded them (Sarwar et al. 2001, Heylighen 1999). This procedure resembles the principle of log analysis: two items are deemed similar if frequently downloaded by the same sets of users. When sufficient usage data are missing to generate dense networks of item relationships, this can be alleviated by associative retrieval techniques (Huang et al., 2004).

This procedure, however, does not address the issue of temporal order. The sequence in which items are accessed or co-occur may be an equally important indicator of relatedness as simple co-occurrence in a user session. Widespread efforts have been made to generate models of the user click stream (Xiao & Dunham 2001) and to apply them to recommender systems (Bollen 2000, Bollen & Rocha 2000), including the use of Hidden Markov Models (Levene & Loizou 1999) and data mining to discover association and sequential rules (Mobasher et al. 2000, Gery & Haddad 2003). Such models can help generate item recommendations based on users' past navigation history or dynamically restructure web sites' link patterns (Yan et al. 1996, Maseglia et al. 1999, Albanese 2004).

A key issue in most analysis of log data concerns data validity: usage logs can be quite noisy (Pitkow 1997). Robots, proxies, caching, and user temperament can all greatly complicate identifying session boundaries. Bollen and Nelson (2002) identify a method to generate self-organizing mental maps by having autonomous information objects monitor their local usage and automatically update their links to other objects, which may be appropriate to turning DYNAMO into a distributed cognitive system.

## 5 Summary and future directions

In response to the lack of a coherent theoretical framework for the distributed cognition approach, and by extension for the domain of SID, we have proposed a relatively simple ‘connectionist’ paradigm that provides a solid foundation for conceiving and building models of concrete distributed cognitive systems, and the processes through which they self-organize and adapt. Subsequently, we have zoomed in on the application of this paradigm to a concrete system: a multimedia platform with associatively connected building projects that aims at the dynamic sharing and exchange of design knowledge by a multitude of architects.

This exchange suffers from at least two thresholds: a physical threshold to feed the platform and a psychological threshold to share ideas and insights with other designers. Both thresholds could be largely overcome by conceiving the platform as an associative network of projects, and exploiting the information and insights that are implicitly available in the project documentation and the data on user (inter)actions to determine and continually update inter-project relationships. This should allow the platform to learn from its previous experiences and progressively increase its ability to satisfactorily support its users.

Further evidence is needed to prove the value of the scenario proposed in this paper. A first step is to define a working strategy to calculate the association matrix on the basis of usage logs and co-occurrence data. Subsequently, we will need to analyze the resulting associations and test whether they are useful in guiding and serving the user. If so, DYNAMO’s interface must be redesigned to visually integrate the degree of associations between projects.

One challenge in implementing the proposed scenario derives from the fact that the algorithm generates the best results if used frequently and recursively. 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 whether all log data are worth 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 necessarily want the algorithm to take into account every single usage, but only those interactions that can help enrich the content of the platform. On the other hand, we can expect that poor-quality data, representing users browsing randomly or not knowing very well what they are doing, will be averaged out by taking all data together, which would leave only the statistically significant trends.

Awaiting the implementation of this scenario, our purpose in presenting the connectionist framework and its envisioned application to a concrete system, is to point out the potential of using a few relatively simple yet well-chosen analogies with dynamic and distributed cognitive processes to better understand social intelligence design, as both the connectionist paradigm and DYNAMO 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.

## References

- Agre P (2003) Information and Institutional Change: The Case of Digital Libraries. In *Digital Library Use. Social Practice in Design and Evaluation*, MIT Press, Cambridge (Mass), 219-240
- Albanese M (2004) Web personalization based on static information and dynamic user behavior, In *Proceedings of the 6<sup>th</sup> ACM international workshop on Web Information and Data Management (WIDM04)*, ACM Press
- Bollen J (2000) Group user models for personalized hyperlink recommendation. In *LNCS 1892-International Conference on Adaptive Hypermedia and Adaptive Web-based Systems (AH2000)*. Springer Verlag, 39-50
- Bollen J (2001) *A Cognitive Model of Adaptive Web Design and Navigation: A Shared Knowledge Perspective*. PhD Thesis, Vrije Universiteit Brussel
- Bollen J & Nelson ML (2002) Adaptive networks of smart objects. In Bollen & Nelson (eds.) *Proceedings of Workshop on Distributed Computing Architectures for Digital Libraries (ICPP2002)*. IEEE, 487-496
- Bollen J & Rocha LM (2000) An adaptive systems approach to the implementation and evaluation of digital library recommendation systems, In *LNCS – 4<sup>th</sup> European Conference on Research and Advanced Technology for Digital Libraries (ECDL2000)*, Springer Verlag
- Bouthors V & Dedieu O (1999) Pharos, a collaborative infrastructure for web knowledge sharing. In Abiteboul & Vercoustre (eds.) *Research and Advanced Technology for Digital Libraries, 3<sup>rd</sup> European Conference, ECDL '99 Proceedings*, vol 1696 of *Lecture Notes in Computer Science*, Springer-Verlag, 215-233
- Clark A (1997) *Being There: putting brain, body, and world together again*. MIT Press, Cambridge (Mass).
- Gery M & Haddad H (2003) Evaluation of web usage mining approaches for user's next request prediction. In *Proceedings of the 5<sup>th</sup> ACM international workshop on Web information and data management*. ACM Press, 74-81
- Harvey CF, Smith P & Lund P (1998) Providing a networked future for interpersonal information retrieval: Infovine and user modeling, *Interacting with Computers* **10**: 195-212
- Hasan H, Gould E & Hyland P (1998) *Information systems and activity theory: tools in context*, University of Wollongong Press, Wollongong.
- Herlocker JL, Konstan JA, Borchers A & Riedl J (1999) An algorithmic framework for performing collaborative filtering, *Proceedings of the 22<sup>nd</sup> Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, 230-237
- Heylighen A & Neuckermans H (2000) DYNAMO: Dynamic Architectural Memory On-line, *Educational Technology and Society* **3**(2): 86-95
- Heylighen A, Neuckermans H & Casaer M (2004a) ICT revisited—from information & communication to integrating curricula?, *ITcon*, **9**: 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): 253-280
- Heylighen F & Bollen J (2002) Hebbian Algorithms for a Digital Library Recommendation System. In *Proceedings 2002 International Conference on Parallel Processing Workshops*, IEEE Computer Society Press.
- Heylighen F, Heath M & Van Overwalle F (2004b) The Emergence of Distributed Cognition: a conceptual framework. In *Proceedings of Collective Intentionality IV*

- Huang Z, Chen H & Zeng D (2004) Applying associative retrieval techniques to alleviate the sparsity problem in collaborative filtering, *ACM Trans. Inf. Syst.* **22**(1): 116-142
- Hutchins E: 1995, *Cognition in the Wild*, MIT Press
- Kirsch D & Maglio P (1994) On distinguishing epistemic from pragmatic action, *Cognitive Science* **18**: 513-549
- Kolodner JL (1993) *Case-Based Reasoning*, Morgan Kaufman, San Mateo
- Konstan JA, Miller BN, Maltz D, Herlocker JL, Gordon LR & Riedl J (1997) GroupLens: applying collaborative filtering to usenet news, *Communications of the ACM* **40**(3): 77-87
- Levene M & Loizou G (1999) A probabilistic approach to navigation in hypertext, *Information Sciences* **114**: 165-186
- Lévy P (1997) *Collective Intelligence*, Plenum
- Masseglia F, Poncet P & Teisseire M (1999) Using data mining techniques on web access logs to dynamically improve hypertext structure, *SIGWEB Newsletter* **8**(3): 13-19
- McLeod P, Plunkett K & Rolls ET (1998) *Introduction to connectionist modeling of cognitive processes*, Oxford University Press, Oxford (UK)
- Mobasher B, Cooley R & Srivastava J (2000) Automatic personalization based on web usage mining, *Communication of the ACM* **43**(8): 142-151
- Mobasher B, Dai H, Luo T & Nakagawa N (2001) Effective personalization based on association rule discovery from web usage data. In *Proceedings of the 3<sup>rd</sup> international workshop on Web information and data management*, ACM Press, 9-15
- Mock KJ & Vemuri VR (1997) Information filtering via hillclimbing, wordnet, and index patterns, *Information Processing and Management* **33**(5): 633-644
- Pitkow J (1997) In search of reliable usage data on the www. In *Proceedings of the 6<sup>th</sup> International WWW Conference*
- Riesbeck CK & Schank RC (1989.) *Inside Case-Based Reasoning*, Lawrence Erlbaum, Hillsdale
- Rucker J & Polanco MJ (1997) Personalized navigation for the web, *Communications of the ACM* **40**: 73-75
- Sarwar B, Karypis G, Konstan J & Reidl J (2001) Itembased collaborative filtering recommendation algorithms. In *Proceedings of the 10<sup>th</sup> international conference on World Wide Web*, ACM Press, 285-295
- Schank RC (1982) *Dynamic Memory—A Theory of Reminding and Learning in Computers and People*, Cambridge University Press, Cambridge (Mass)
- Susi T & Ziemke T (2001) Social Cognition, Artefacts, and Stigmergy: A Comparative Analysis of Theoretical Frameworks for the Understanding of Artefact-mediated Collaborative Activity, *Cognitive Systems Research* **2**(4): 273-290
- Takano H & Winograd T (1998) Dynamic bookmarks. Managing personal navigation space by analysis of link structure and user behavior. In *Hypertext 98*, ACM, 297-298
- Twidale MB, Nichols DM & Paice CD (1997) Browsing is a collaborative process, *Information Processing and Management* **33**(6): 761-783
- Van House NA (2003) Digital Libraries and Collaborative Knowledge Instruction. In *Digital Library Use. Social Practice in Design and Evaluation*, MIT Press, Cambridge (Mass), 271-296
- Van Overwalle F. & Heylighen F. (2006): "[Talking Nets](#): A Multi-Agent Connectionist Approach to Communication and Trust between Individuals", *Psychological Review* (in press)
- Prabhu (eds.), *Handbook of Human-Computer Interaction*, Elsevier Science, Amsterdam, 878-914
- Wasfi AMA (1999) Collecting user access patterns for building user profiles and collaborative filtering. In *Proceedings of the 4<sup>th</sup> international conference on Intelligent user interfaces*, ACM Press, 57-64
- Xiao YQ & Dunham MH (2001) Efficient mining of traversal patterns, *DATA AND KNOWLEDGE ENGINEERING* **39**(2): 191-214.
- Yan TW, Jacobsen M, Garcia-Molina H & Umeshwar Dayal U (1996) From user access patterns to dynamic hypertext linking, *Computer Networks and ISDN Systems* **28**: 1007-1014